

Hewlett Packard
Enterprise

# SAN Design Reference Guide

**Abstract**

This reference document provides information about HPE 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 HPE SAN.

Part Number: P10035-001
Published: July 2018
Edition: 96

https://t.me/learningnets

© Copyright 2001, 2018 Hewlett Packard Enterprise Development LP

## Notices

The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise 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. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Confidential computer software. Valid license from Hewlett Packard Enterprise 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.

Links to third-party websites take you outside the Hewlett Packard Enterprise website. Hewlett Packard Enterprise has no control over and is not responsible for information outside the Hewlett Packard Enterprise website.

## Acknowledgments

Itanium® is a trademark of Intel Corporation in the U.S. and other countries.

Microsoft® and Windows® are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

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

UNIX® is a registered trademark of The Open Group.

https://t.me/learningnets

# Contents


**New and changed information in this edition..... 14**

**Architecture..... 15**

**SAN design overview..... 16**

SAN solutions..... 16
HPE SAN implementations..... 17
SAN components..... 17
Fibre Channel technology..... 18
Storage area networks..... 18
SAN infrastructure..... 19
Fabrics..... 19
SAN scaling..... 19
Fibre Channel switches..... 20
SAN design approaches..... 20
SAN design considerations..... 21

**SAN fabric topologies..... 24**

Fabric topologies..... 24
Routed SAN fabrics..... 25
FCoE SAN fabrics..... 25
Benefits..... 25
Single-switch fabric..... 25
Switch models..... 26
Benefits..... 26
Cascaded fabric..... 26
Switch models..... 27
Benefits..... 27
Meshed fabric..... 28
Switch models..... 29
Benefits..... 29
Ring fabric..... 29
Switch models..... 31
Benefits..... 31
Core-edge fabric..... 31
Core-edge fabric types..... 32
Switch models..... 34
Benefits..... 34
Topology data access..... 34
Topology maximums..... 35
B-series switches..... 36
C-series switches..... 36
H-series switches..... 36
Routed fabric topologies..... 37
B-series Meta SAN..... 37
B-series Virtual Fabrics with IFR..... 38




https://t.me/learningnets

C-series VSANs with IVR....39
H-series switches with TR....39
FCoE fabric topologies....40
Data availability....40
Levels....41
Considerations....43
Topology migration....44
Nondisruptive migration....44
Migrating a cascaded fabric SAN....44
Migrating a meshed fabric SAN....45
Migrating a ring fabric SAN....45

**Fibre Channel routing....46**
Fibre Channel routing overview....46
Fabric, Virtual Fabric, and VSAN independence....47
Fabric services....47
Worldwide Name....47
Import and export....47
Routing table....48
SAN scaling and routing....48
Switch scaling....48
Scaling by routing....49
Fibre Channel routing implementations....49
Fibre Channel routing techniques....50
B-series fabric groups....51
B-series fabric partitioning using Virtual Fabrics....51
B-series Virtual Fabrics architecture....51
C-series fabric partitioning....52
H-series switch fabric routing....52
B-series, C-series, and H-series routing differences....52
Fabric redundancy and routing....55
Supported routing configurations....57
Routing and core-edge fabrics....57
Routing through an IP network....59
High-availability router configurations....59
H-series switches with TR configurations....61
Routing use cases....61

**Fibre Channel over Ethernet....65**
HPE FCoE solutions overview....65
HPE FCoE converged switch technology....67
Converged network switch ports....67
CN Switch Interoperability....70
HPE FCoE products....70
Converged network switches and blades....71
HPE FlexFabric 5900CP Switch....71
HPE FlexFabric 5930 32QSFP+ Switch....72
HPE FlexFabric 5930 4-slot Modular Switch and HPE FlexFabric 5930 2QSFP+ 2-slot Modular Switches....73
HPE FlexFabric 5950 4-slot Modular Switch and HPE FlexFabric 5950 48SFP28 8QSFP28 Switch....74
HPE FlexFabric 5940 4-slot Modular Switch and HPE FlexFabric 5940 2-slot Modular Switch....75
HPE 5820 Ethernet/FCoE Switch....76



https://t.me/learningnets

HPE FlexFabric 5700 Switch....77
HPE StorageWorks DC SAN Director Switch 10/24 FCoE Blade....78
HPE StoreFabric SN8600B 64/32 Short Wave 8QSFP Integrated Fibre Channel Blade....78
Cisco Nexus 5548UP/5596UP Converged Network Switches....79
Cisco Nexus 6001P switch....80
Cisco Nexus 6004EF switch....80
Cisco Fabric Extender for HPE BladeSystem....80
Converged network adapters....81
FCoE storage systems....85
HPE 3PAR StoreServ 7200/7400/7450/8200/8400/8440/8450 AFA/10400/10800/20450/20800/20850 with 10GbE FCoE host ports....85
XP7 with the FCoE 10-GbE channel adapter (8 ports per adapter)....85
P9500 with the FCoE 10-GbE channel adapter (4 ports per adapter)....86
P63xx/P65xx with FCoE/iSCSI 10-GbE front end host ports (2 ports per controller)....87
FCoE configuration rules....91
Server support....96
HPE FlexFabric 5900CP Switch and HPE 5900AF Switch....97
HPE 5820 Converged Network Switch....97
Usage....97
FCoE Converged Network Switch fabric rules....97
C-series and Cisco FCoE Converged Network switches....98
Cisco Nexus 5548UP FCoE Converged Network Switch....98
Cisco Nexus 5596UP FCoE Converged Network Switch....98
Expansion Module Options for the Cisco Nexus 5500 Series switches....98
Expansion Modules for the Cisco Nexus 6004EF switch....99
Models....100
FCoE Converged Network switch fabric rules....101
Operating systems and storage products....101
Boot from SAN support....101
Multipathing software support....101

**Fabric infrastructure rules....102**

**HPE FlexFabric switches and storage support....103**
HPE FlexFabric converged switches....103
Model numbering....103
Model naming....104
Switch models....104
Features and storage support....104

**B-series switches and fabric rules....105**
B-series Fibre Channel switches, HPE StoreFabric SN4000B SAN Extension Switch, HPE StoreFabric SN8600B 32Gb SAN Extension Blade, 1606 Extension SAN Switches, and MP Router Blade....105
Model naming....106
Switch models....106
Features....109
Usage....116
Fibre Channel switch fabric rules....118
Operating systems and storage products....120
B-series Fibre Channel switch and fabric rules....120



https://t.me/learningnets

B-series Encryption Switch fabric rules..... 124
HPE StoreFabric SN8600B 32Gb SAN Extension Blade and SN4000B SAN Extension Switch fabric rules..... 125
1606 Extension SAN Switch and DC Dir Switch MP Extension Blade fabric rules.....127
Zoning limits and enforcement..... 133
Zoning guidelines for B-series Fibre Channel switches..... 134

**C-series switches and fabric rules..... 137**
C-series Fibre Channel switches..... 137
Model naming.....139
Switch models..... 140
Features..... 141
Usage.....143
Fibre Channel switch fabric rules.....145
Operating systems and storage products.....145
Fabric rules for C-series Fibre Channel switches.....146
ISL maximums.....147
Zoning.....148
Zoning limits and enforcement..... 149
C-series VSAN high availability.....149

**H-series switches and fabric rules..... 151**
H-series switches.....151
Model numbering.....152
Model naming.....152
Switch models..... 152
Features..... 153
Usage.....155
Fabric rules.....156
Servers, operating systems, and storage products..... 156
Fabric rules for H-series switches..... 157
ISL maximums.....157
Fabric rules for H-series switches with TR..... 157
Zoning limits and enforcement..... 160

**McDATA/M-series switches ..... 162**

**SAN fabric connectivity and switch interoperability rules..... 163**
SAN fabric connectivity rules..... 163
Switch port interfaces.....163
Device port interfaces.....163
Fiber optic cables..... 164
Fiber optic cable loss budgets.....165
Storage product interface, switches, and transport distance rules.....169
SAN fabric switch interoperability rules.....174
Dual interoperable, heterogeneous SAN fabrics.....174
Interoperable, heterogeneous switch fabrics.....174
B-Series, C-Series, and H-Series switches..... 174
HPE FlexFabric switches.....175
Third-party switch support.....175
SAN performance considerations..... 175



https://t.me/learningnets

Infrastructure factors....175
Performance guidelines....176

**Host and storage system rules....178**

**Heterogeneous server rules....179**
SAN platform rules....180
Heterogeneous storage system support....180
HPE FC Switches for the c-Class BladeSystem and HPE Synergy server environment....181
HPE Virtual Connect for the c-Class BladeSystem server environment....181
HPE Virtual Connect 20/40 F8....182
HPE Virtual Connect FlexFabric 10 Gb/24-Port Module for c-Class BladeSystem....182
HPE Virtual Connect Flex-10/10D Ethernet Module for c-Class BladeSystem....183
HPE Virtual Connect 8 Gb 20-Port Fibre Channel Module for c-Class BladeSystem....183
HPE Virtual Connect 8 Gb 24-Port Fibre Channel Module for c-Class BladeSystem....183
HPE Virtual Connect FC connectivity guidelines....184
BladeSystem with Brocade Access Gateway mode....184
Failover policy and failback policy....185
AG mode considerations....186
AG mode connectivity guidelines....186
BladeSystem with Cisco N_Port Virtualization mode....187
Failover policy....188
NPV mode considerations....188
NPV mode connectivity guidelines....189
NPV with FlexAttach....190
HPE BladeSystem c3000 enclosure considerations....191
HBA N_Port ID Virtualization....192
HBA NPIV considerations....192
HBA NPIV connectivity guidelines....192
NonStop servers (XP only)....193
HP-UX SAN rules....207
HPE OpenVMS SAN rules....210
Host-based volume shadowing....212
HPE Tru64 UNIX SAN rules....212
Apple Mac OS X SAN rules....212
IBM AIX SAN rules....214
Linux SAN rules....216
Linux multipath software coexistence support....217
Microsoft Windows SAN rules....218
Oracle Solaris SAN rules....219
VMware ESX SAN rules....222
VMware ESX HBA and multipath software coexistence support....222
Citrix Xen SAN rules....223
Heterogeneous SAN storage system coexistence....224
Common SAN storage coexistence....224
Common server access, different storage system types....224
Common server, common HBA....225
Common server, different HBAs....225
Server zoning rules....226



https://t.me/learningnets

## MSA storage system rules ..... 227
HPE MSA storage system configurations ..... 227
Heterogeneous SAN support ..... 228
Configuration rules ..... 231
Zoning ..... 231
Maximums ..... 232
Management software support ..... 233

## HPE StoreVirtual 3200 and 4000 Storage System Rules ..... 234
Fibre Channel on HPE StoreVirtual 3200 Storage ..... 234
Fibre Channel on HPE StoreVirtual 4000 Storage ..... 234
Campus cluster support ..... 235
Heterogeneous SAN support ..... 235
Configuration rules ..... 236
Configuration parameters ..... 237
Data migration ..... 237
Zoning ..... 237
Boot from SAN support ..... 237

## P6000 storage system rules ..... 238
P6000/EVA storage ..... 238
Heterogeneous SAN support ..... 239
Configuration rules ..... 239
Configuration parameters ..... 242
P6000/EVA data migration ..... 243
MPX200 Multifunction Router with data migration ..... 243
Data migration considerations ..... 246
HPE P6000 Continuous Access SAN integration ..... 247
Zoning ..... 248
Tape storage ..... 248
P6000/EVA SAN boot support ..... 249
Storage management server integration ..... 249
Cabling ..... 250
Dual-channel HBA configurations ..... 250

## P9000/XP storage system rules ..... 252
P9000/XP storage systems ..... 252
Heterogeneous SAN support ..... 253
Configuration rules ..... 254
Zoning ..... 255
Tape storage ..... 255
P9000/XP SAN boot support ..... 256
LUN Configuration and Security Manager XP support ..... 257
P9000/XP data migration ..... 257
MPX200 Multifunction Router with data migration ..... 257
Data migration considerations ..... 259

## HPE 3PAR StoreServ storage rules ..... 260
HPE 3PAR StoreServ storage ..... 260
Heterogeneous SAN support ..... 260




https://t.me/learningnets

HPE 3PAR StoreServ Host Interfaces....261
Configuration rules....262
Configuration parameters....264
Virtual Connect Flat SAN (Direct-attach) Fibre Channel for 3PAR Storage....264
HPE 3PAR data migration....265
HPE 3PAR Online Import Utility....265
MPX200 Multifunction Router with data migration....265
Zoning....269
3PAR zoning best practices....269
Cabling....270
Tape storage....270
HPE 3PAR SAN boot support....271
HPE 3PAR storage management....271

**HPE Data Availability, Protection and Retention....272**

**SAN extension and bridging....273**

**SAN extension....274**
SAN extension overview....274
SAN extension technology....275
SAN-iSCSI bridging technology....276
Fibre Channel long-distance technology....277
Fiber optic transceivers....277
Wavelength division multiplexing....279
WDM overview....279
WDM network implementation....279
WDM system architectures....279
WDM system characteristics....280
HPE coarse wave division multiplexing....281
Third-party WDM products....282
Extended fabric settings for Fibre Channel switches....283
B-series switch settings....283
B-series trunking and WDM....283
C-series switch settings....284
H-series switch settings....284
Multi-protocol long-distance technology....284
Fibre Channel over Internet Protocol....285
FCIP mechanisms....285
FCIP link configurations....285
FCIP network considerations....286
FCIP bandwidth considerations....287
FCIP gateways....289
Third-party QoS and data encryption FCIP products....289
WAN accelerator products....289
Fibre Channel over SONET....289
FC-SONET IP link configurations....290
FC-SONET network considerations....291
Third-party SONET gateways....291
Fibre Channel over ATM....291
HPE multi-protocol long-distance products....291
HPE SAN extension products summary and usage....291
HPE IP Distance Gateway (mpx110)....293



https://t.me/learningnets

IP Distance Gateway configuration examples....294
Configuration rules....301
MPX200 Multifunction Router with FCIP....302
MPX200 Multifunction Router FCIP configuration examples....303
FCIP Configuration rules....305
HPE StoreFabric SN4000B SAN Extension Switch....306
HPE StoreFabric SN4000B SAN Extension Switch features and requirements....307
HPE StoreFabric SN4000B SAN Extension Switch configuration examples....308
B-series 1606 Extension SAN Switch and DC Dir Switch MP Extension Blade..310
1606 Extension SAN Switch and DC Dir Switch MP Extension Blade features and requirements....311
1606 Extension SAN Switch configuration examples....312
C-series MDS 9222i, IPS-4, IPS-8, 18/4 Multiservice Modules....314
HPE storage replication products....315
SAN extension best practices for HPE P6000 Continuous Access....316
HPE P6000 Continuous Access with XCS 11x, XCS 10x, or XCS 09x....316
HPE P6000 Continuous Access with XCS 6.x....319
HPE P6000 Continuous Access with VCS 4.x....321
HPE P9000 (XP) Continuous Access....323
OpenVMS host-based volume shadowing....339
Certified third-party WDM, iFCP, and SONET products....340
Certified third-party WDM products....340

**iSCSI storage....341**
iSCSI overview....341
iSCSI and Fibre Channel....341
iSCSI concepts....342
Initiator and target devices....342
iSCSI naming....342
Discovery mechanisms....343
Service Location Protocol....343
Static configuration....343
SendTargets command....343
Internet Storage Name Service....343
Sessions and logins....344
Security....344
Software and hardware iSCSI initiators....345
iSCSI boot....345
iSCSI storage network requirements....345
HPE Native iSCSI products....346
HPE 3PAR StoreServ 20000 and StoreServ 8000....346
HPE 3PAR StoreServ 20000 and 8000 iSCSI overview....346
HPE 3PAR StoreServ 20000 and HPE StoreServ 8000 10 GbE iSCSI support....346
Operating system and multipath software support....346
HPE 3PAR StoreServ10000 and 7000....347
HPE 3PAR StoreServ 10000 iSCSI overview....347
HPE 3PAR StoreServ 10000 and StoreServ 7000 10 GbE iSCSI support....347
HPE 3PAR F-Class, T-Class....348
HPE 3PAR F-Class, T-Class iSCSI overview....348
HPE 3PAR F-Class, T-Class iSCSI support....348
P6300/P6350/P6500/P6550 EVA....348


https://t.me/learningnets

P6300/P6350/P6500/P6550 EVA overview....349
P6300/P6350/P6500/P6550 EVA iSCSI support....349
HPE StorageWorks MSA family of iSCSI SAN arrays....349
**HPE MSA 2050/2052 SAN Overview** ....350
**HPE MSA 1050 iSCSI SAN Overview** ....350
HPE MSA 2040 SAN overview....350
HPE MSA 1040 iSCSI overview....350
P2000 G3 FC/iSCSI, P2000 G3 10Gb iSCSI, P2000 G3 iSCSI overview....351
MSA iSCSI storage family maximum configurations ....352
HPE StoreVirtual Storage....355
HPE StoreVirtual Storage overview....356
HPE StoreVirtual 3200 Storage support....357
HPE StoreVirtual 4000 Storage support....357
HPE iSCSI bridge products....358
Bridging and routing....358
iSCSI bridge to Fibre Channel....359
MPX200 Multifunction Router with iSCSI for P6000/EVA storage....359
MPX200 simultaneous operation....359
MPX200 configuration options....360
MPX200 iSCSI rules and supported maximums ....363
MPX200 blade configurations....364
P6000/EVA storage system rules and guidelines ....365
HPE P6000 Command View and MPX200 management rules and guidelines....365
P6000/EVA storage system software....366
Fibre Channel switch and fabric support....366
Operating system and multipath software support....366
MPX200 Multifunction Router with iSCSI for 3PAR StoreServ Storage....366
MPX200 configuration options....367
MPX200 iSCSI rules and supported maximums....369
3PAR storage system rules and guidelines....369
MPX200 Multifunction Router with iSCSI for XP storage....370
MPX200 configuration options....371
MPX200 iSCSI rules and supported maximums....373
XP storage system rules and guidelines....373
Operating system and multipath software support....374
B-series iSCSI Director Blade....374
Blade overview....374
Hardware support....375
Software support....375
Scalability rules....376
C-series iSCSI....377
Modules overview....377
Hardware support....378
Software support....379
Configuration rules....380
HPE ProLiant Storage Server iSCSI Feature Pack....381
Overview....381
HPE ProLiant Storage Server iSCSI Feature Pack support....381
HPE ProLiant Storage Server iSCSI license upgrade options....382
Designing a Microsoft Exchange solution with iSCSI Feature Pack....383
Sample iSCSI NAS Microsoft Exchange Server 2003 configuration....385

**HPE SAN Product Programs and Applications ....387**



https://t.me/learningnets

**HPE Complete Program ..... 388**

**HPE Smart SAN for 3PAR Application..... 389**

**Storage security, best practices, and support information..... 390**

**Storage security..... 391**
Storage security threats..... 391
Storage security compliance..... 393
Security technologies..... 393
IP SAN security technologies..... 394
Fibre Channel SAN security technologies..... 394
Encryption security technologies..... 395
Key management..... 395
Organizational security policies..... 395
HPE security strategy..... 395
HPE Secure Advantage..... 396
Resource protection..... 396
Data protection..... 396
Security validation..... 397
Storage security best practices..... 397
Assessing security risks..... 398
Managing organizational risks..... 398
Data security implementations..... 398
HPE storage security solutions..... 398
C-series Storage Media Encryption..... 399
C-series NX-OS security..... 399
C-series IP SAN security..... 400
B-series Encryption Switch and Encryption FC Blade security..... 401
B-series Fabric OS security..... 402
Resource protection..... 402
Data protection..... 404
Security validation..... 405
Key management..... 405

**Best practices..... 407**
SAN planning..... 407
Design specification..... 408
SAN topology..... 408
Multi-fabric SANs..... 408
SAN and fabric monitoring..... 408
Failover protection..... 409
Data access patterns..... 409
ISL ratio..... 409
Incremental SAN expansion..... 409
SAN configuration ..... 409
Fibre Channel switch configuration..... 410
Server setup..... 410
Storage system configuration..... 410
Storage-based LUN masking..... 411



https://t.me/learningnets

Zoning....411
Zoning enforcement....411
Zoning guidelines....412
Zoning by operating system....412
Zoning by HBA....413
Zoning by HBA port....413
Zoning by NPIV port....413
Zoning with 3PAR persistent ports....413
Zoning by application....413
Zoning by port allocation....413
EBS zoning....414
Zone naming....414
Naming by identifier type....414
Case sensitivity of fabric identifiers....414
Server naming....414
Storage system naming....414
FCoE switch configuration quick-setup instructions....415
HPE 5820 FCoE Converged Network Switch quick setup....415
DC SAN Director Switch 10/24 FCoE Blade quick setup....418
HPE C-series Nexus 5548UP/5596UP Converged Network Switch quick setup....421
SAN scaling....426
Cascaded fabric expansion....427
Meshed fabric expansion....427
Ring fabric expansion....427
Core-edge fabric expansion....427
SAN fabric merging....427
Fabric segmentation errors....428
Switch configuration parameters....429
Independent fabric merge....429
High-availability redundant fabric merge....429
SAN infrastructure monitoring....430

**Support and other resources....431**
Accessing Hewlett Packard Enterprise Support....431
Accessing updates....431
Customer self repair....432
Remote support....432
Warranty information....432
Regulatory information....433
Documentation feedback....433

**Related information....434**



https://t.me/learningnets

# New and changed information in this edition

Changes to the July 2018 version of this guide include:

- Removed EOSL FCoE products
- Added information for HPE FlexFabric 5940 and 5950 switches
- Added information to the B-series Switches and Fabric rules chapter
- Added information to the C-series Switches and Fabric rules chapter



https://t.me/learningnets

# Architecture

SAN architecture is presented in these chapters:


- **SAN design overview** on page 16
- **SAN fabric topologies** on page 24
- **Fibre Channel routing** on page 46
- **Fibre Channel over Ethernet** on page 65




https://t.me/learningnets

# 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. HPE 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** on page 16
- **HPE SAN implementations** on page 17
- **SAN components** on page 17
- **Fibre Channel technology** on page 18
- **Storage area networks** on page 18
- **SAN infrastructure** on page 19
- **Fibre Channel switches** on page 20
- **SAN design approaches** on page 20
- **SAN design considerations** on page 21


## 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.



https://t.me/learningnets

- **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.

# HPE SAN implementations

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

- **Flexible design and deployment**

  Hewlett Packard Enterprise 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**

  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**

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

- **Geographically dispersed installations**

  Hewlett Packard Enterprise provides components to meet local and long-distance connectivity requirements.

For information about SAN infrastructure solutions, see the Hewlett Packard Enterprise Storage Networking website at **http://www.hpe.com/info/StoreFabric**

# SAN components

A SAN consists of the following hardware and software components:



https://t.me/learningnets

- **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**

  Hewlett Packard Enterprise 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



https://t.me/learningnets

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. Hewlett Packard Enterprise 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 Hewlett Packard Enterprise-supported switch and port count fabric maximums, see:

  - **B-series switches and fabric rules** on page 105
  - **C-series switches and fabric rules** on page 137
  - **H-series switches and fabric rules** on page 151

- **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.



https://t.me/learningnets

# 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, HPE 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

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

- **HPE standard design**

  Hewlett Packard Enterprise 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. Hewlett Packard Enterprise recommends this approach for users who are designing a SAN for the first time.

- **Modified HPE 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. Hewlett Packard Enterprise recommends this approach for users with an intermediate level of SAN experience.

- **Custom design using the HPE 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. Hewlett Packard Enterprise recommends this approach for users with an intermediate or advanced level of SAN experience.

For information about:

- Standard SAN designs, see **SAN fabric topologies** on page 24
- Customizing a SAN design, see:
  - **B-series switches and fabric rules** on page 105
  - **C-series switches and fabric rules** on page 137
  - **H-series switches and fabric rules** on page 151
- Heterogeneous SAN design, see:
  - **Heterogeneous server rules** on page 179
  - **MSA storage system rules** on page 227



https://t.me/learningnets

- **HPE StoreVirtual 3200 and 4000 Storage System Rules** on page 234
- **P6000 storage system rules** on page 238
- **P9000/XP storage system rules** on page 252
- **HPE 3PAR StoreServ storage rules** on page 260

- Recommended SAN solutions and conventions, see **Best practices** on page 407

# 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** on page 105.

- **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** on page 40.

- **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** on page 24.

- **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:

  - **MSA storage system rules** on page 227
  - **HPE StoreVirtual 3200 and 4000 Storage System Rules** on page 234
  - **P6000 storage system rules** on page 238
  - **P9000/XP storage system rules** on page 252
  - **HPE 3PAR StoreServ storage rules** on page 260

- **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.



https://t.me/learningnets

For information about configuring systems in a heterogeneous environment, see **Heterogeneous server rules** on page 179.

- **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** on page 407.

- **Backup and restore**

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

  For information about centralized backup, see **HPE Data Availability, Protection and Retention** on page 272.

- **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** on page 274.

- **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** on page 105
  - **C-series switches and fabric rules** on page 137
  - **H-series switches and fabric rules** on page 151

- **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 **Core-edge fabric types** on page 32.

- **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** on page 24.

- **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.


https://t.me/learningnets

For information about zoning, see:

- **B-series switches and fabric rules** on page 105
- **C-series switches and fabric rules** on page 137
- **H-series switches and fabric rules** on page 151

- **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** on page 24.

- **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** on page 24.



https://t.me/learningnets

# SAN fabric topologies

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


- **Fabric topologies** on page 24
- **Single-switch fabric** on page 25
- **Cascaded fabric** on page 26
- **Meshed fabric** on page 28
- **Ring fabric** on page 29
- **Core-edge fabric** on page 31
- **Topology data access** on page 34
- **Topology maximums** on page 35
- **Routed fabric topologies** on page 37
- **FCoE fabric topologies** on page 40
- **Data availability** on page 40
- **Topology migration** on page 44


## Fabric topologies

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

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

- A HPE standard SAN fabric topology design
- A subset or variation of a HPE 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** on page 105
- **C-series switches and fabric rules** on page 137
- **H-series switches and fabric rules** on page 151
- **SAN fabric connectivity and switch interoperability rules** on page 163
- **Heterogeneous server rules** on page 179
- **MSA storage system rules** on page 227
- **HPE StoreVirtual 3200 and 4000 Storage System Rules** on page 234
- **P6000 storage system rules** on page 238




https://t.me/learningnets

- **P9000/XP storage system rules** on page 252
- **HPE 3PAR StoreServ storage rules** on page 260

## Routed SAN fabrics

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

- **B-series Meta SAN** on page 37
- **C-series VSANs with IVR** on page 39
- **H-series switches with TR** on page 39

## FCoE SAN fabrics

HPE 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** on page 40.

## Benefits

With HPE 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** on page 163 and **SAN extension** on page 274.
- Connect multiple SAN fabrics using routing technology, see **B-series Meta SAN** on page 37.
- Deploy multiple logical fabrics using the Virtual Fabrics feature.
- Deploy multiple VSANs, see **C-series VSANs with IVR** on page 39.
- Incorporate a range of SAN availability levels, see **Data availability** on page 40.

# Single-switch fabric

A single-switch fabric consists of a Fibre Channel switch, server, and storage system (**Figure 1: Single-switch fabric** on page 26). This topology forms the basis for all HPE 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.



https://t.me/learningnets

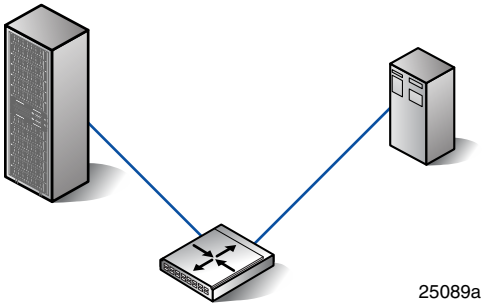



**Figure 1: Single-switch fabric**

## Switch models

For a small, single-switch SAN fabric, use a SAN, Fabric, or Edge switch (4, 8, 16, 24, 28, or 32 ports). For a larger single-switch SAN fabric, use a SAN, Fabric, or Edge switch (32 to 96 ports), or a Core or Director switch (48 to 512 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: Cascaded fabric** on page 27). 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. Hewlett Packard Enterprise 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



https://t.me/learningnets

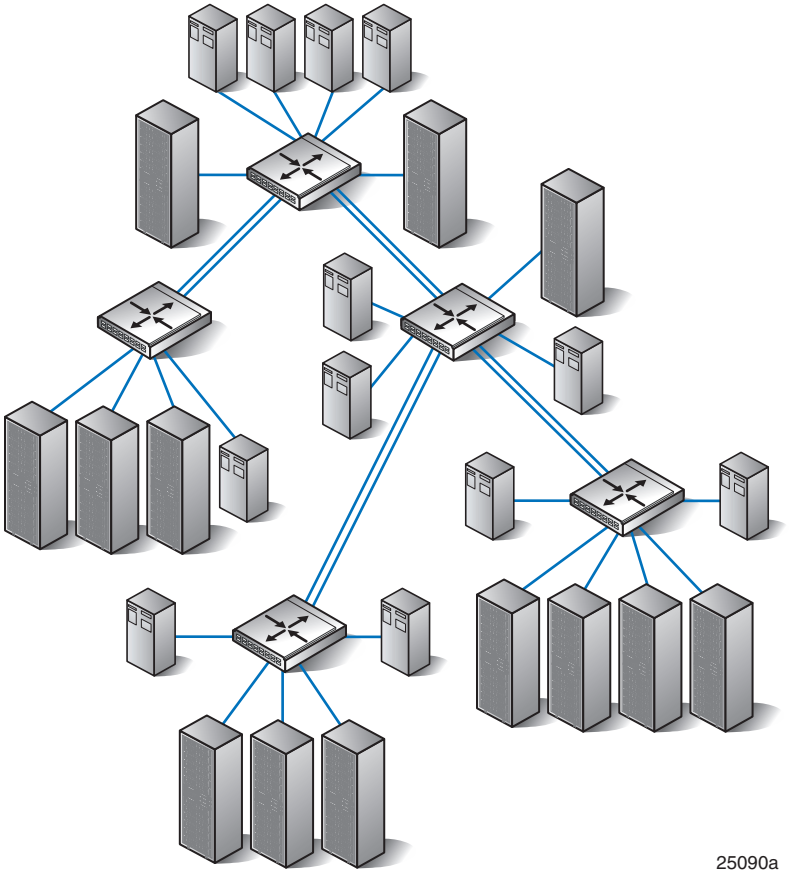



**Figure 2: Cascaded fabric**

## Switch models

All HPE Fibre Channel switches are supported for use in a cascaded fabric topology. Cascaded fabric topologies typically use 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. For additional switch hop information, see:

- **B-series switches and fabric rules** on page 105
- **C-series switches and fabric rules** on page 137
- **H-series switches and fabric rules** on page 151

## 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



https://t.me/learningnets

- 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: Meshed fabric** on page 28). 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.

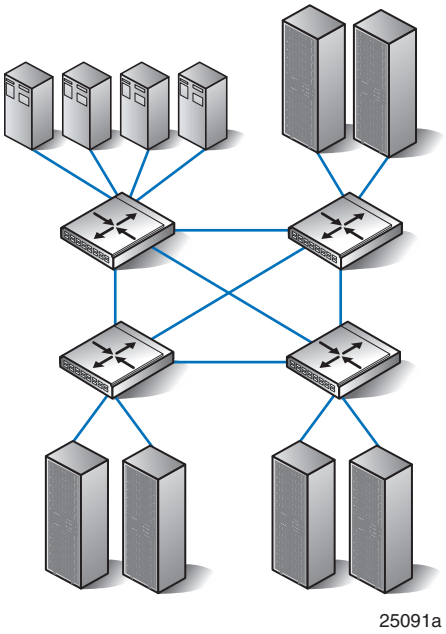



**Figure 3: Meshed fabric**

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: ISL connections in a meshed fabric** on page 29). 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.



https://t.me/learningnets

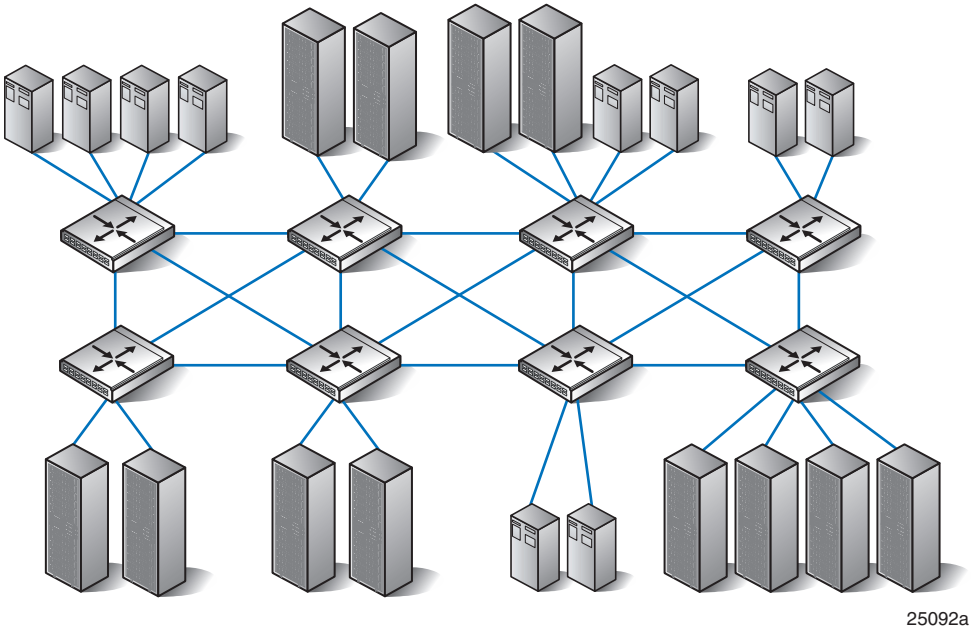



**Figure 4: ISL connections in a meshed fabric**

## Switch models

All Fibre Channel switches are supported for use in a meshed fabric topology. Meshed fabric topologies typically use 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** on page 34)
- Less impact on performance due to intraswitch traffic

# Ring fabric

A ring fabric is a ring of interconnected switches (**Figure 5: Ring fabric** on page 30). 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.



https://t.me/learningnets

**NOTE:**

Hewlett Packard Enterprise does not recommend the ring fabric for applications requiring many-to-many connectivity.

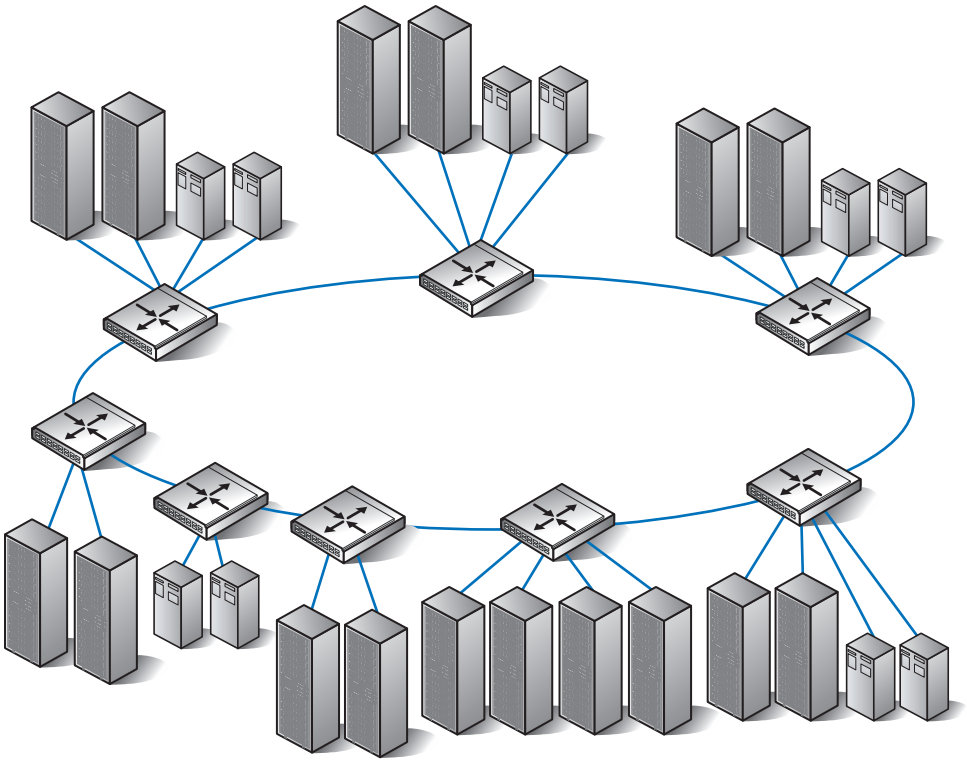



**Figure 5: Ring fabric**

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: Ring fabric with satellite switches** on page 31). 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** on page 105
- **C-series switches and fabric rules** on page 137
- **H-series switches and fabric rules** on page 151



https://t.me/learningnets

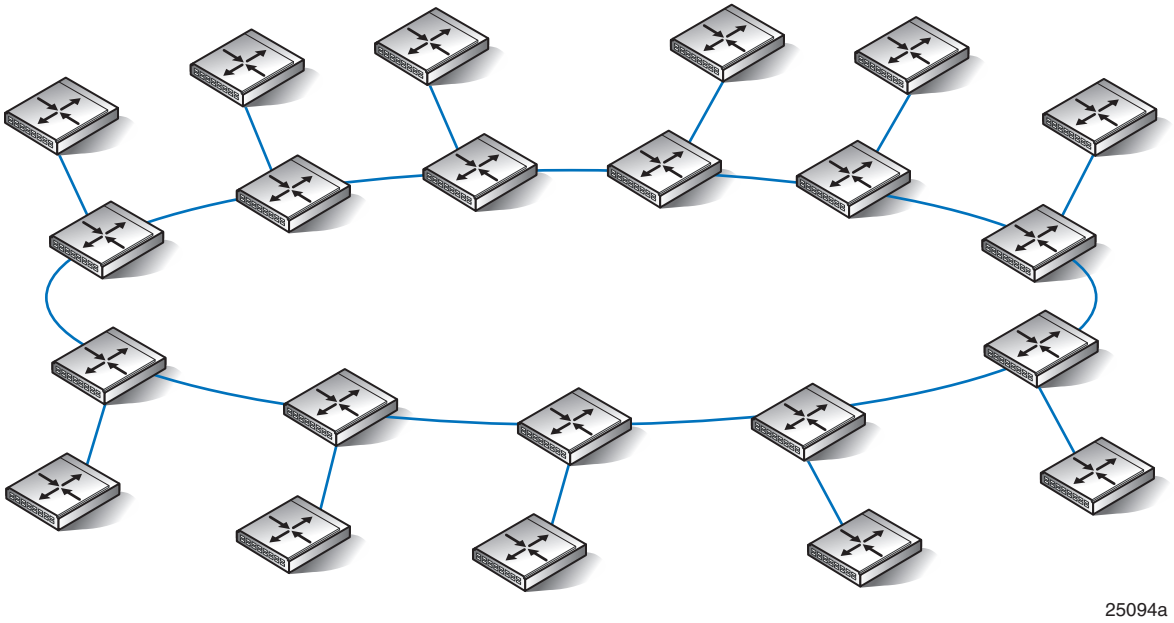



**Figure 6: Ring fabric with satellite switches**

## Switch models

All Fibre Channel switches are supported for use in a ring fabric topology. Ring fabric topologies typically use 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

Hewlett Packard Enterprise 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: Core-edge fabric (typical depiction)** on page 32). 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



https://t.me/learningnets

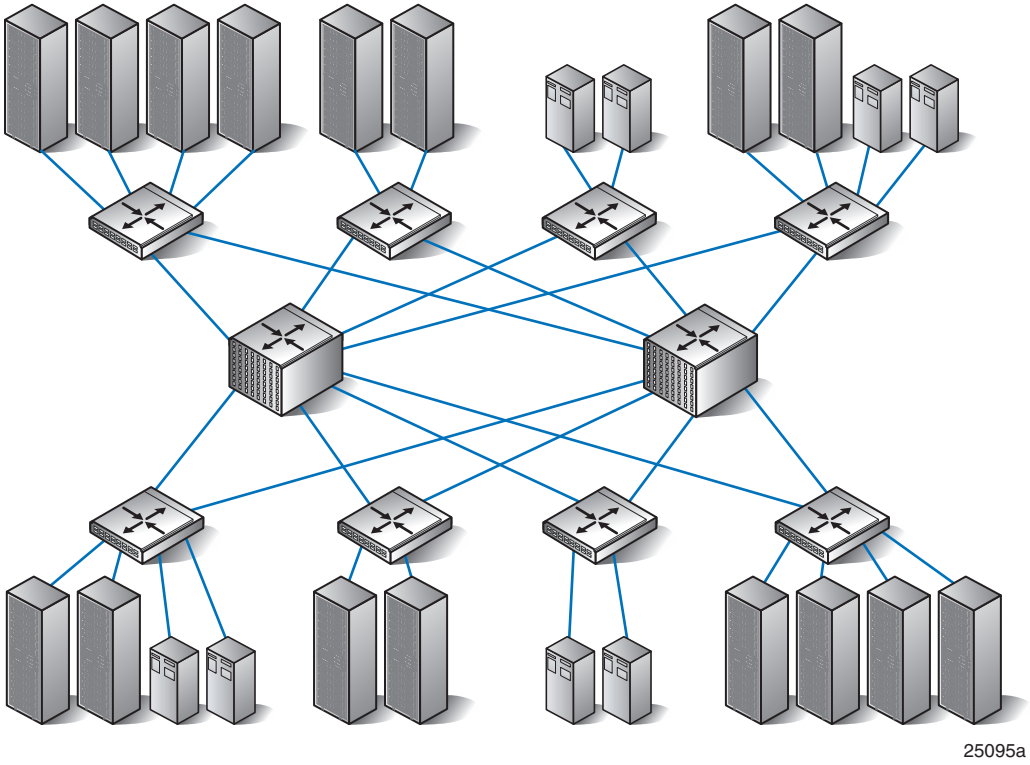



**Figure 7: Core-edge fabric (typical depiction)**

Core-edge fabric topologies are typically depicted as shown in **Figure 7: Core-edge fabric (typical depiction)** on page 32, but can also be depicted hierarchically as shown in **Figure 8: Core-edge fabric (hierarchical depiction)** on page 32. 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.

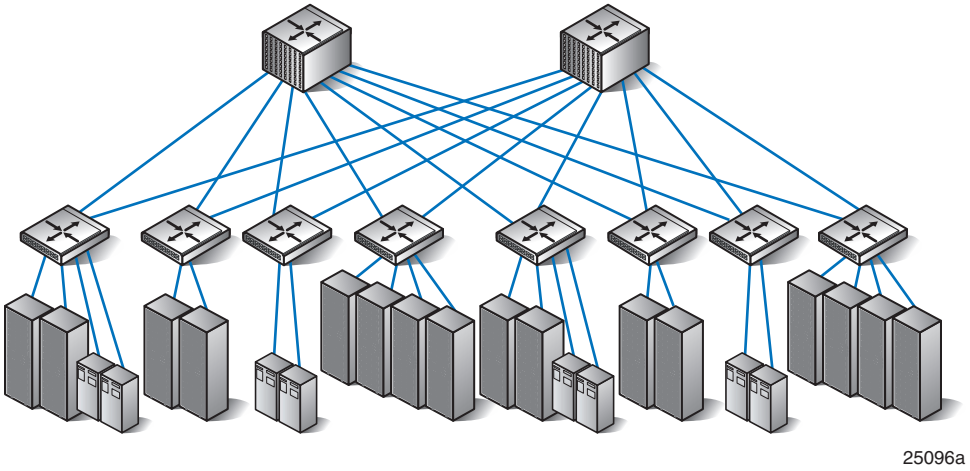



**Figure 8: Core-edge fabric (hierarchical depiction)**

## 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



https://t.me/learningnets

There are two core-edge fabric topology types: fat tree and skinny tree. Core-edge fabric topology types table 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** on page 34)
- Server application performance requirements

Recommended core-edge fabric ISL ratios table 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, > 560 MB/s at 16 Gb/s, > 1120 MB/s at 32 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, < 560 MB/s at 16 Gb/s, < 1120 MB/s at 32 Gb/s) | 7:1 to 15:1 |

**NOTE:** Hewlett Packard Enterprise 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 **n**1 x **n**2, where **n**1 represents the number of core switches and **n**2 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.



https://t.me/learningnets

**Figure 9: Core-edge fabric (4 x 12)** on page 34 shows a 4 x 12 core-edge fabric with 4 core switches and 12 edge switches, each connected to the core with 4 ISLs.

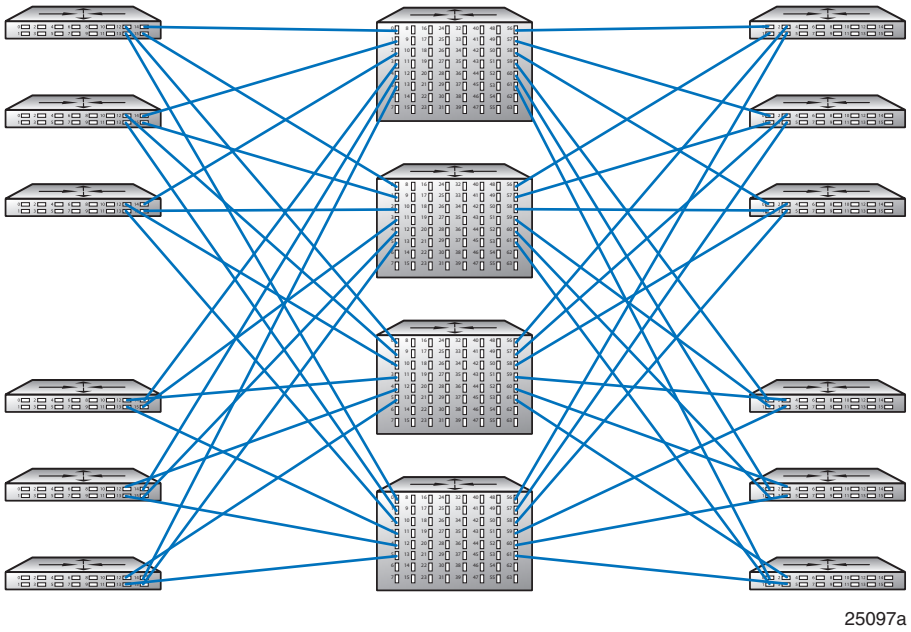



**Figure 9: Core-edge fabric (4 x 12)**

## Switch models

All Fibre Channel switches are supported for use in a core-edge fabric topology. Core-edge topologies typically use SAN, Fabric, or Edge switches on the edge; and Core and Director switches in the core. When using switches with different Fibre Channel maximum speed capabilities (such as 1 Gb/s, 2 Gb/s, 4 Gb/s, 8 Gb/s, 16 Gb/s, or 32 Gb/s), Hewlett Packard Enterprise 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:



https://t.me/learningnets

- **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: 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

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** on page 105
  - **C-series switches and fabric rules** on page 137
  - **H-series switches and fabric rules** on page 151
- If you connect a Storage Management Appliance to the fabric, this further reduces the number of ports available for server and storage connections.



https://t.me/learningnets

## B-series switches

### 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: 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<br>(maximum of 12 Director switches) | 3,500<br>(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

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.



https://t.me/learningnets

**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 | | | |

# Routed fabric topologies

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

This section describes the following Fibre Channel routed fabric technologies:

- **B-series Meta SAN** on page 37 — 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, Multi-protocol Router Blade (MP Router Blade FR4-18i), or HPE StoreFabric SN4000B SAN Extension Switch, which provide selective Fibre Channel routing connectivity between multiple B-series fabrics.
- **B-series Virtual Fabrics with IFR** on page 38 — 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 partions.
- **C-series VSANs with IVR** on page 39 — 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** on page 39 — 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, StoreFabric SN4000B SAN Extension Switch, 1606 Extension SAN Switches, DC Director Switch MP Extension Blades, or MP Router Blade. The 8 Gb/s switches with Fibre Channel routing, StoreFabric SN4000B SAN Extension Switch, 1606 Extension SAN Switches, DC Dir Switch MP Extension Blades, or MP Router Blade 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.



https://t.me/learningnets

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** on page 46.

**B-series Virtual Fabrics architecture**

In B-series configurations, a switched fabric is partitioned into several logical switches and logical fabrics (**Figure 15: Basic Virtual Fabric IFR configuration** on page 50) 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 5.2**x** (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**

Hewlett Packard Enterprise supports Meta SANs with all B-series switches in either HPE 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** on page 105.

**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



https://t.me/learningnets

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**

Hewlett Packard Enterprise supports Virtual Fabrics in either HPE standard or customized topologies. You must follow all B-series fabric rules. For more information, see **B-series switches and fabric rules** on page 105.

**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**

Hewlett Packard Enterprise supports VSANs with all C-series switches in either HPE 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** on page 137.

**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.



https://t.me/learningnets

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**

Hewlett Packard Enterprise supports the SN6000 Fibre Channel Switch or the 8/20q Fibre Channel Switch with TR with B-series or C-series switches in either HPE 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 table. You must follow all H-series switch fabric rules. For more information, see **H-series switches and fabric rules** on page 151.

**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 Hewlett Packard Enterprise-supported Fibre Channel topologies are supported integrated with FCoE at the Fibre Channel fabric edge. Hewlett Packard Enterprise also supports end-to-end FCoE solutions for the P63**xx**/P65**xx** EVA storage systems with the FCoE target interface, HPE 3PAR StoreServ Storage systems with the FCoE target interface, 3PAR StoreServ Storage systems using a fabric connection between the MPX200 and Fibre Channel storage, and P6000 EVA storage using the MPX200 Multifunction Router. For more information about FCoE, see **Fibre Channel over Ethernet** on page 65.

# 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 HPE B-series MP routing function or C-series VSANs, see **High-availability router configurations** on page 59.

Several factors affect SAN data availability:

- Application software
- Server operating systems



https://t.me/learningnets

- 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: Level 1: single connectivity fabric** on page 41). Each server and storage system has one path to the fabric.

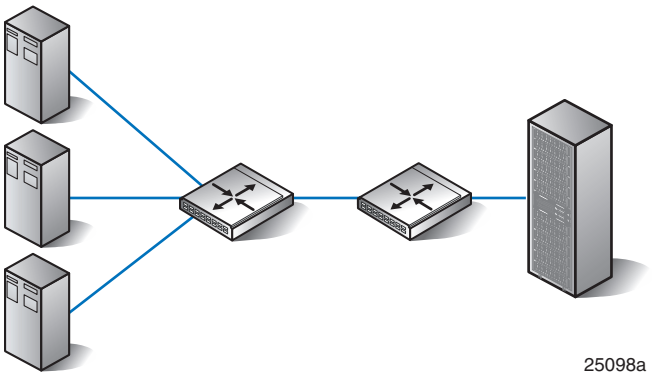



**Figure 10: Level 1: single connectivity fabric**

**Level 2: single resilient fabric**

Level 2 provides fabric path redundancy by using multiple ISLs between switches (**Figure 11: Level 2: single resilient fabric** on page 42). 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.



https://t.me/learningnets

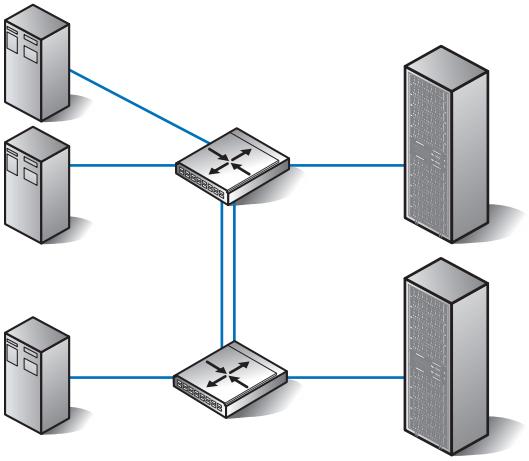



**Figure 11: Level 2: single resilient fabric**

### 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: Level 3: single resilient fabric with multiple device paths** on page 42). 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, Hewlett Packard Enterprise 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.

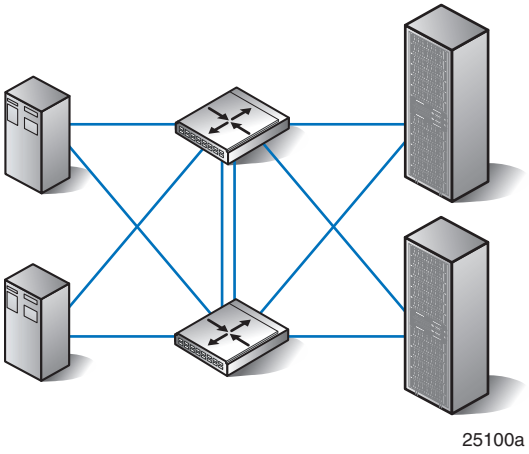



**Figure 12: Level 3: single resilient fabric with multiple device paths**

### 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: Level 4: multiple fabrics and device paths (NSPOF)** on page 43). 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.



https://t.me/learningnets

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.

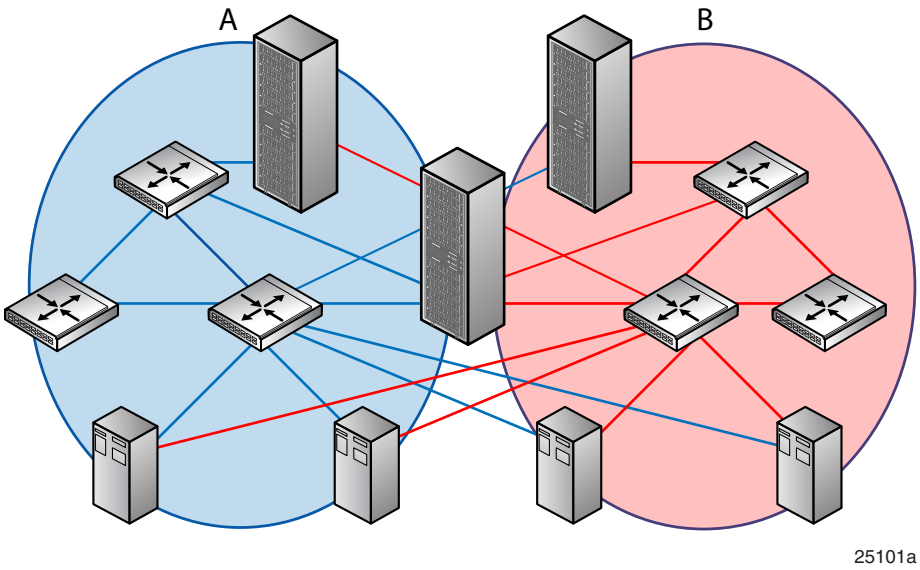



**Figure 13: Level 4: multiple fabrics and device paths (NSPOF)**

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, Hewlett Packard Enterprise 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 7: 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 |

*Table Continued*



https://t.me/learningnets

| Fabric design | Availability level | SAN topologies |
|---|---|---|
| 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 8: Calculating for data availability levels**

| Fabric design | Hardware cost | Number of available ports |
|---|---|---|
| Level 1: single connectivity fabric | **x** | # ports = **n** – number of ISL ports |
| 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 |
| Level 4: multiple fabrics and device paths (NSPOF) | **x** + additional ISLs + additional HBAs + additional switches | # ports = 2**n** – number of ISL ports – additional HBA ports |

# 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**



https://t.me/learningnets

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 **Migrating a cascaded fabric SAN** on page 44.

## 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 **Migrating a cascaded fabric SAN** on page 44.



https://t.me/learningnets

# Fibre Channel routing

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


- **Fibre Channel routing overview** on page 46
- **SAN scaling and routing** on page 48
- **Fibre Channel routing implementations** on page 49
- **Fabric redundancy and routing** on page 55
- **Supported routing configurations** on page 57


## 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. Hewlett Packard Enterprise does not support using the B-series Fibre Channel routing, HPE StoreFabric SN4000B SAN Extension Switch, 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.
- Hewlett Packard Enterprise 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



https://t.me/learningnets

- **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.



https://t.me/learningnets

## 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** on page 105
- **C-series switches and fabric rules** on page 137
- **H-series switches and fabric rules** on page 151

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, Hewlett Packard Enterprise 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:



https://t.me/learningnets

- 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



https://t.me/learningnets

## Fibre Channel routing techniques

This section describes the following Fibre Channel routing techniques:

- B-series routing connects independent fabrics (SAN islands), as shown in **Basic 1606 Extension Switch configuration**. Hewlett Packard Enterprise-supported B-series router products include:
  - 8 Gb/s switch models with license enabled integrated Fibre Channel routing
  - StoreFabric SN4000B SAN Extension Switch
  - HPE StorageWorks 1606 Extension SAN Switch (1606 Extension SAN Switch)
  - HPE StorageWorks DC SAN Director Multi-protocol Extension Blade (DC Dir Switch MP Extension Blade)
  - Multi-protocol Router Blade (MP Router Blade for the 4/256 SAN Director)
- A B-series switch with IFR connects multiple Virtual Fabrics, as shown in **Figure 15: Basic Virtual Fabric IFR configuration** on page 50.
- A C-series switch with IVR connects multiple VSANs, as shown in **Figure 16: Basic VSAN IVR configuration** on page 51.
- An H-series switch with TR connects to other B-series or C-series fabrics, as shown in **Figure 17: Basic TR configuration** on page 51.

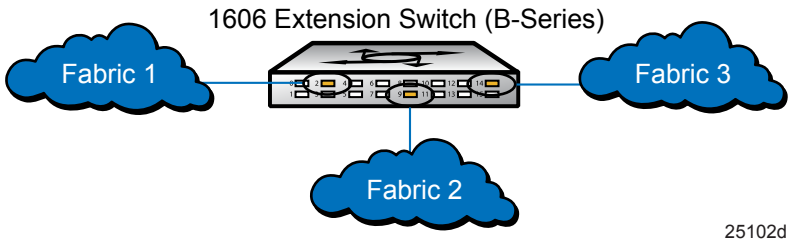



**Figure 14: Basic 1606 Extension Switch configuration**

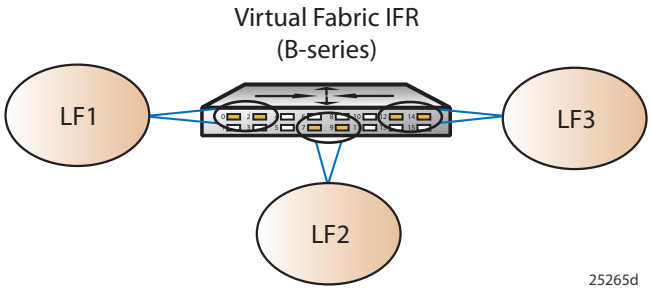



**Figure 15: Basic Virtual Fabric IFR configuration**



https://t.me/learningnets

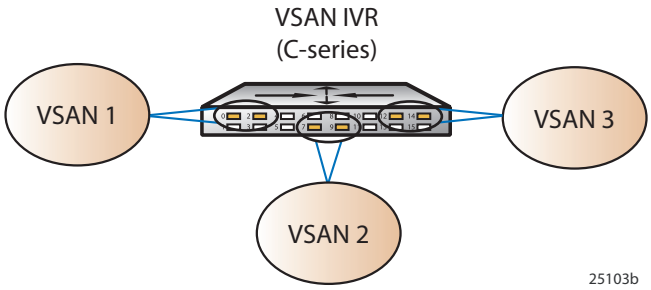



**Figure 16: Basic VSAN IVR configuration**

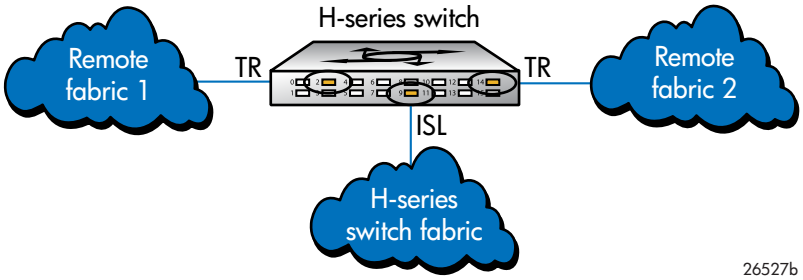



**Figure 17: Basic TR configuration**

## 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.

**Basic 1606 Extension Switch configuration** 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** on page 105.

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** on page 127 and **HPE StoreFabric SN8600B 32Gb SAN Extension Blade and SN4000B SAN Extension Switch fabric rules** on page 125. The fabrics can have identical domain names and zoning definitions.

The StoreFabric SN4000B SAN Extension Switch, 1606 Extension SAN Switches, or DC Dir Switch MP Extension Blades, and MP Router Blade also provide FCIP capabilities, allowing implementation of Fibre Channel routing and FCIP SAN extension, see **Routing use cases** on page 61.

## B-series fabric partitioning using Virtual Fabrics

In B-series configurations, a switched fabric is partitioned into several Virtual Fabrics (**Figure 15: Basic Virtual Fabric IFR configuration** on page 50). 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 5.2**x** (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: Basic Virtual Fabric IFR configuration** on page 50.



https://t.me/learningnets

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: Basic VSAN IVR configuration** on page 51 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. Hewlett Packard Enterprise currently supports connection to B-series and C-series remote fabrics.

**Figure 17: Basic TR configuration** on page 51 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** on page 157. 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, StoreFabric SN4000B SAN Extension Switch, 1606 Extension SAN Switches or DC Dir Switch MP Extension Blades, or MP Router Blade, 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 Blade 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: B-series routing** on page 53, an LSAN can include devices connected to different fabrics (for example, the LSAN Zone connects devices from Fabric 1 and Fabric 2).



https://t.me/learningnets

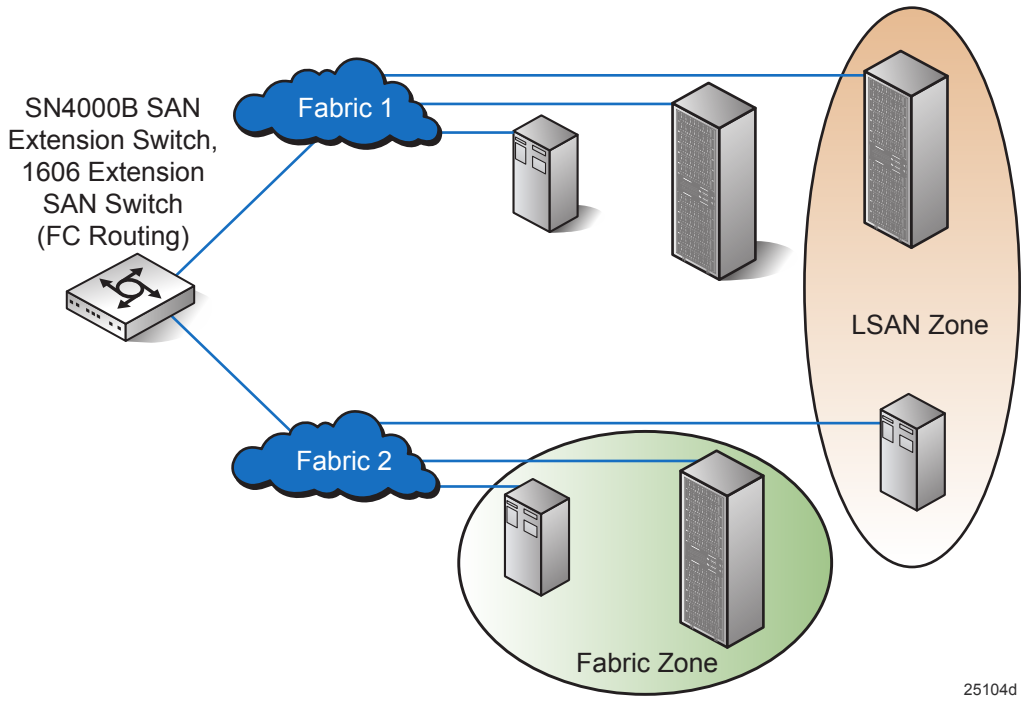



**Figure 18: B-series routing**

**Figure 19: B-series Virtual Fabric IFR** on page 53 and **Figure 20: C-series VSAN IVR** on page 54 show the differences between B-series routing and C-series routing.

**Figure 19: B-series Virtual Fabric IFR** on page 53 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.

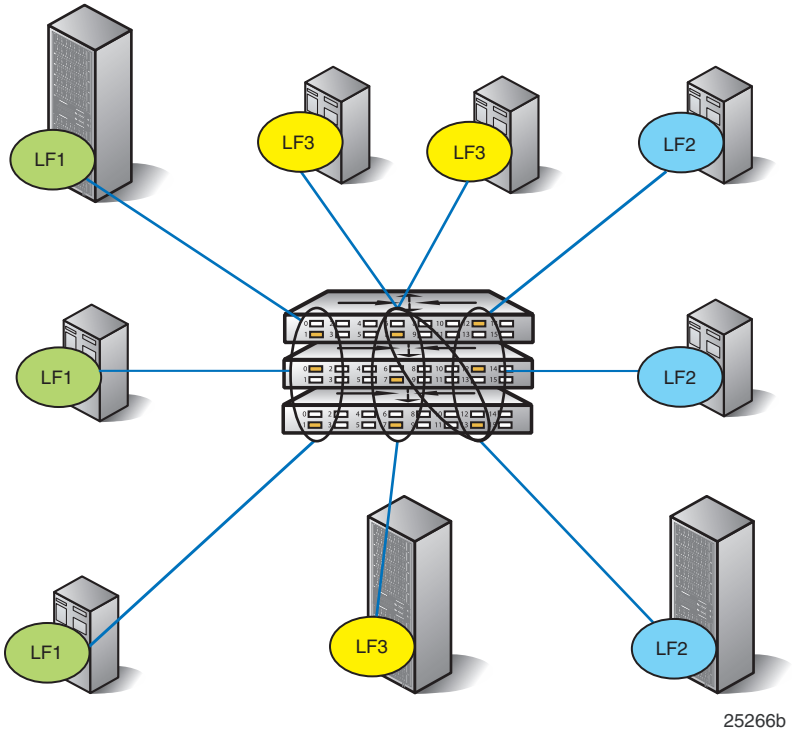



**Figure 19: B-series Virtual Fabric IFR**



https://t.me/learningnets

As shown in **Figure 20: C-series VSAN IVR** on page 54, 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.

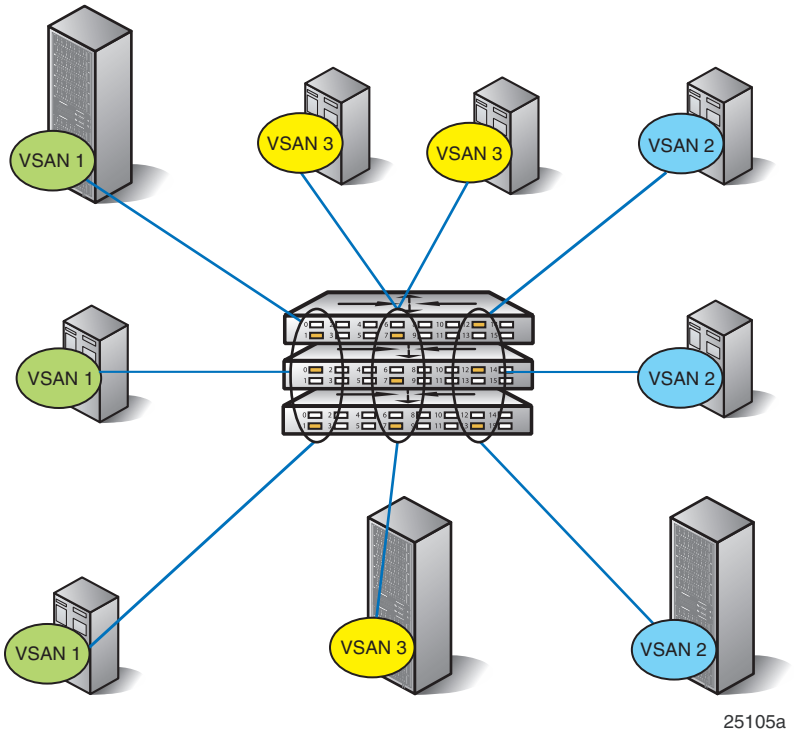



**Figure 20: C-series VSAN IVR**

**Figure 21: H-series switch routing** on page 55 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.



https://t.me/learningnets

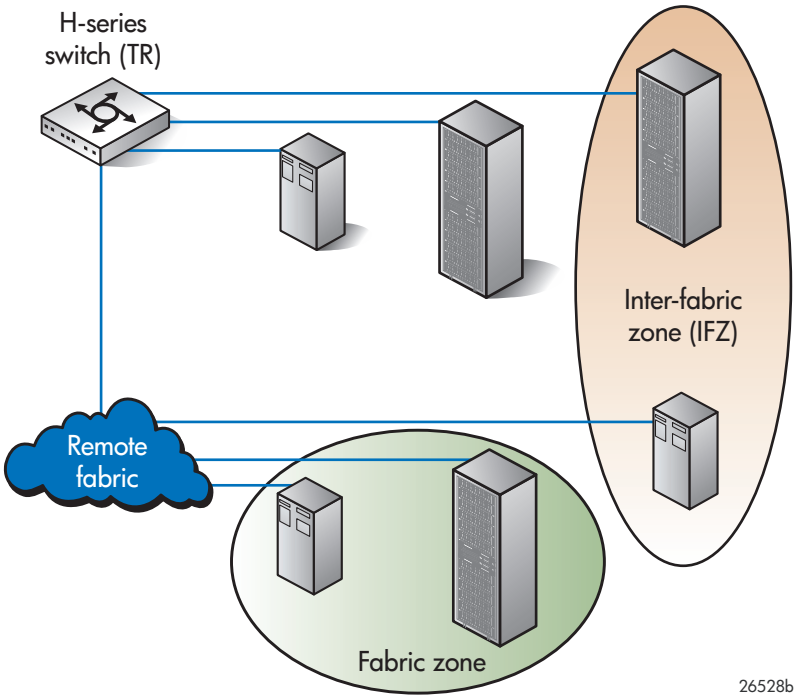



**Figure 21: H-series switch routing**

# Fabric redundancy and routing

B-series 8 Gb/s switches with integrated Fibre Channel routing, StoreFabric SN4000B SAN Extension Switch, 1606 Extension SAN Switches or DC Dir Switch MP Extension Blades, or MP Router Blade 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 (**Dual-redundant Meta SAN with HPE StoreFabric SN4000B SAN Extension Switch, 1606 Extension SAN Switch, or MP Router Blade**), or two groups of Virtual Fabrics or VSANs with IFR (**Figure 23: Dual-redundant Virtual Fabric** on page 56) or IVR (**Figure 24: Dual-redundant VSAN** on page 57).

For more information, see **High-availability router configurations** on page 59.



https://t.me/learningnets

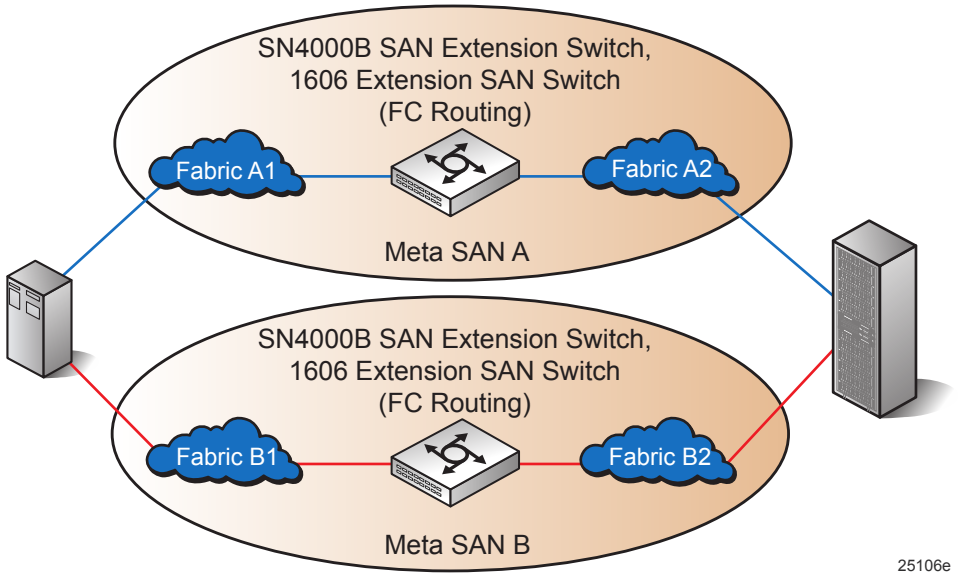



**Figure 22: Dual-redundant Meta SAN**

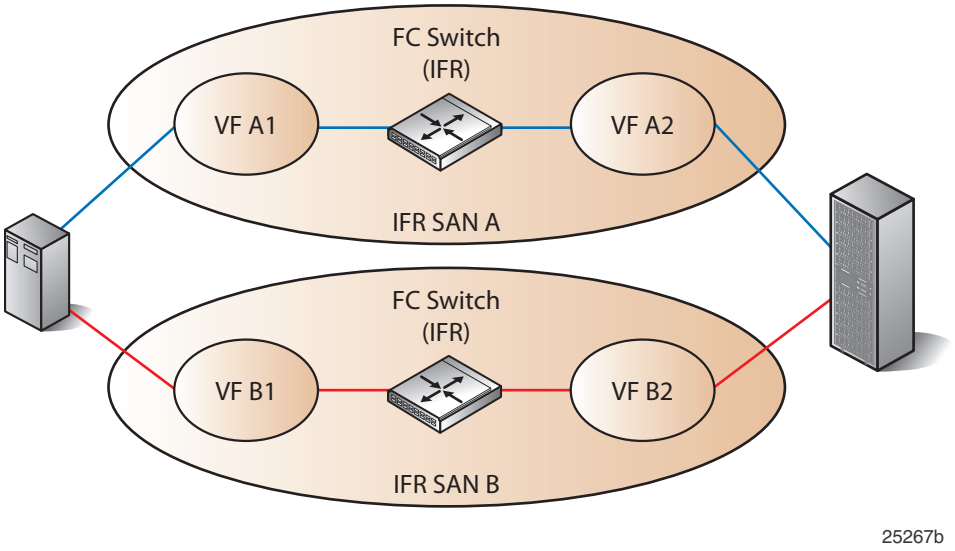



**Figure 23: Dual-redundant Virtual Fabric**



https://t.me/learningnets

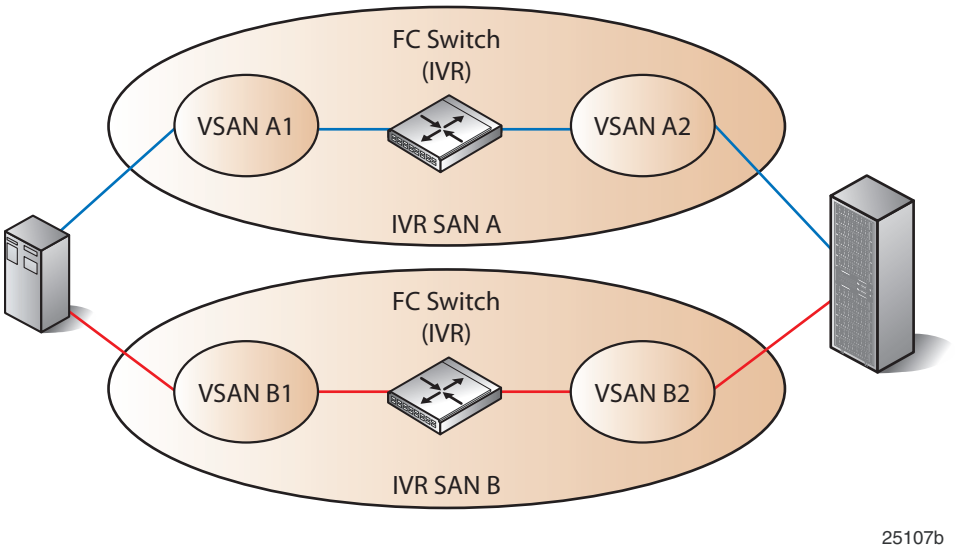



**Figure 24: Dual-redundant VSAN**

# Supported routing configurations

Routing requires additional configuration rules for fabrics. For more information about routing configuration rules, see **C-series switches and fabric rules** on page 137 and **H-series switches and fabric rules** on page 151.

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 **HPE StoreFabric SN4000B SAN Extension Switch or 1606 Extension SAN Switch connecting core switches**. A MP Router Blade 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:

- B-series Fibre Channel routing (**Figure 25: Connecting core switches** on page 57)
- B-series Virtual Fabrics with IFR (**Figure 26: Virtual Fabrics connecting core switches** on page 58)
- C-series VSANs with IVR (**Figure 27: VSANs connecting core switches** on page 58)

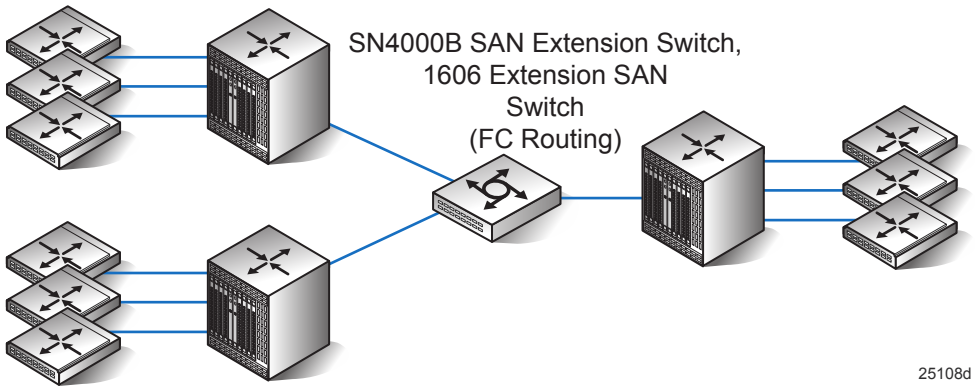



**Figure 25: Connecting core switches**



https://t.me/learningnets

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.

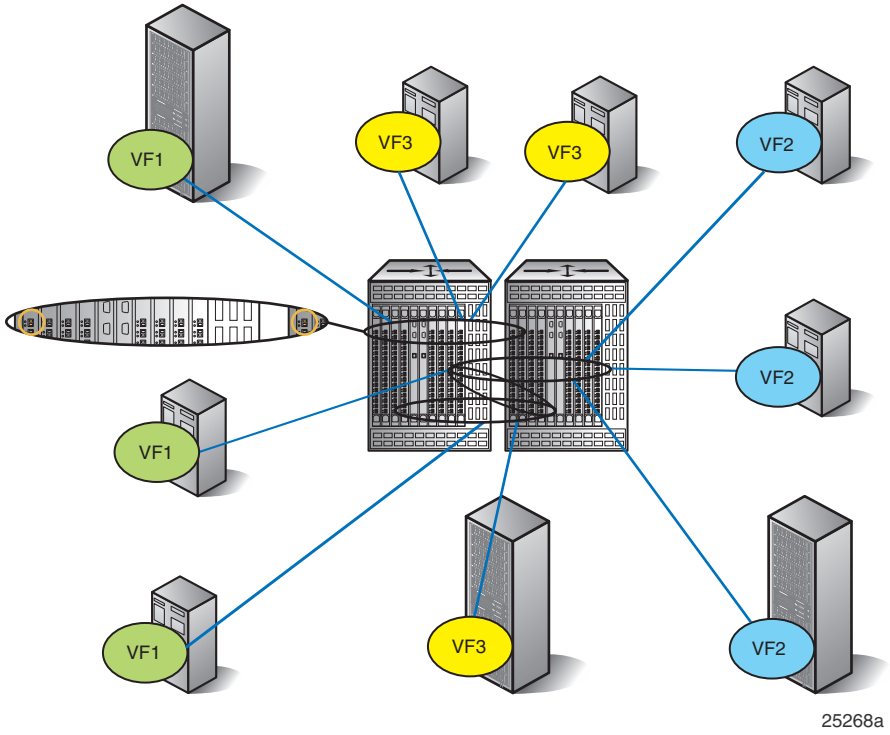



**Figure 26: Virtual Fabrics connecting core switches**

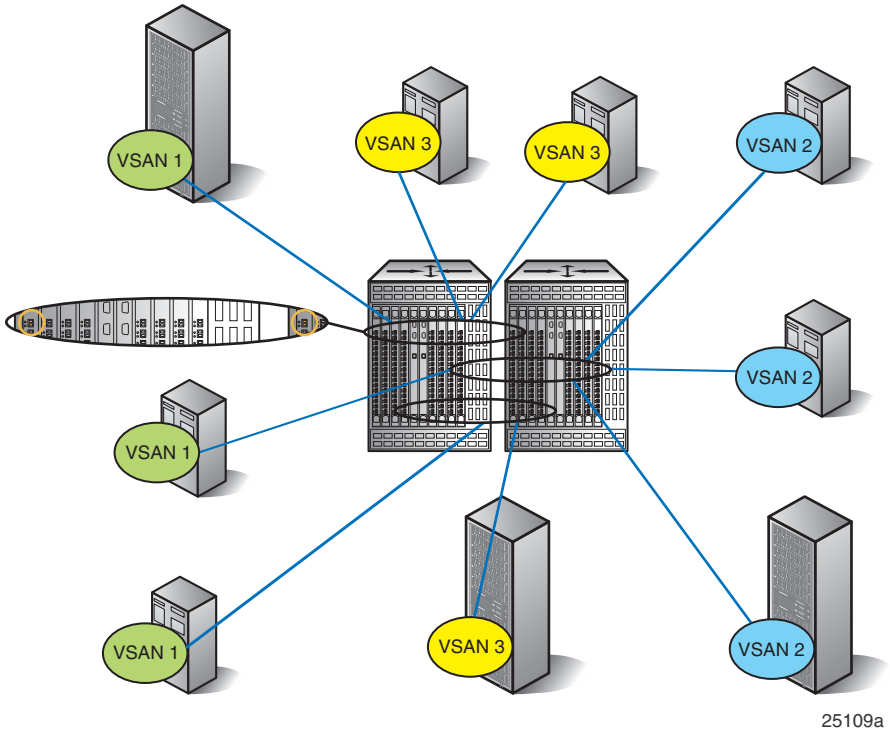



**Figure 27: VSANs connecting core switches**



https://t.me/learningnets

## Routing through an IP network

When connecting fabrics through IP, StoreFabric SN4000B SAN Extension Switch, 1606 Extension SAN Switches or DC Dir Switch MP Extension Blades, or MP Router Blade serve as FCIP gateways with Fibre Channel routing. Routers that communicate with the FCIP protocol must be installed in pairs (**Figure 28: HPE StoreFabric SN4000B SAN Extension Switch, 1606 Extension SAN Switches connecting fabrics through IP** on page 59).

Hewlett Packard Enterprise supports FCIP configurations in which StoreFabric SN4000B SAN Extension Switch, 1606 Extension SAN Switches or DC Dir Switch MP Extension Blades, 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 StoreFabric SN4000B SAN Extension Switch, 1606 Extension SAN Switches or DC Dir Switch MP Extension Blades (**Figure 28: HPE StoreFabric SN4000B SAN Extension Switch, 1606 Extension SAN Switches connecting fabrics through IP** on page 59), or MP Router Blade.

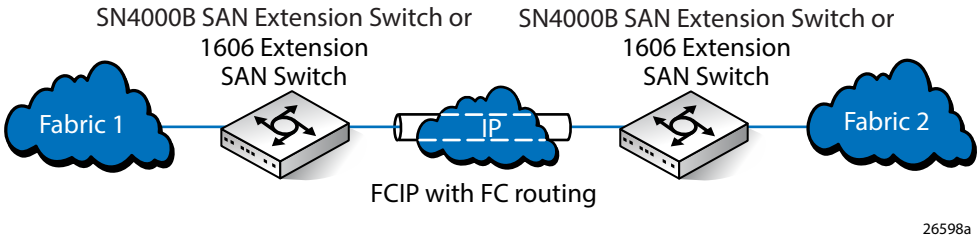



**Figure 28: HPE StoreFabric SN4000B SAN Extension Switch, 1606 Extension SAN Switches connecting fabrics through IP**

VSANs can be connected through IP using the FCIP functionality of the C-series Fibre Channel switches, see **C-series Fibre Channel switches** on page 137.

**Figure 29: IVR connecting VSANs through IP** on page 59 shows IVR connecting VSANs through IP.

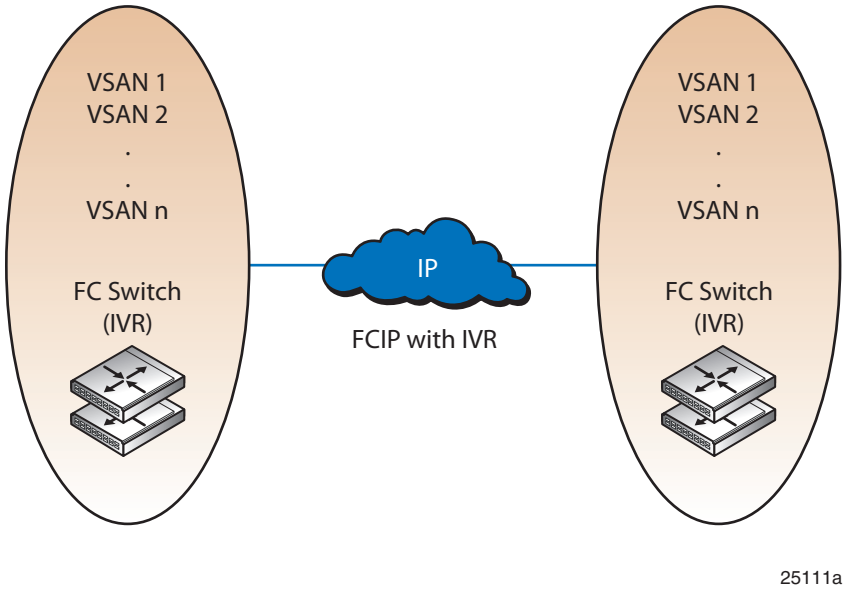



**Figure 29: IVR connecting VSANs through IP**

## High-availability router configurations

In high-availability configurations, use pairs of 8 Gb/s switches with integrated Fibre Channel routing, StoreFabric SN4000B SAN Extension Switch, 1606 Extension SAN Switches, routers, or IFL pairs to



https://t.me/learningnets

provide redundant paths between fabrics. **Figure 30: High-availability configurations** on page 60 shows valid configuration.

- 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** on page 40.
- 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 StoreFabric SN4000B SAN Extension Switch or 1606 Extension SAN Switch 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** on page 105.

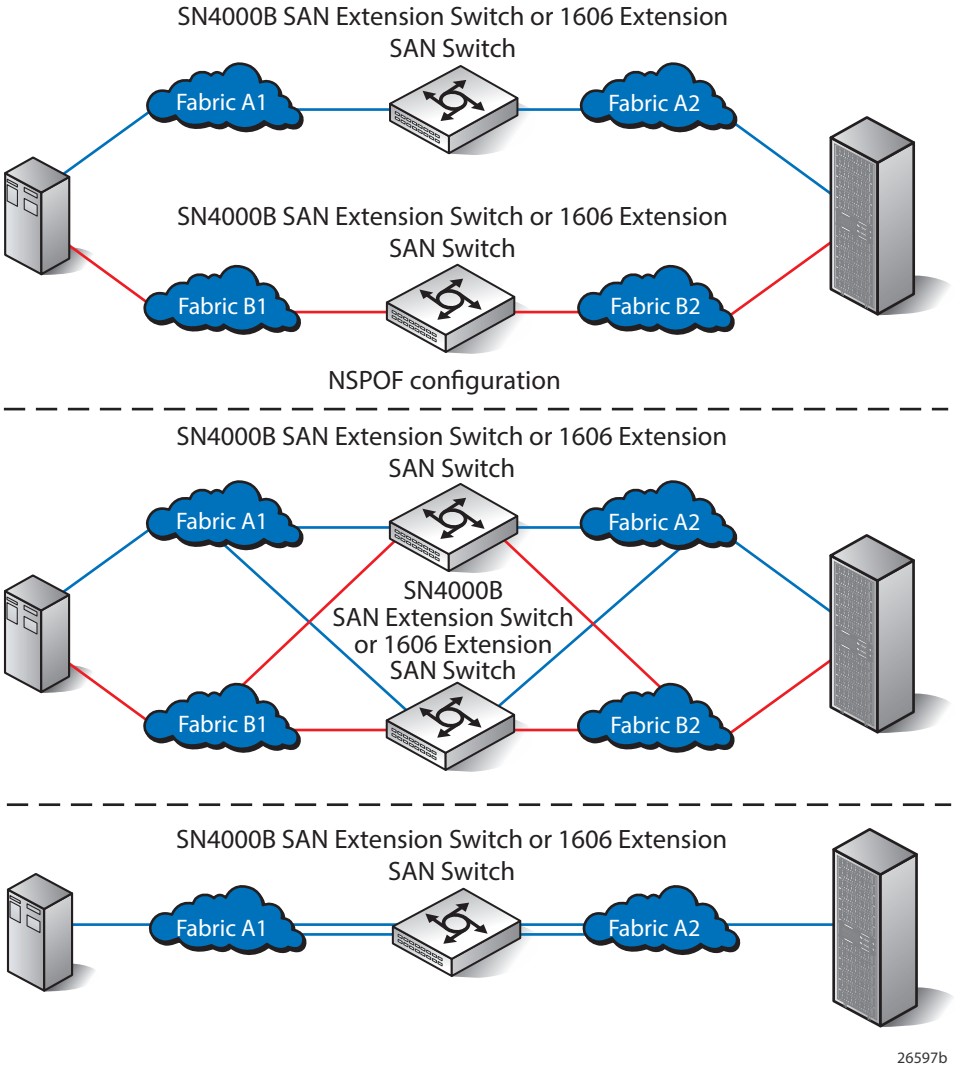



**Figure 30: High-availability configurations**



https://t.me/learningnets

## 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: Basic TR configuration** on page 51 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: H-series switch routing** on page 55 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, StoreFabric SN4000B SAN Extension Switch, 1606 Extension SAN Switches or DC Dir Switch MP Extension Blades, or MP Router Blade enables sharing of devices located anywhere in the Meta SAN.
- Ability to connect B-series fabrics using Secure Fabric OS.

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. Hewlett Packard Enterprise recommends the use of Fabric Manager to simplify management procedures when implementing a Meta SAN.



https://t.me/learningnets

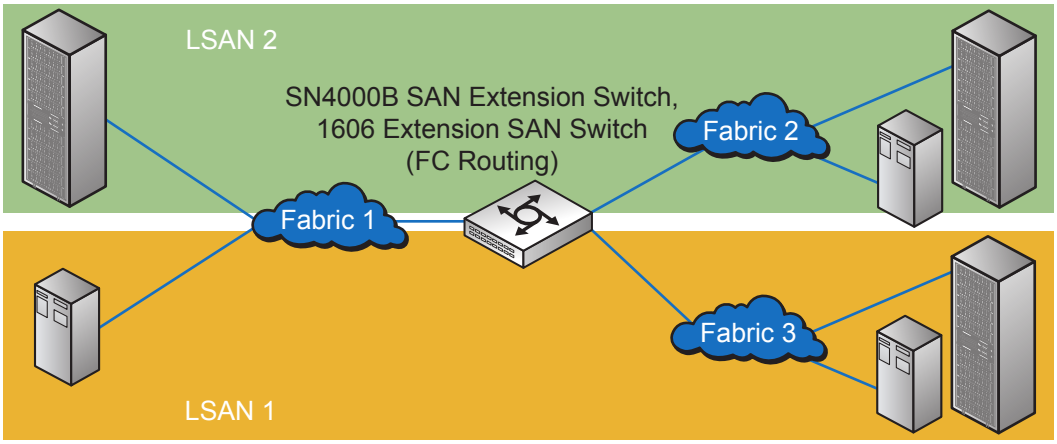



**Figure 31: Consolidating SAN islands**

### Integration of Fibre Channel routing and FCIP

You can use the StoreFabric SN4000B SAN Extension Switch, HPE 1606 Extension SAN Switch, or MP Router Blades integrated FCIP capability to extend disaster-tolerant applications such as HPE 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 32: NSPOF configuration with HPE StoreFabric SN4000B SAN Extension Switch or 1606 Extension SAN Switch providing Fibre Channel routing and FCIP with direct connect devices** on page 63 and **Figure 33: NSPOF configuration with HPE StoreFabric SN4000B SAN Extension Switch or 1606 Extension SAN Switch providing FC routing and FCIP with direct connect devices** on page 63 shows typical HPE Continuous Access NSPOF configurations in which the StoreFabric SN4000B SAN Extension Switch or 1606 Extension SAN Switch provides Fibre Channel routing and FCIP. **Figure 33: NSPOF configuration with HPE StoreFabric SN4000B SAN Extension Switch or 1606 Extension SAN Switch providing FC routing and FCIP with direct connect devices** on page 63 shows a configuration in which the StoreFabric SN4000B SAN Extension Switch or 1606 Extension SAN Switches are used as both an FCIP gateway and Fibre Channel switch.

**NOTE:** The StoreFabric SN4000B SAN Extension Switches need to be installed in pairs for FCIP. They cannot be used for FCIP tunnels with any other platform.



https://t.me/learningnets

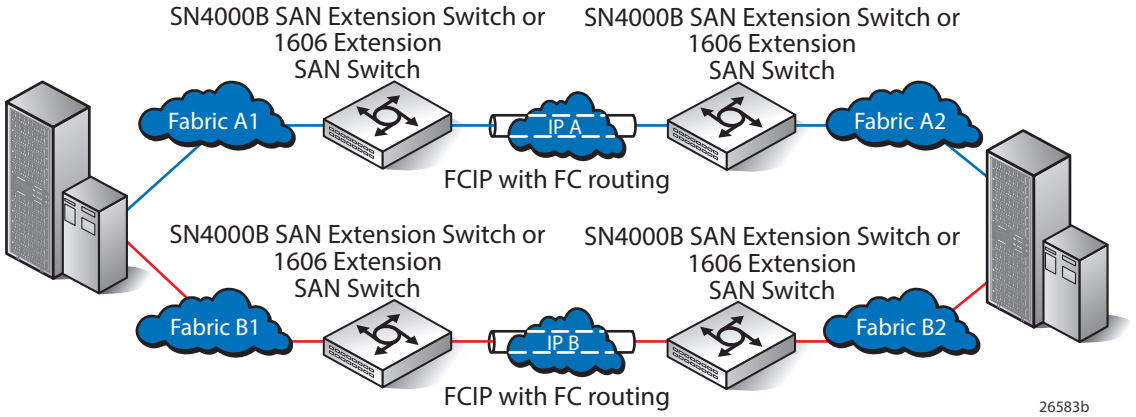



**Figure 32: NSPOF configuration with HPE StoreFabric SN4000B SAN Extension Switch or 1606 Extension SAN Switch providing Fibre Channel routing and FCIP with direct connect devices**

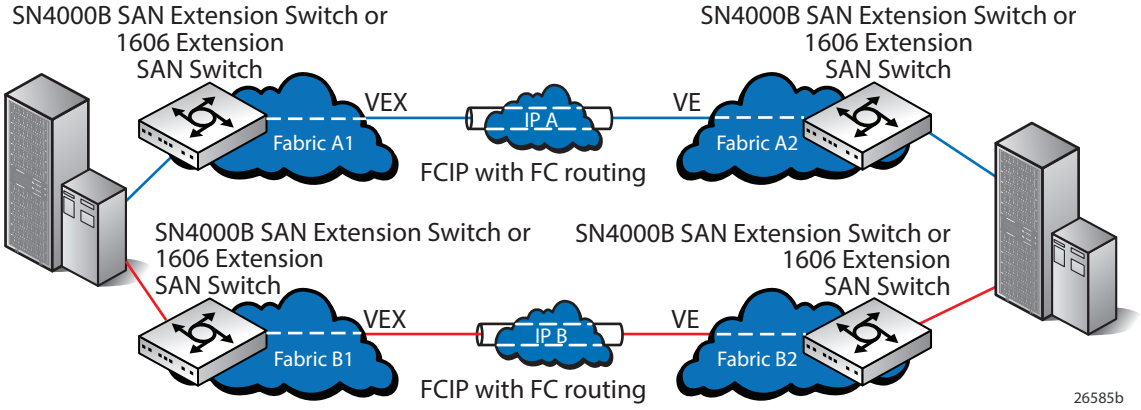



**Figure 33: NSPOF configuration with HPE StoreFabric SN4000B SAN Extension Switch or 1606 Extension SAN Switch providing FC routing and FCIP with direct connect devices**

### Tape backup consolidation

The StoreFabric SN4000B SAN Extension Switch or 1606 SAN Extension Switch 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.

**HPE StoreFabric SN4000B SAN Extension Switch or 1606 Extension SAN Switch consolidating tape backup in a Meta SAN** shows a configuration in which a StoreFabric SN4000B SAN Extension Switch or 1606 SAN Extension Switch consolidates tape backup in a Meta SAN. For information about supported HPE tape products and backup applications, see the *HP Enterprise Backup Solution Design Guide* and the *EBS compatibility matrix*. These documents are available at **http://www.hpe.com/info/ebs**.



https://t.me/learningnets

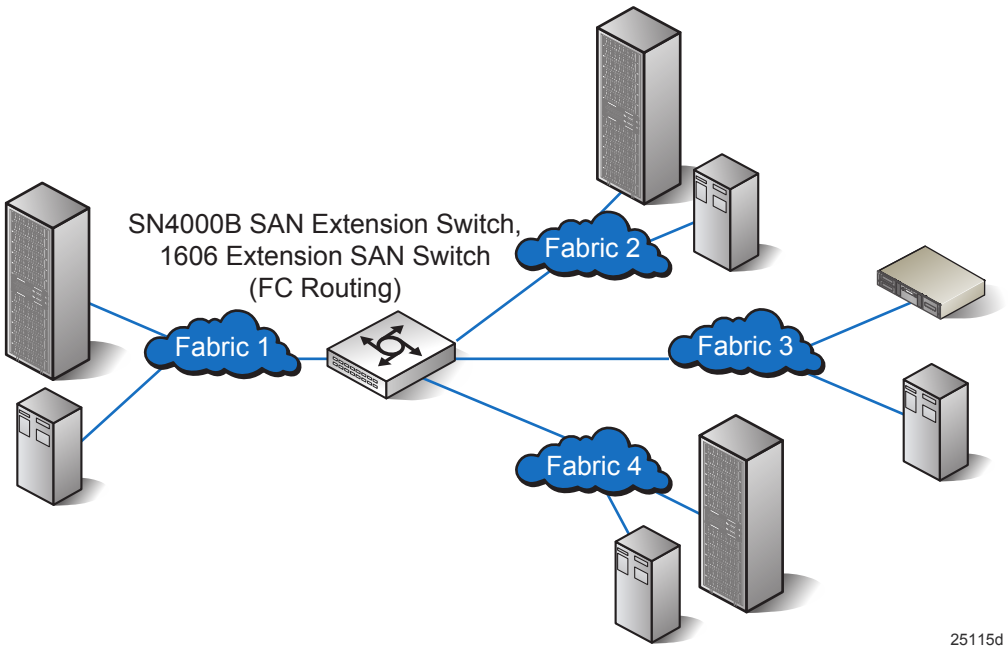



**Figure 34: Consolidating tape backup in a Meta SAN**

Independent fabrics connected through a StoreFabric SN4000B SAN Extension Switch or 1606 Extension SAN Switch cannot have a direct ISL connection (**Unsupported HPE StoreFabric SN4000B SAN Extension Switch or 1606 Extension SAN Switch configuration**). A direct ISL connection between fabrics bypasses the FC routing function of the switch, resulting in a full fabric merge.

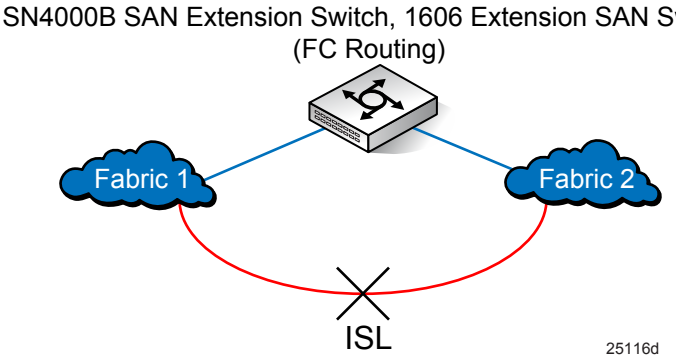



**Figure 35: Unsupported Switch configuration**



https://t.me/learningnets

# 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. Simultaneously, the 10-GbE network supports Ethernet data thus creating a single converged network for Ethernet and Fibre Channel.

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:


- **HPE FCoE solutions overview** on page 65
- **HPE FCoE converged switch technology** on page 67
- **HPE FCoE products** on page 70
- **FCoE configuration rules** on page 91


## HPE FCoE solutions overview

Hewlett Packard Enterprise offers multiple FCoE solutions:

- HPE 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 HPE Fibre Channel SANs.
- HPE end-to-end FCoE solutions for 3PAR StoreServ, XP7, P9500, and P6000/EVA storage provide an FCoE technology implementation from server to storage. The solutions consist of FCoE CNAs, CN switches, and FCoE storage targets. The following FCoE end-to-end configurations are available:
  - **HPE 3PAR StoreServ 7200/7400/7450/8200/8400/8440/8450 AFA/ 10400/10800/20450/20800/20850 FCoE/iSCSI target** — Provides 10GbE FCoE and iSCSI ports. The number of available ports varies depending on the 3PAR StoreServ model. See the **HPE Native iSCSI products** on page 346 and the *HPE 3PAR Feature Availability Matrix* located on the HPE SPOCK website at **http://www.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access.
  - **XP7 FCoE target** — Provides eight 10 GbE FCoE ports on each channel adapter. See **FCoE storage systems** on page 85 and the *HPE XP7 Owner Guide*.
  - **P9500 FCoE target** — Provides four 10 GbE FCoE ports on each channel adapter. See **FCoE storage systems** on page 85 and the *HPE P9500 Owner Guide*.
  - **P63xx/P65xx FCoE/iSCSI target** — Provides two 10 GbE FCoE/iSCSI ports on each storage controller. See **Figure 39: P63xx/P65xx FCoE/iSCSI end-to-end configuration** on page 89, **P6300/P6350/P6500/P6550 EVA** on page 348, and the iSCSI or iSCSI/FCoE configuration rules and guidelines chapter in the *HPE P63x0/P65x0 Enterprise Virtual Array User Guide*.



https://t.me/learningnets

- **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, 3PAR StoreServ, or XP24000/12000 (iSCSI only) Fibre Channel storage systems using a fabric connection between the MPX200 and the Fibre Channel storage.

The FCoE module and switch product set consists of the following products:

- HPE c-Class BladeSystem Virtual Connect FCoE capable modules and switches, all supporting both dual hop and multi-hop FCoE configurations:
  - **HPE Virtual Connect 20/40 F8 on page 182**
  - **HPE Virtual Connect FlexFabric 10 Gb/24-Port Module for c-Class BladeSystem on page 182**
  - **HPE Virtual Connect Flex-10/10D Ethernet Module for c-Class BladeSystem on page 183**
  - **HPE 6127XLG and 6125XLG Ethernet Blade Switches** — Provides a converged fabric solution that supports Ethernet, native FCoE, and iSCSI protocols that enables connectivity for FCoE or iSCSI storage solutions.

    The 6127XLG has 16 internal 1/10/20 Gb downlinks, the 6125XLG has 16 internal 1/10 Gb downlinks. Both switches have 4 external 40 Gb ports and 8 external 10Gb ports.

- **HPE FlexFabric 5900CP-48XG-4QSFP+ Switch, HPE FlexFabric 5930 4-slot and 2-slot Modular Switches, HPE FlexFabric 5940 4-slot Switch and 2-slot Modular Switches, HPE FlexFabric 5950 4-slot Switch** — Provides a bridge function from a converged network to separate Ethernet and FC fabrics, or support end-to-end FCoE when used with CNAs and FCoE storage targets, and end-to-end FC when used with HBAs and FC storage targets.
- **HPE FlexFabric 5950 48SFP28 8QSFP28 Switch, HPE 5900AF-48XG-4QSFP+ Switch, HPE FlexFabric 5930 32QSFP+ Switch, and HPE FlexFabric 5700 Switch** — Provide Ethernet and end-to-end FCoE when used with CNAs and FCoE storage targets.
- **HPE 5820X-14XG-SFP+ Switch** — Provides a bridge function from a converged network to separate Ethernet and FC fabrics.
- **HPE 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 FCoE storage targets.
- **HPE StoreFabric CN1200E, CN1100E, CN1100R, CN1000E, CN1000Q CNAs, and 650, 630, 556, 554, NC553, NC551, 536, 534 and 526 FlexFabric Adapters** — Converge Ethernet and FC technologies in the server over 10-GbE links to converged network switches.
- **HPE StoreFabric CN1200E-T, CN1100R-T, 556FLR-T and 553FLR-T 10GBASE-T CNAs and FFAs**
- **HPE 3PAR StoreServ, XP7, P9500, and P63xx/P65xx EVA storage systems with FCoE** — Provides native FCoE target connectivity within the storage system.
- **HPE MPX200 Multifunction Router iSCSI/FCoE** — Provides FCoE target functionality when integrated with P6000/EVA or 3PAR StoreServ Fibre Channel storage systems. For more information, see the *HPE MPX200 Multifunction Router User Guide*.



https://t.me/learningnets

# HPE FCoE converged switch technology

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

FCoE CN switches have multiple types of physical ports:

- 10-GbE ports
- 40-GbE (FlexFabric 5900CP, 5900AF)
- 4-Gb or 8-Gb FC ports
- Converged ports: 10-GbE, 2-Gb, 4-Gb, or 8Gb FC ports (FlexFabric 5900CP)

Converged ports on the HPE FlexFabric 5900CP switch can be used for Ethernet/FCoE or FC device connectivity, including NPV/NPIV. 40GbE ports on the FlexFabric 5900CP and 5900AF switches can be used for Ethernet/FCoE connections. Fibre Channel ports on the HPE 5820 CN switch can be used for NPV/NPIV connections.

Fibre Channel ports on the B-series FCoE CN switches and HPE C-series CN switches can be used for ISL connections to an existing Fibre Channel fabric, or to a Fibre Channel storage system or HBA.

### Table 9: Number of ports per switch

| HPE switch | Number of 10-GbE ports | Number of FC ports |
|---|---|---|
| HPE FlexFabric 5900CP Switch | 0 to 48 (10GbE/FCoE or 2Gb/s, 4Gb/s, 8Gb/s FC)<br>0 to 4 (40GbE/FCoE or 16x10GbE/FCoE) | |
| HPE 5900AF Switch | 0 to 48 (10GbE/FCoE)<br>0 to 4 (40GbE/FCoE or 16x10GbE/FCoE) | None |

*Table Continued*



https://t.me/learningnets

| HPE switch | Number of 10-GbE ports | Number of FC ports |
|---|---|---|
| HPE FlexFabric 5950 4-slot Switch<br>Hot swappable Modules available:<br>• HPE 5950 16-port QSFP+ Module JH405A<br>• HPE 5950 8-port QSFP28 Module JH406A<br>• HPE FlexFabric 5950 24-port SFP28 and 2-port QSFP28 Module JH450A<br>• HPE 5930 24-port SFP+ and 2-port QSFP+ Module JH180A<br>• HPE 5930 24-port SFP+ and 2-port QSFP+ with MACsec Module JH181A<br>• HPE 5930 24-port Converged Port and 2-port QSFP+ Module JH184A<br>• HPE 5930 8-port QSFP+ Module JH183A | 10 GbE, 25GbE, 40GbE, 100Gb and FC ports available are based on module types utilized in the switches. For additional information, see the HPE FlexFabric Switch Series Storage Solutions Configuration Guide and the SPOCK website at **http://www.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access. | |
| HPE FlexFabric 5950 48SFP28 8QSFP28 Switch JH402A | 48 - 1GbE/10GbE/25GbE<br>8 - 40GbE/100GbE | None |

*Table Continued*



https://t.me/learningnets

| HPE switch | Number of 10-GbE ports | Number of FC ports |
|---|---|---|
| HPE FlexFabric 5940 4-slot Switch<br>HPE FlexFabric 5940 2-slot Switch<br>Hot swappable Modules available:<br>• HPE 5930 24-port SFP+ and 2-port QSFP+ Module JH180A<br>• HPE 5930 24-port SFP+ and 2-port QSFP+ with MACsec Module JH181A<br>• HPE 5930 24-port 10GBASE-T and 2-port QSFP+ with MACsec Module JH182A<br>• HPE 5930 24-port Converged Port and 2-port QSFP+ Module JH184A<br>• HPE 5930 8-port QSFP+ Module JH183A<br>• HPE 5940 2-port QSFP+ and 2-port QSFP28 Module JH409A<br>• HPE FlexNetwork 5930 24-port 10GbE SFP/SFP+ and 2-port 40GbE QSFP+ Module JH689A<br>• HPE FlexNetwork 5930 24-port 10GBASE-T and 2-port 40GbE QSFP+ MACsec Module JH690A | 10 GbE, 40GbE, and FC ports available are based on module types utilized in the switches. For additional information, see the HPE FlexFabric Switch Series Storage Solutions Configuration Guide and the SPOCK website at **http://www.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access. | |
| HPE FlexFabric 5930 4-slot Switch<br>HPE FlexFabric 5930 2-slot Switch<br>Hot swappable Modules available:<br>• HPE 5930 24-port SFP+ and 2-port QSFP+ Module JH180A<br>• HPE 5930 24-port SFP+ and 2-port QSFP+ with MACsec Module JH181A<br>• HPE 5930 24-port 10GBASE-T and 2-port QSFP+ with MACsec Module JH182A<br>• HPE 5930 24-port Converged Port and 2-port QSFP+ Module JH184A<br>• HPE 5930 8-port QSFP+ Module JH183A | 10 GbE and FC ports available are based on module types utilized in the switches. For additional information, see the *HPE FlexFabric Switch Series Storage Solutions Configuration Guide* and the SPOCK website at **http://www.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access. | |
| HPE 5820 Converged Network Switch | 14 to 18 | 4 to 8 (8 Gb/s) |

*Table Continued*



https://t.me/learningnets

| HPE switch | Number of 10-GbE ports | Number of FC ports |
|---|---|---|
| HPE FlexFabric 5700 Switch | 42 or 54 (model dependent)<br>1 to 4 (40GbE/FCoE or 8x 10GbE/FCoE) | None |
| DC SAN Director Switch 10/24 FCoE Blade | 24 | 0 external ports<br>32 (8 Gb/s) backplane ports |
| Cisco Nexus 5548UP Converged Network Switch | 0 to 48 with expansion module | 0 to 48 with expansion module |
| Cisco Nexus 5596UP Converged Network Switch | 0 to 96 with expansion modules | 0 to 96 with expansion modules |

**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. Hewlett Packard Enterprise recommends that you implement a dual path configuration with redundant FCoE CN switches to minimize disruption during firmware updates.

### CN Switch Interoperability

The FlexFabric 5900CP Switch, 5950 Switches, 5940 Switches, 5930 Switch, and 5820 Converged Network Switch interoperate with certain other series HPE FC switches using NPV/NPIV uplinks.

**NOTE:** For switch series and switch model NPV/NPIV support with the FlexFabric 5900CP switch, see the *HPE FlexFabric Switch Series Storage Solutions Configuration Guide* and the SPOCK website at **http://www.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access.

Hewlett Packard Enterprise does not support E_Port interoperability between FlexFabric 5900CP, B-series, and C-series CN and FC switches:

- 10/24 FCoE Blade operates with B-series FC switches only.
- Cisco Nexus 5548UP/5596UP Converged Network Switches operate with other C-series MDS FC switches only.

# HPE FCoE products

This section describes the following FCoE CN switches and CNAs.

- **HPE FlexFabric 5900CP Switch**
- **HPE FlexFabric 5930 32QSFP+ Switch**
- **HPE FlexFabric 5930 4-slot Modular Switch and HPE FlexFabric 5930 2QSFP+ 2-slot Modular Switches**


https://t.me/learningnets

- **HPE FlexFabric 5950 4-slot Modular Switch and HPE FlexFabric 5950 48SFP28 8QSFP28 Switch**
- **HPE FlexFabric 5940 4-slot Modular Switch and HPE FlexFabric 5940 2-slot Modular Switch**
- **HPE 5820 Ethernet/FCoE Switch**
- **HPE FlexFabric 5700 Switch**
- **HPE StorageWorks DC SAN Director Switch 10/24 FCoE Blade**
- **HPE StoreFabric SN8600B 64/32 Short Wave 8QSFP Integrated Fibre Channel Blade**
- **Cisco Nexus 5548UP/5596UP Converged Network Switches**
- **FCoE storage systems**

# Converged network switches and blades

## HPE FlexFabric 5900CP Switch

The FlexFabric 5900CP Switch is a ToR converged L2/L3 switch with 48 converged ports and 4-40GbE QSFP+ ports. This switch provides seamless integration of 10GbE networks into existing server edge environments (FCoE DCB, Fibre Channel) and the HPE Server/Storage Ecosystem. The switch is full featured and no software licensing is required. Switches can be added in a 5900 series IRF (Intelligent Resilient Framework) fabric. An IRF fabric appears as one node and is accessible at a single IP address on the network. You can use this IP address to log in at any member device to manage all the members of the IRF fabric.

**NOTE:** For more information about IRF usage in a storage configuration, see the *HPE FlexFabric Switch Series Storage Solutions Configuration Guide*.

**Features and benefits:**

- Ethernet/FCoE support on 10GbE converged ports and 40GbE ports
- Fibre Channel 8Gb/4Gb/2Gb support on converged ports
- DCB/FCoE-FCF/NPV/TRILL/IPv6 support (QCN ready)
- Forty-eight converged ports that can use a single transceiver which supports 10GbE/FCoE or 8Gb/4Gb Fibre Channel
- FCoE/FC NPV gateway support
- Dual-hop support with Virtual Connect blade switches and the HPE 6125-XLG
- Multi-hop support using FCoE VE_Port (7 hops) or Fibre Channel E_Port (3 hops) ISLs
- 5900 Series IRF support (up to nine switches with Ethernet, two switches with storage)
- iSCSI support
- Front-to-back or back-to-front airflow
- Comware OS version 7.1
- No additional feature licenses. The following four switch modes are available:
  - Standard — Configurable DCB switch
  - FCF — FC/FCoE initiator; target; FlexFabric SAN switch; F, VF, VE, and E port connectivity (4,000 zones); FSPF



https://t.me/learningnets

- NPV — Gateway for FC/FCoE multivendor connectivity
- Transit — FIP-snooping DCB aggregation switch

- 1.28 Tb/s switching capacity
- 952.32 million PPS throughput, integrated 9 MB packet buffer
- 10GbE cut-through latency < 1.5 µs (64-byte packets)
- CLI and iMC/VAN/VFM (Intelligent Management Center) fabric management
- HPE IIAS (Intelligent Infrastructure Analyzer Software) support
- L2/L3, IPv4/IPv6 dual stack, TRILL, VEPA
- Cloud and SDN ready (OpenFlow 1.3.1 support)

## HPE FlexFabric 5930 32QSFP+ Switch

The HPE 5930 32QSFP+ switch has 32 QSFP+ ports, dual hot-pluggable power supplies and fan trays, and ultra-low-latency. It is ideally suited for deployment at the spine and server access layer in large enterprise data centers.

**Features and benefits:**

The FlexFabric 5930 32QSFP+ switch is available in a 1RU 32-port 40GbE QSFP+ form factor. 40GbE ports 5 through 28 may be split into four 10GbE ports each for a total of 96 10GbE ports with eight 40GbE Uplinks per switch.

The FlexFabric 5930 32QSFP+ switch provides the following features:

- FCoE support on 10GbE and 40GbE ports
- DCB/QCN
- Delivers up to 2.56 Tb/s switching capacity for the most demanding applications
- Supports up to 1492 MPPS throughput for data-intensive environments
- Low latency, under a 1.5µs 10GbE latency, gives your business agility
- iSCSI support; optional lossless support
- VXLAN support for network virtualization and overlay solutions
- Delivers IRF <50 msec convergence time enabling faster application response time
- The In Service Software Update (ISSU) enables high availability with updates accomplished without a reboot or power cycle, in the background
- Simplifies switch management by up to 88% with 9 unit Intelligent Resilient Framework (IRF) (Ethernet only)
- All switch ports are active and ready to use without need for activation licenses
- Automate tedious tasks with a Software-defined Network (SDN) and reclaim wasted resources


https://t.me/learningnets

## HPE FlexFabric 5930 4-slot Modular Switch and HPE FlexFabric 5930 2QSFP+ 2-slot Modular Switches

The HPE FlexFabric 5930 4-slot switch has four module slots, dual hot-pluggable power supplies and fan trays, and ultra-low-latency. Ideally suited for deployment at spine and server access layers in large enterprise data centers.

The HPE FlexFabric 5930 2QSFP+ 2-slot switch has two module slots, 2 QSFP+ ports, dual hotpluggable power supplies and fan trays, and ultra-low-latency. Ideally suited for deployment at spine and server access layers in large enterprise data centers. For additional information about HPE FlexFabric 5930 Modular switches, see **http://www.hpe.com/networking/5930**.

- Ports:
  - (4) module slots
  - (2) module slots
  - (2) QSFP+ 40GbE ports
  - (2) QSFP+ 40GbE ports on chassis; each QSFP+ port operable as 4x10GbE.
- Hot swappable Modules available for 4-slot and 2-slot 5930:
  - HPE 5930 24-port SFP+ and 2-port QSFP+ Module JH180A
  - HPE 5930 24-port SFP+ and 2-port QSFP+ with MACsec Module JH181A
  - HPE 5930 24-port SFP+ and 2-port QSFP+ with MACsec Module JH181A
  - HPE 5930 24-port 10GBASE-T and 2-port QSFP+ with MACsec Module JH182A
  - HPE 5930 24-port Converged Port and 2-port QSFP+ Module JH184A
  - HPE 5930 8-port QSFP+ Module JH183A
- Latency:
  - 10 Gbps Latency: < 1 µs
- Throughput:
  - Up to 1429 Mbps
- Routing/switching capacity:
  - 2560 Gbps
  - Fibre Channel 8Gb/4Gb/2Gb support on any of the 24 converged ports of the JH184A module
  - FCoE support on 10GbE and 40GbE ports
- Stacking capabilities:
  - IRF with 9 switches (Ethernet), 2 switches (storage)
  - CLI, OneView, and Intelligent Management Center (iMC/VAN/VFM) fabric management



https://t.me/learningnets

## HPE FlexFabric 5950 4-slot Modular Switch and HPE FlexFabric 5950 48SFP28 8QSFP28 Switch

The HPE FlexFabric 5950 Switch Series is a family of high-density, ultra-low-latency, top-of-rack (ToR) switches that is part of the Hewlett Packard Enterprise (HPE) FlexNetwork architecture's HPE FlexFabric solution. These switches are ideally suited for deployment at the aggregation or server access layer of large enterprise data centers, or at the core layer of medium-sized enterprises.

The HPE FlexFabric 5950 4-slot switch has four module slots, four hot-pluggable power supply slots and two fan tray slots. The HPE FlexFabric 5950 48SFP28 8QSFP28 Switch includes 48 SFP28 I/O ports, 8 QSFP28 I/O ports, and 2 SFP I/O ports, two power supply slots and six fan tray slots.

For additional information about HPE FlexFabric 5950 Switches, see **http://www.hpe.com/networking/flexfabric**.

- Ports:

  HPE FlexFabric 5950 4-slot Switch JH404A

  - (4) module slots

  HPE FlexFabric 5950 48SFP28 8QSFP28 Switch JH402A

  - 48 SFP28 I/O ports
  - 8 QSFP28 I/O ports
  - 2 SFP I/O ports

- Supported Hot swappable Modules available for 4-slot 5950:

  - HPE 5950 16-port QSFP+ Module JH405A
  - HPE 5950 8-port QSFP28 Module JH406A
  - HPE FlexFabric 5950 24-port SFP28 and 2-port QSFP28 Module JH450A
  - HPE 5930 24-port SFP+ and 2-port QSFP+ Module JH180A
  - HPE 5930 24-port SFP+ and 2-port QSFP+ with MACsec Module JH181A
  - HPE 5930 8-port QSFP+ Module JH183A
  - HPE 5930 24-port Converged Port and 2-port QSFP+ Module JH184A

- Latency:

  - 10 Gbps Latency: < 1 µs

- Throughput:

  - Up to 3169 Mpps

- Routing/switching capacity:



https://t.me/learningnets

- 3200 Gbps
- Fibre Channel 8Gb/4Gb/2Gb support on any of the 24 converged ports of the JH184A module
- FCoE support

- Stacking capabilities:
  - IRF with up to ten switches (Ethernet), two switches (storage)
  - CLI, and Intelligent Management Center (iMC/VAN/VFM) fabric management

## HPE FlexFabric 5940 4-slot Modular Switch and HPE FlexFabric 5940 2-slot Modular Switch

The HPE FlexFabric 5940 4-slot switch has four module slots, dual hot-pluggable power supplies and fan trays, and ultra-low-latency. Ideally suited for deployment at spine and server access layers in large enterprise data centers, or at the core layer of medium-sized enterprises.

The HPE FlexFabric 5940 2-slot switch has two module slots, 2 QSFP+ ports, dual hotpluggable power supplies and fan trays, and ultra-low-latency. Ideally suited for deployment at spine and server access layers in large enterprise data centers.

For additional information about HPE FlexFabric 5940 Modular switches, see **www.hpe.com/networking/FlexFabric**.

- Ports:

  HPE FlexFabric 5940 4-slot Switch JH398A
  - (4) module slots

  HPE FlexFabric 5940 2-slot Switch JH397A
  - (2) module slots
  - (2) QSFP+ 40GbE ports

- Supported Hot swappable Modules available for 4-slot and 2-slot 5940:
  - HPE 5930 24-port SFP+ and 2-port QSFP+ Module JH180A
  - HPE 5930 24-port SFP+ and 2-port QSFP+ with MACsec Module JH181A
  - HPE 5930 24-port 10GBASE-T and 2-port QSFP+ with MACsec Module JH182A
  - HPE 5930 8-port QSFP+ Module JH183A
  - HPE 5930 24-port Converged Port and 2-port QSFP+ Module JH184A
  - HPE 5940 2-port QSFP+ and 2-port QSFP28 Module JH409A
  - HPE FlexNetwork 5930 24-port 10GbE SFP/SFP+ and 2-port 40GbE QSFP+ Module JH689A
  - HPE FlexNetwork 5930 24-port 10GBASE-T and 2-port 40GbE QSFP+ MACsec Module JH690A



https://t.me/learningnets

- Latency:
  - 10 Gbps Latency: < 1 µs
- Throughput:
  - 4-slot, up to 1429 Mpps; 2-slot up to 1071 Mpps
- Routing/switching capacity:
  - 4-slot, 2560 Gbps; 2-slot, 1440 Gbps
  - Fibre Channel 8Gb/4Gb/2Gb support on any of the 24 converged ports of the JH184A module
  - FCoE support on 10GbE and 40GbE ports
- Stacking capabilities:
  - IRF with 9 switches (Ethernet), 2 switches (storage)
  - CLI, and Intelligent Management Center (iMC/VAN/VFM) fabric management

## HPE 5820 Ethernet/FCoE Switch

The 5820 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:


https://t.me/learningnets

- HPE 4-port 8/4/2 Gb/s FCoE SFP+ 5820 Module
- HPE 5820X/A5800 4-port 10-GbE SFP+ Module
- HPE 5820X/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: Converged SAN fabric configuration** on page 92 and **Figure 44: FCoE integrated with FC SAN fabric** on page 94.
  - Access to native FCoE target devices not supported via 10-GbE ports

## HPE FlexFabric 5700 Switch

The HPE FlexFabric 5700 Switch Series is a family of cost-effective, high-density, ultra-low-latency, topofrack (ToR) switches. They are ideally suited for deployment at the server access layer of large enterprise data centers.

The FlexFabric 5700 switch series includes the following models:

- **5700-40XG-2QSFP+ (JG896A) 42 Port Switch** — It comes with 40x fixed 1000 / 10000 Mb/s SFP+ ports, and 2x QSFP+ for 40-GbE ports, which can optionally be split to 4x10GbE ports.
- **5700-48G-4XG-2QSFP+ (JG894A) 54 Port Switch** — It comes with 48x 10/100/1000 Mb/s ports, 4x fixed 1000 / 10000 Mb/s SFP+ ports, and 2x QSFP+ for 40-GbE ports, which can optionally be split to 4x10GbE ports.

  **NOTE:** FCoE is supported over 10GbE and 40GbE links, and iSCSI over 1GbE and 10GbE links is supported.

- **5700-32XGT-8XG-2QSFP+ (JG898A) 42 Port Switch** — It comes with 32x RJ-45 10GBASE-T ports, 8x fixed 1000 / 10000 Mb/s SFP+ ports, and 2x QSFP+ for 40-GbE ports, which can optionally be split to 4x10GbE ports.

**Features and benefits:**

The FlexFabric 5700 switch series products provides the following features:

- FCoE support on 10GbE and 40GbE ports
- Cut-through with ultra-low-latency and wire speed
- HPE Intelligent Resilient Framework (IRF) for virtualization and two-tier architecture
- High 1GbE/10GbE ToR port density with 10 GbE and 40 GbE uplinks
- Layer-2 and Light Layer-3 Features with Static Routing and RIP
- Convergence ready with DCB/QCN, FCoE, and TRILL
- iSCSI support; optional DCB lossless support
- CLI, OneView, and Intelligent Management Center (iMC/VAN/VFM) fabric management
- IPv6 Host/Management



https://t.me/learningnets

- IPv6 Routing
- OpenFlow (SDN)
- Redundant Power
- Replaceable Fans
- Stacking
- Fibre Channel over Ethernet (FCoE)
- OpenFlow now - v1.3.1

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

The HPE 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.

## HPE StoreFabric SN8600B 64/32 Short Wave 8QSFP Integrated Fibre Channel Blade

Introducing HPE StoreFabric SN8600B 32Gb 64-port FC Blade providing 64 ports of 32Gbps Fibre Channel or 10Gb/25Gb/40Gb Fibre Channel over Ethernet (FCoE) connectivity via 16 front QSFP ports.


https://t.me/learningnets

- Scales the SN8600B Director to 512 ports while maximizing space utilization with 33 percent more device connectivity in a high-density blade
- Increases agility by enabling flexible architectures with concurrent Fibre Channel, NVMe, or FCoE connectivity
- Handles more workloads with Gen6 Fibre Channel performance and by increasing total system bandwidth up to 20 Tbps
- Reduces complexity with 75 percent fewer cables and increases port density with Q-Flex connections
- Drives IT innovation by optimizing NVMe performance through integrated real-time latency monitoring with IO Insight
- With the SN8600B Brocade FC32-64 Port Blade and SN8600B Director you can optimize the performance of NVMe over Fibre Channel at a granular level by leveraging integrated, non-intrusive, real-time monitoring and alerts. This proactive monitoring of NVMe over Fibre Channel traffic and VMs too provides administrators with key insights for maintaining optimal network health and performance.

## Cisco Nexus 5548UP/5596UP Converged Network Switches

Hewlett Packard Enterprise supports the interoperability through Cisco Nexus 5548UP and 5596UP switches to HPE storage systems and CNAs. For more information about supported configurations, see the SPOCK website at **http://www.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access.

**NOTE:** Hewlett Packard Enterprise neither sells nor provides HPE Services support for Cisco Nexus 5548UP and 5596UP switches.

**Features**

- L2 switch ports, VLAN trunks, VLAN remapping
- IEEE DCB and FCoE support
- L3 support with up to 4,096 VLANs
- Support up to 16 fabric extenders on each Nexus 5500 with L3 module
- Cisco Nexus 5548UP (1 RU)
  - Up to 48 ports — 32 unified ports and 1 16-port expansion module
  - 32 unified ports can be configured as 1- and 10-GbE and FCoE, or 8/4/2/1 Gb/s FC ports
- Cisco Nexus 5596UP (2 RU)
  - Up to 96 ports — 48 unified ports and 3 16-port expansion modules
  - 48 unified ports can be configured as 1- and 10-GbE and FCoE, or 8/4/2/1 Gb/s FC ports
- Optional expansion modules include:
  - 16-port Unified Port Expansion Module — Up to 16 of 1/10 GbE FCoE ports or 16 of 8/4/2/1 Gb/s FC ports
  - 8-port FC (8/4/2/1-Gb) + 8-port Ethernet and FCoE (1/10 GB) module



https://t.me/learningnets

## Cisco Nexus 6001P switch

Cisco Nexus 6001P - The Cisco Nexus 6001P (Figure 1) is a 1RU, 10- and 40- Gigabit Ethernet switch offering wire-speed performance for up to sixty-four 10 Gigabit Ethernet ports (using Quad Small Form-Factor Pluggable [QSFP] breakout cables) for Ethernet and FCoE traffic, with an overall throughput of 1.28 terabits per second (Tbps). The Cisco Nexus 6001P offers 48 fixed 1/10 Gigabit Ethernet Enhanced Small Form-Factor Pluggable (SFP+) ports and four 40 Gigabit Ethernet QSFP+ ports. Each 40 Gigabit Ethernet port can be split into four 10 Gigabit Ethernet ports using a QSFP+ breakout cable.

## Cisco Nexus 6004EF switch

The Cisco Nexus 6004EF is a 4-rack-unit (4RU) 10 and 40 Gigabit Ethernet switch with the industry's highest port density in a compact, energy-efficient form factor. It offers eight line-card expansion module (LEM) slots to support up to 96 ports of 40-Gbps throughput. Each LEM supports 12 ports of 40 Gigabit Ethernet in a Quad Small Form-Factor Pluggable (QSFP) footprint. Each 40 Gigabit Ethernet port can also be split into four line-rate 10 Gigabit Ethernet ports using QSFP breakout cables. This flexibility makes the Cisco Nexus 6004EF the only fully extensible fixed 10 and 40 Gigabit Ethernet platform in the industry. The Cisco Nexus 6004EF also offers a unified expansion module for 1 and 10 Gigabit Ethernet and Fibre Channel and FCoE support.

## Cisco Fabric Extender for HPE BladeSystem

The Cisco Fabric Extender for BladeSystem (Model B22HP) provides an extension of the Cisco Nexus switch fabric to the HPE server edge (see **Figure 36: FC and FCoE storage with Cisco Fabic Extender for HPE BladeSystem configuration** on page 81). Logically, it behaves like a remote line card to a parent Cisco Nexus 5000 series or 5500 series switch, with the Fabric Extender and the parent Nexus together forming a distributed modular system. When combined with HPE server embedded CNA or mezzanine adapters, it provides total network flexibility with 1Gb/10Gb Ethernet, Data Center Bridging (FCoE) and iSCSI network connections.

**Features**

- Sixteen 1Gb/10Gb downlinks
- Each Ethernet downlink supports Fibre Channel over Ethernet (FCoE), Datacenter Bridging (DCB), and iSCSI
- Each Ethernet downlink can autonegotiate to 1Gb to support legacy NICs and/or Wake On LAN
- Eight 10Gb uplinks:
  - Each Fabric uplink carries both Ethernet and Fibre Channel over Ethernet (FCoE) traffic
  - Fabric uplinks support .5m/1m/3m/5m/7m/10m twinax, Fabric Extender Transceiver (FET-10G), 10Gbase-SR, and 10Gbase-LR optics
- Up to 24 Fabric Extenders can be managed by a single Nexus 5000 or 5500 series switch
- Ports support Ethernet, iSCSI and DCB protocols



https://t.me/learningnets

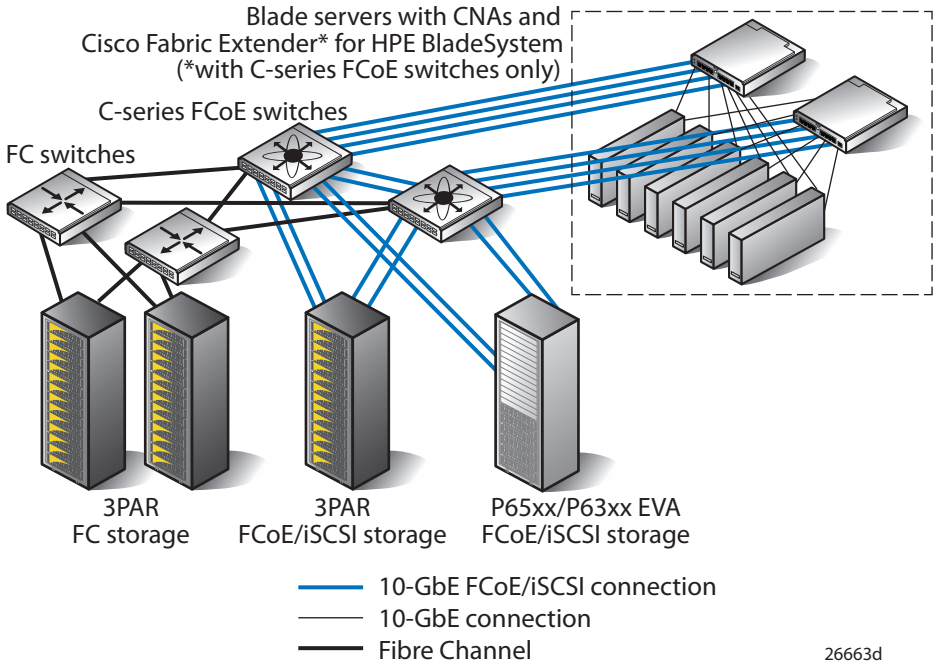



**Figure 36: FC and FCoE storage with Cisco Fabic Extender for HPE BladeSystem configuration**

## Converged network adapters

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

**HPE CNA part numbers** contains relevant information on all CNAs, standalone and embedded, currently supported by HPE.

### Table 10: HPE CNA part numbers

| HPE model | Port Type | Bus | Ports | Port Speed | Protocols | Advanced features | Compliances |
|---|---|---|---|---|---|---|---|
| HPE CN1000E Dual Port<br>HPE P/N:<br>AW520A<br>AW520B | SFP+ | PCIe 2.0 x8 | Dual | 10GbE | NIC, ISCSI, FCoE | 802.10 VLAN tagging<br>MSI-X<br>NIC teaming (bonding)<br>Receive side scaling (RSS)<br>Jumbo frames<br>ISCSI/FCoE Offload<br>PXE boot<br>NPIV, AER, TOE | ANSI INCITSTII FC-BB-5 2.0, FC-PI-2, FC-GS-4, FC-TAPE, and FCP-3 Advanced Error Reporting (AER)<br>IEEE 802.3ae, 802.1q, 802.3ad, 802.3x, 802.1p, 802.1Qaz , 802.1Qaz, 802.1Qbb , 802.1ab |
| HPE CN1100E Dual Port<br>HPE P/N:<br>BK835A | | | | | | | |

*Table Continued*



https://t.me/learningnets

| HPE model | Port Type | Bus | Ports | Port Speed | Protocols | Advanced features | Compliances |
|---|---|---|---|---|---|---|---|
| HPE CN1200E Dual Port<br>HPE P/N:<br>E7Y06A | | PCIe 3.0 x8 | | | | | |
| HPE StoreFabric CN1200E-T 10GBASE-T<br>HPE P/N:<br>N3U51A | RJ-45 | | | | | | |
| HPE FlexFabric 10gb 2-port 554M Adapter<br>HPE P/N:<br>647590-B21 | N/A | PCIe 2.0 x8 Type A Mezz | | | | 802.1Q VLAN tagging<br>NIC teaming (bonding)<br>Receive side scaling (RSS)<br>Jumbo frames<br>ISCSI / FCoE Offload<br>PXE boot, TOE, SR-IOV<br>Flex-10 support | IEEE: 802.1p, 802.1q, 802.1qau, 802.3ad, 802.3ae, 802.3ap, 802.3x |
| HPE FlexFabric 10gb 2-port 556LR-T Adapter<br>HPE P/N:<br>794525-B21 | SFP+ | PCIe 3.0 x8 FlexLOM | | | | 802.1Q VLAN tagging<br>NIC teaming (bonding)<br>Receive side scaling (RSS)<br>Jumbo frames<br>ISCSI / FCoE Offload<br>PXE boot, TOE, SR-IOV<br>RoCE | IEEE: 802.3ae, 802.1Q, 802.3x, 802.1p, 802.3ad/ LACP, 802.1AB, 802.1Qbg, 802.1Qbb, 802.1Qaz, 802.3an |

*Table Continued*



https://t.me/learningnets

| HPE model | Port Type | Bus | Ports | Port Speed | Protocols | Advanced features | Compliances |
|---|---|---|---|---|---|---|---|
| HPE FlexFabric 20gb 2-port 650M Adapter<br>HPE P/N:<br>700767-B21 | N/A | PCIe 3.0 x8 Type A Mezz | | 10/20GbE | | MSI-X<br>802.1Q VLAN tagging<br>NIC teaming (bonding)<br>Receive side scaling (RSS)<br>Jumbo frames<br>ISCSI / FCoE Offload<br>PXE boot, SR-IOV<br>Flex-20 | IEEE: 802.3ae, 802.1Q, 802.3x, 802.1p, 802.3ad/ LACP, 802. 1AB, 802.1Qbg, 802.1Qbb, 802.1Qaz, 802.3ap |
| HP StoreFabric CN11OOR-2P<br>HPE P/N:<br>QW990A | SFP+ | PCIe 2.0 x8 | | 10Gbps | | 802.1Q VLAN tagging<br>MSI-X<br>NIC teaming (bonding)<br>Receive side scaling (RSS)<br>Jumbo frames<br>ISCSI / FCoE Offload<br>PXE boot<br>NVG RE/ VXLAN<br>tunneling offloads<br>NPAR (Not on 536FLR-T)<br>SR-IOV | x8 PCI Express V2.0 (5 GT/s) support<br>PCI Bus Power Management<br>Interlace (rev 1.2)<br>ACPI v2.0, SMBus 2.0, IPv4 (RFQ 791), IPv6 (RFC 2460)<br>IEEE: 802.3x, 802.3ad/802.1ax,<br>802.3ae, 802.1q, 802.1as/1588, 802.1Qaz, 802.1Qbb, 802.1Qbg, 802.1Qbh |
| HPE StoreFabric CN1100R-T 10GBase-T<br>HPE P/N:<br>N3U52A | RJ-45 | | | | | | |
| HP FlexFabric 10Gb 2-Port 533FLR-T Adapter<br>HPE P/N:<br>700759-B21 | RJ-45 | Flexible LOM Rack PCIe 2.0 x8 | | | | AER<br>TOE | 0 IPv4 (RFQ 791), IPv6 (RFC 2460)<br>IEEE: 802.3x, 802.3ad/802.1ax, 802.3ae, 802.1q, 802.1as/1588, 802.1Qaz, 802.1Qbb, 802.1Qbg, 802.1Qbh |

*Table Continued*



https://t.me/learningnets

| HPE model | Port Type | Bus | Ports | Port Speed | Protocols | Advanced features | Compliances |
|---|---|---|---|---|---|---|---|
| HP FlexFabric 10Gb 2-Port 534FLR-SFP+ Adapter<br>HPE P/N:<br>700751-B21 | SFP+ | | | | | | IEEE:802.3, 802.3ae, 802.3x, 802.2x, 802.3ad, 802.1Qaz, 802.1Qau, 802.1Qbb, 802.1Qbg, 802.1ax |
| HP FlexFabric 10Gb 4-port 536FLR-T Adapter<br>HPE P/N:<br>764302-B21 | RJ-45 | PCIe 3.0 x8 | Quad | | | | IEEE:802.3, 802.3x, 802.2x, 802.3ad, 802.1Qaz, 802.1Qau, 802.1Qbb, 802.1Qbg, 802.1ax |
| HP FlexFabric 2-Port 534M Adapter<br>HPE P/N:<br>700748-B21 | N/A | PCIe 2.0 Type A Mezz | Dual | | | | IEEE:802.3, 802.3ab, 802.3u, 802.3x, 802.3ad, 802.3p, 802.1q, 802.3ae, 802.3ap |
| HP FlexFabric 10Gb 2-Port 536FLB Adapter<br>HPE P/N:<br>766490-B21 | | PCIe 3.0 x8 FlexLOM Blade | | | | | |

*Table Continued*



https://t.me/learningnets

| HPE model | Port Type | Bus | Ports | Port Speed | Protocols | Advanced features | Compliances |
|---|---|---|---|---|---|---|---|
| HP FlexFabric 20Gb 2-Port 630FLB Adapter<br>HPE P/N:<br>700065-B21 | | | | 10/20Gbps | | | IEEE:802.3, 802.1ab, 802.3x, 802.3ad, 802.3p, 802.1q, 802.3ae, 802.1Qau, 802.3ap |
| HP FlexFabric 20Gb 2-Port 630M Adapter<br>HPE P/N:<br>700076-B21 | | PCIe 3.0 Type A Mezz | | | | | IEEE:802.3, 802.3ab, 802.3u, 802.3x, 802.3ad, 802.3p, 802.1q, 802.3ae, 802.3ap, 1588 |

# FCoE storage systems

Hewlett Packard Enterprise supports the following FCoE storage systems:

- **HPE 3PAR StoreServ 7200/7400/7450/8200/8400/8440/8450 AFA/10400/10800/20450/20800/20850 with 10GbE FCoE host ports** on page 85
- **XP7 with the FCoE 10-GbE channel adapter (8 ports per adapter)** on page 85
- **P9500 with the FCoE 10-GbE channel adapter (4 ports per adapter)** on page 86
- **P63xx/P65xx with FCoE/iSCSI 10-GbE front end host ports (2 ports per controller)** on page 87

## HPE 3PAR StoreServ 7200/7400/7450/8200/8400/8440/8450 AFA/ 10400/10800/20450/20800/20850 with 10GbE FCoE host ports

### FCoE host ports

For more information, see the **HPE Native iSCSI products** on page 346 and the HPE 3PAR StoreServ FCoE product documentation.

### Operating system and multipath software support

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

## XP7 with the FCoE 10-GbE channel adapter (8 ports per adapter)

### FCoE channel adapter

The FCoE channel adapter can be configured with redundancy using multiple ports on a single channel adapter or using multiple channel adapters, see **Figure 38: XP7 and P9500 FCoE end-to-end configuration** on page 87.

### Operating system and multipath software support

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



https://t.me/learningnets

### Table 11: XP7 10-GbE FCoE operating system and multipath software support

| Operating system | Multipath software | Clusters XP7 storage system |
|---|---|---|
| Red Hat Linux, SUSE Linux | Device Mapper | None |
| VMware | VMware embedded | None |
| Microsoft Windows Server 2008, 2012, 2016 | MPIO with Microsoft DSM | Failover Clustering, MSCS |

## P9500 with the FCoE 10-GbE channel adapter (4 ports per adapter)

#### FCoE channel adapter

The FCoE channel adapter can be configured with redundancy using multiple ports on a single channel adapter or using multiple channel adapters, see **Figure 38: XP7 and P9500 FCoE end-to-end configuration** on page 87.

#### Operating system and multipath software support

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

### Table 12: P9500 10-GbE FCoE operating system and multipath software support

| Operating system | Multipath software | Clusters P9500 storage system |
|---|---|---|
| Red Hat Linux, SUSE Linux | Device Mapper | None |
| VMware | VMware embedded | None |
| Microsoft Windows Server 2008, 2012, 2016 | MPIO with Microsoft DSM | Failover Clustering, MSCS |



https://t.me/learningnets

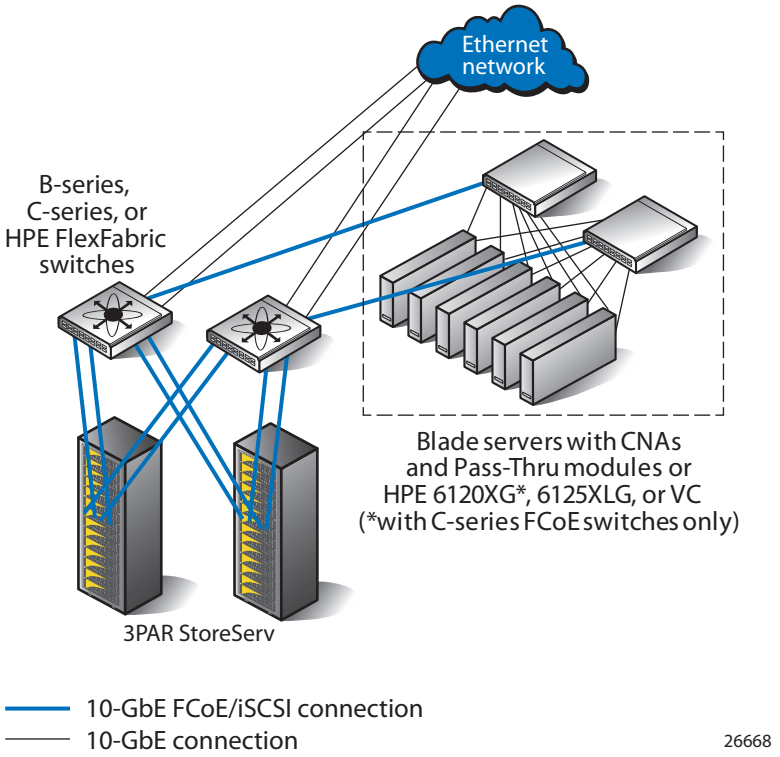



**Figure 37: 3PAR StoreServ FCoE end-to-end configuration**

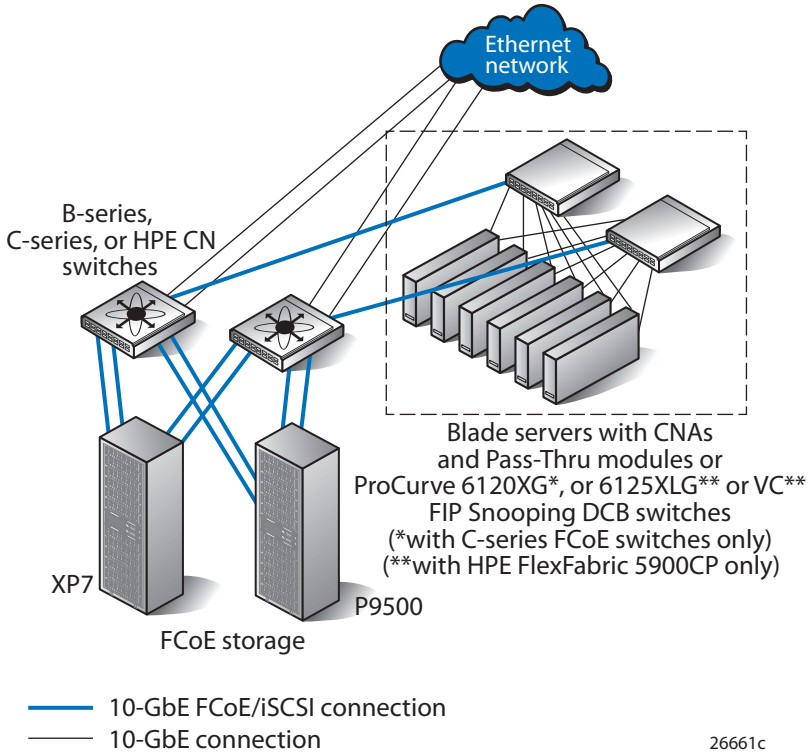



**Figure 38: XP7 and P9500 FCoE end-to-end configuration**

## P63xx/P65xx with FCoE/iSCSI 10-GbE front end host ports (2 ports per controller)

### iSCSI/FCoE module



https://t.me/learningnets

The iSCSI/FCoE module is configured in a dual-controller configuration in the HPE P63**xx**/P65**xx** EVA, see **Figure 39: P63xx/P65xx FCoE/iSCSI end-to-end configuration** on page 89. 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 P63**xx**/P65**xx** EVA Fibre Channel support, see **P6000 storage system rules** on page 238. For information about P63**xx**/P65**xx** EVA 1-GbE iSCSI support, see **iSCSI storage** on page 341.

**Operating system and multipath software support**

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

### Table 13: P63xx/P65xx 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 DSM (supports both CN1000E and CN1100E) | MPIO with Microsoft DSM (supports CN1000E only) | Failover Clustering, MSCS |
| Microsoft Windows Server 2003 | MPIO with DSM (supports CN1000E and CN1100E) | MPIO with Microsoft DSM (supports CN1000E only) | |
| Red Hat Linux, SUSE Linux | Device Mapper | | None |
| VMware | VMware MPxIO | | None |
| Microsoft Windows Server 2012 | MPIO with DSM<br>(supports CN1000E, CN1100E, CN1100R, CN1000Q) | MPIO with Microsoft DSM<br>(supports CN1000E, CN1100E, CN1100R, CN1000Q) | Failover Clustering, MSCS |

For more information, see the iSCSI or iSCSI/FCoE configuration rules and guidelines chapter in the *HPE P63x0/P65x0 Enterprise Virtual Array User Guide* .

**Figure 39: P63xx/P65xx FCoE/iSCSI end-to-end configuration** on page 89 describes the P63**xx**/P65**xx** EVA 10-GbE iSCSI/FCoE configuration.



https://t.me/learningnets

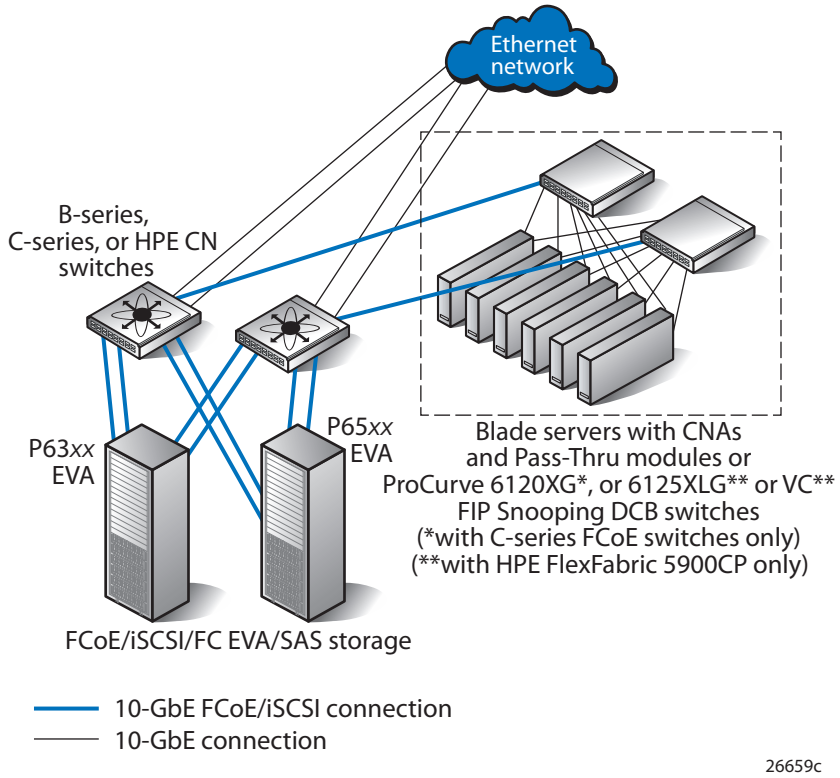



**Figure 39: P63xx/P65xx FCoE/iSCSI end-to-end configuration**

**Figure 40: Blade or rack server and mixed FC/FCoE storage with B-series, C-series, or HPE FlexFabric 5900CP FCoE switches** on page 90 and **Figure 41: Blade server mixed FC/FCoE storage with HPE 5820 switches** on page 91 show mixed FC and FCoE configurations using FC and FCoE storage targets.

**Figure 40: Blade or rack server and mixed FC/FCoE storage with B-series, C-series, or HPE FlexFabric 5900CP FCoE switches** on page 90 depicts blade server FC/FCoE configurations with B-series, C-series, or FlexFabric FCoE switches and either Pass-Thru modules, 6120XG or 6125XLG DCB switches, or Virtual Connect modules.

- If Pass-Thru modules are used, this configuration is supported with B-series or C-series FCoE and corresponding FC switches.
- If 6120XG switches are used, this configuration is supported with C-series FCoE switches and FC switches only.
- If 6125XLG switches are used with FlexFabric 5900CP switches, the 6125XLG switch must be in NPV mode.



https://t.me/learningnets

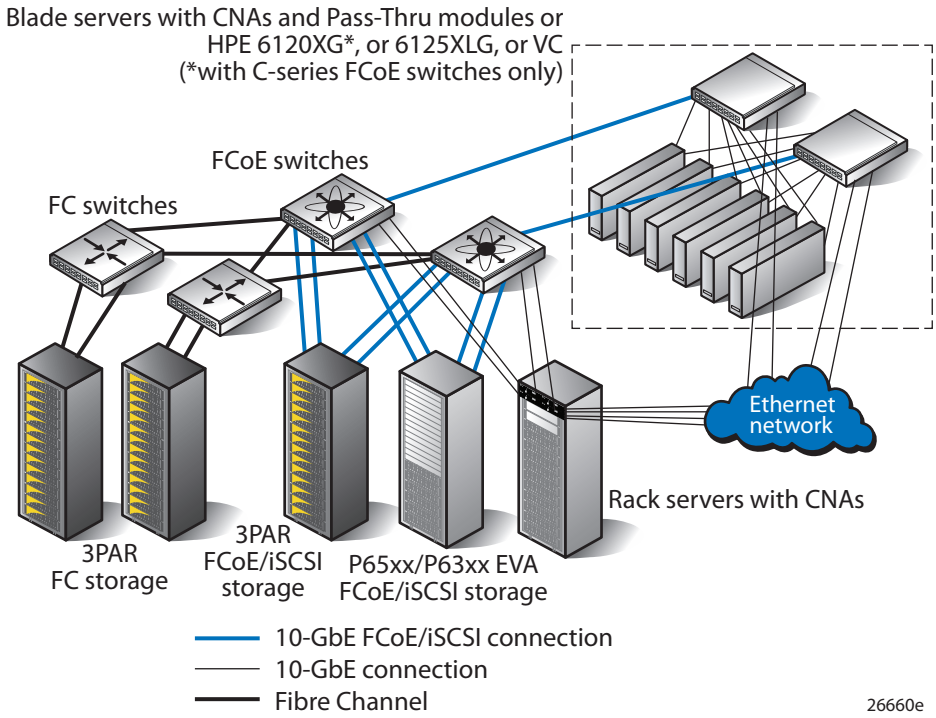



**Figure 40: Blade or rack server and mixed FC/FCoE storage with B-series, C-series, or HPE FlexFabric 5900CP FCoE switches**

**NOTE:** For switch series and switch model NPV/NPIV support with the FlexFabric 5900CP switch, see the *HPE FlexFabric 5900CP Solutions Configuration Guide* and the SPOCK website at **http://www.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access.

**Figure 41: Blade server mixed FC/FCoE storage with HPE 5820 switches** on page 91 depicts blade server FC/FCoE configurations with 5820 switches and either Pass-Thru modules or 6120XG DCB switches. Connectivity to B-series, C-series or H-series FC switches is supported with 5820 switches.



https://t.me/learningnets

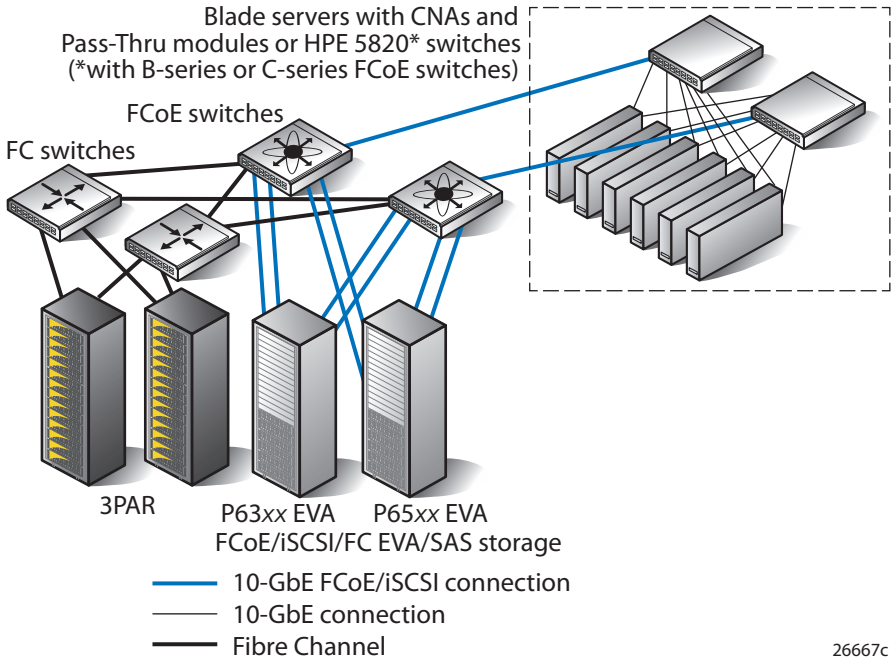



**Figure 41: Blade server mixed FC/FCoE storage with HPE 5820 switches**

# FCoE configuration rules

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

**Figure 42: Converged SAN fabric configuration** on page 92 shows a converged SAN fabric with servers using CNAs only.

**NOTE:** Fibre Channel HBAs can also be connected to B-series, C-series, or FlexFabric 5900CP converged ports as F_Ports.



https://t.me/learningnets

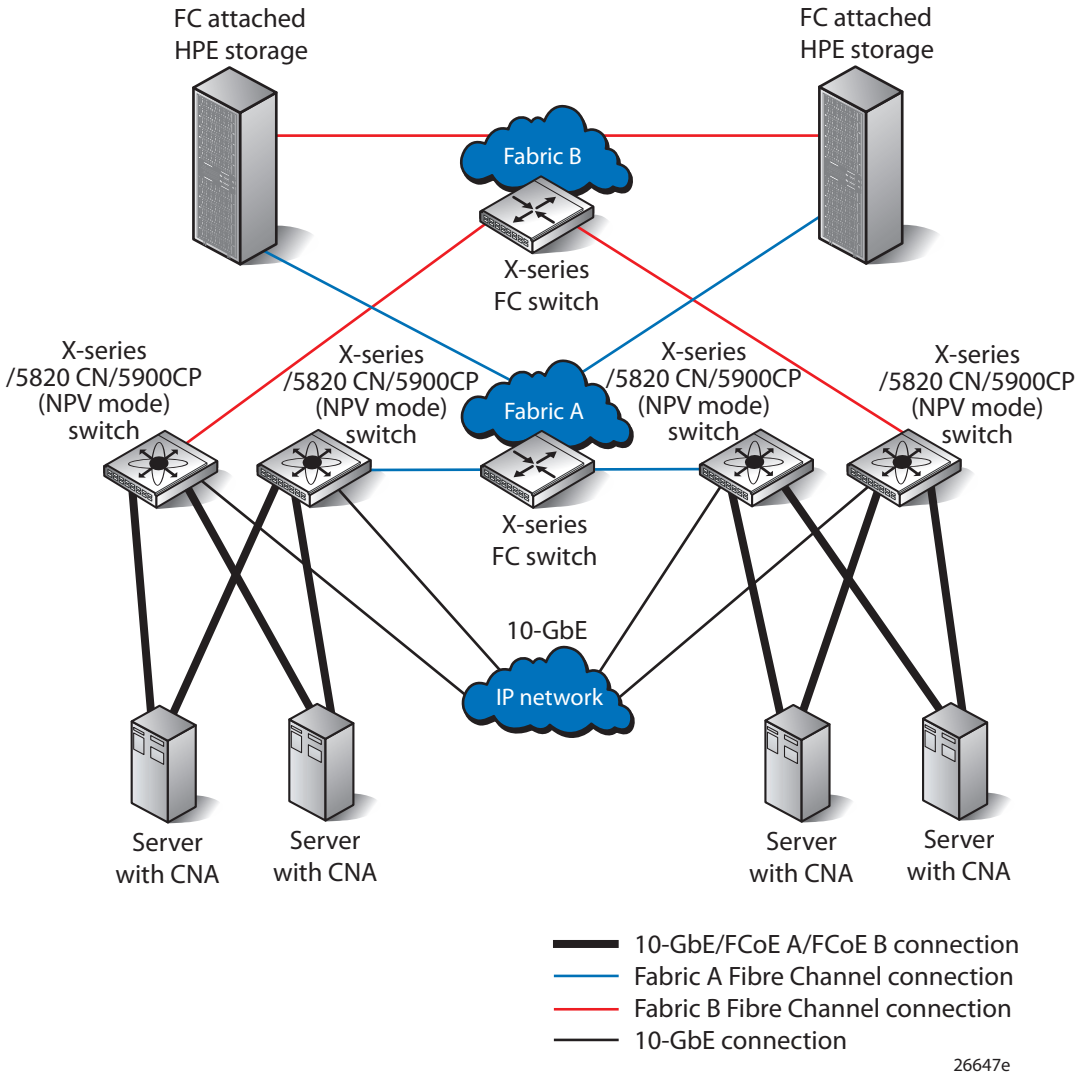



**Figure 42: Converged SAN fabric configuration**

**NOTE:** For switch series and switch model NPV/NPIV support with the FlexFabric 5900CP switch, see the *HPE FlexFabric 5900CP Solutions Configuration Guide* and the SPOCK website at **http://www.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access.

**Figure 43: Converged SAN fabric with a B-series FCoE blade** on page 93 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.



https://t.me/learningnets

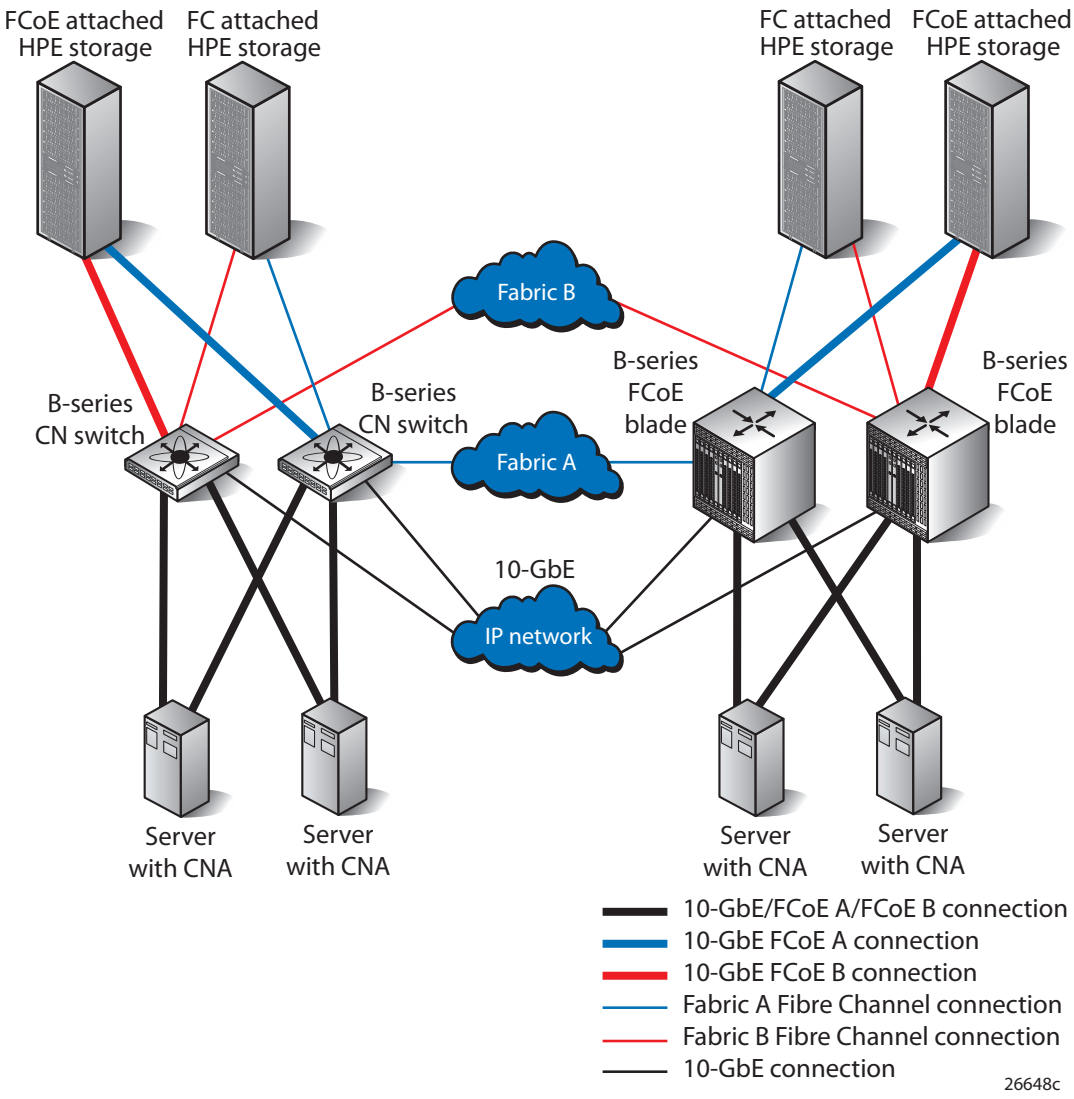



**Figure 43: Converged SAN fabric with a B-series FCoE blade**

**NOTE:** Consider the following for a converged SAN fabric with an FCoE blade:

- End-to-end FCoE is supported through FC ISLs.
- FCoE traffic can be carried across multiple FC ISL hops, eventually terminating in FC or FCoE storage.

As an example, a CNA connected to the B-series CN switch on the left side can communicate with an FCoE target connected to the B-series FCoE blade on the right, going through Fabric A or B depending on where it is connected.

**Figure 44: FCoE integrated with FC SAN fabric** on page 94 shows a converged fabric that has servers with CNAs and servers with FC HBAs and NICs.

**NOTE:** Fibre Channel HBAs can also be connected to B-series, C-series, or FlexFabric 5900CP Converged Network switch F_Ports.



https://t.me/learningnets

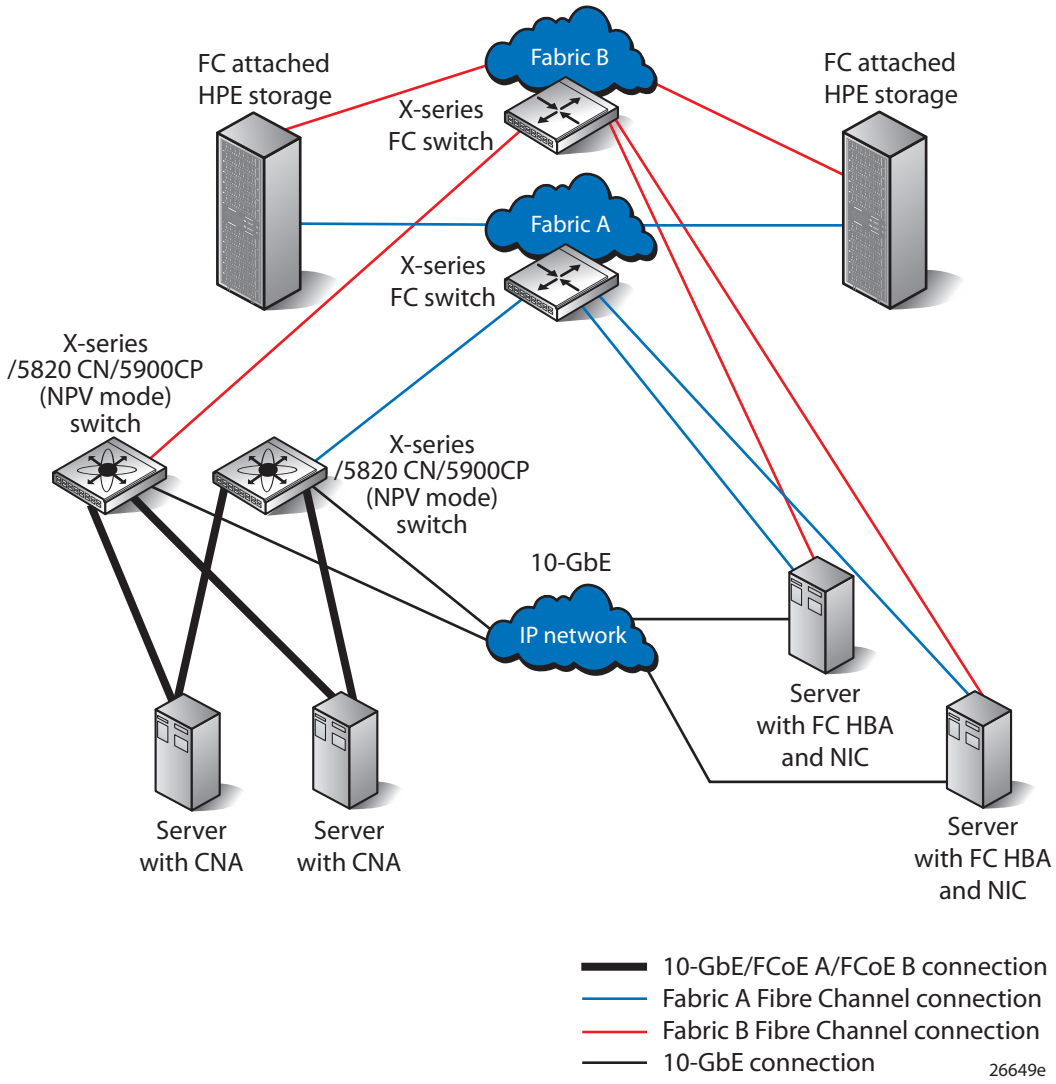



**Figure 44: FCoE integrated with FC SAN fabric**

**NOTE:** For switch series and switch model NPV/NPIV support with the FlexFabric 5900CP switch, see the *HPE FlexFabric 5900CP Solutions Configuration Guide* and the SPOCK website at **http://www.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access.

**Figure 45: FCoE end-to-end MPX200 direct-connect P6000/EVA FC storage configuration** on page 95 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.



https://t.me/learningnets

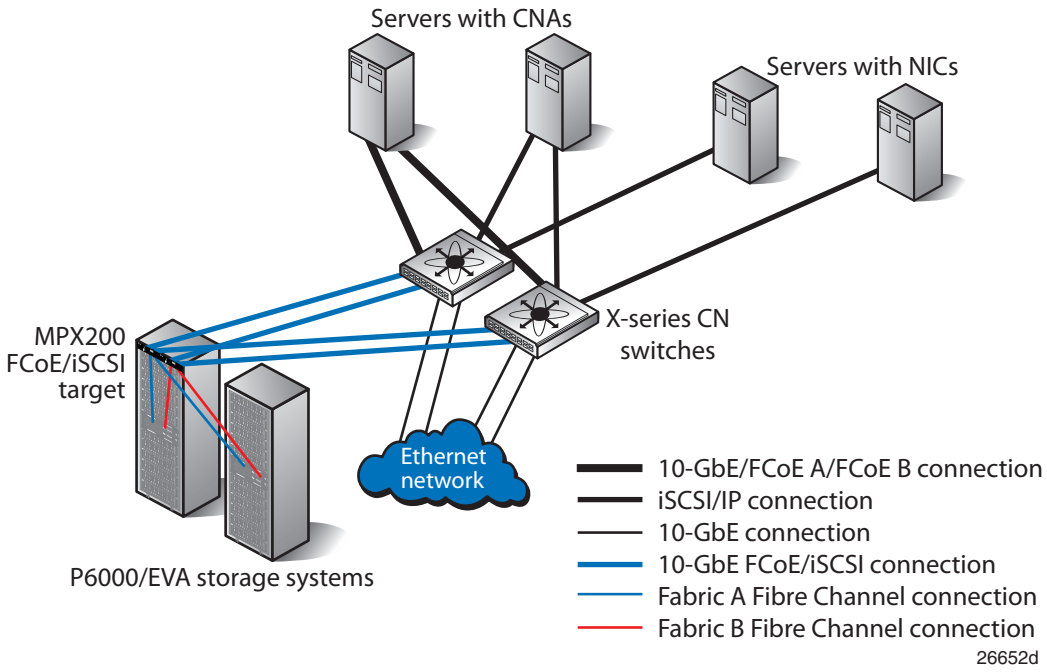



**Figure 45: FCoE end-to-end MPX200 direct-connect P6000/EVA FC storage configuration**

The configuration shown in **Figure 46: FCoE end-to-end MPX200 fabric-connect P6000/EVA FC storage configuration** on page 96 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.



https://t.me/learningnets

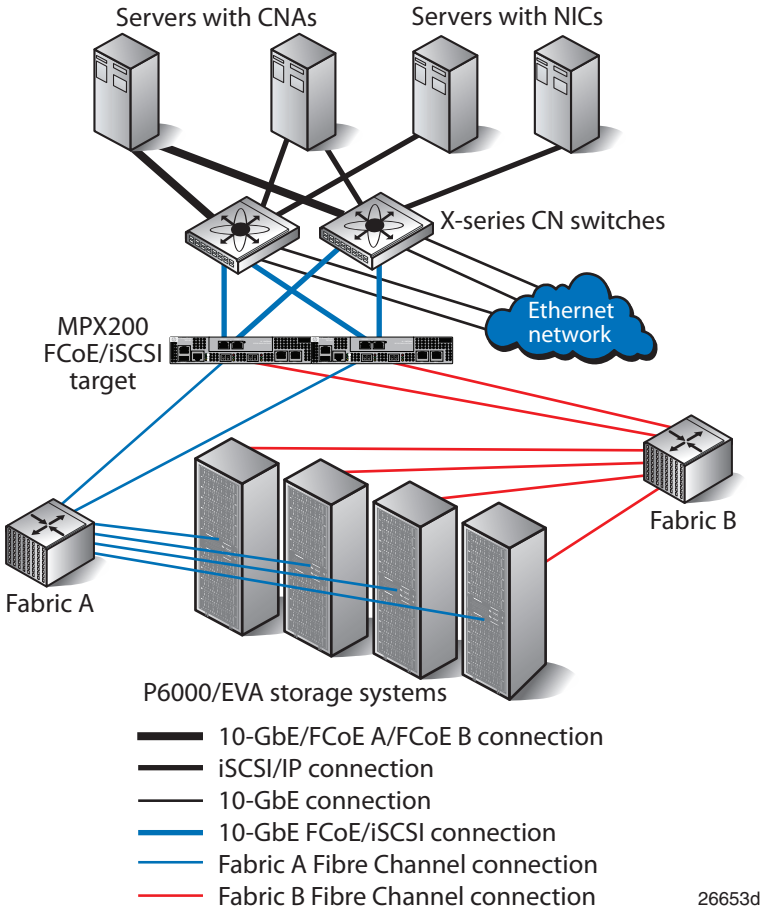



**Figure 46: FCoE end-to-end MPX200 fabric-connect P6000/EVA FC storage configuration**

**CNA configuration rules**

### Table 14: HPE CNA requirements

| Item | Description |
|---|---|
| Number of CNAs per server (maximum) | 2 |
| Cable types and distances | • SFP+ optical: CN1100E, CN1100R, and CN1000Q 100m; CN1000E 40m<br>• Copper: CN1100R 5m; CN1100E, CN1000E, and CN1000Q 7m and 10m (with C-series FCoE CN switches only) |

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

## Server support

FCoE is supported with ProLiant G5, G6, G7, and Gen8 servers.



https://t.me/learningnets

For current server support, see the server QuickSpecs at **http://www.hpe.com/info/quickspecs/**.

## HPE FlexFabric 5900CP Switch and HPE 5900AF Switch

For information about the FlexFabric 5900CP switch, see the *HPE FlexFabric 5900CP Solution Configuration Guide*. This document is available at **www.hpe.com/info/flexfabric5900cp-docs**.

For the current FlexFabric 5900CP and 5900AF product support, see the SPOCK website at **http://www.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access.

## HPE 5820 Converged Network Switch

This section describes support for the 5820X-14XG FCoE CN switch and FCoE fabric rules. For information on configuring the 5820X-14XG FCoE CN switch, see **FCoE switch configuration quick-setup instructions** on page 415.

### Usage

The 5820X 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** on page 40.

### FCoE Converged Network Switch fabric rules

This section describes the fabric rules for the 5820X FCoE CN switch and other factors you should consider when building 5820X FCoE CN fabrics. For Fibre Channel switch fabric rules pertaining to the specific family of FC switches being used, see **Fabric infrastructure rules** on page 102.

#### Table 15: HPE 5820X CN switch fabric rules

| Rule number | Description |
|---|---|
| 1 | At least one 4-port 8/4/2 Gb/s FCoE SFP+ 5820 module is required for the 5820X-14XG to operate as an FCoE CN switch. |
| 2 | At least one FC port of the 5280 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 5820 FCoE module. |
| 4 | The 5280 FCoE module FC ports only support FC switch connections. |
| 5 | Native FCoE storage is not supported using the 5820 10-GbE ports. |

*Table Continued*



https://t.me/learningnets

| Rule number | Description |
|---|---|
| 6 | Up to two 5820 FCoE modules are supported in a single 5820X switch. However, when using two A5280 FCoE modules, consider the following:<br>• When two 5820 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.<br>• The current Emulex CNA firmware does not support the load balancing method used when two 5820 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. |
| 7 | Connection to Brocade 8 Gb/s switches and ARB(FF)/ARB(FF) mode setting: Hewlett Packard Enterprise recommends all 8 Gb/s Brocade switches running FOS V6.3.1a and later be set to mode 3 using the `portcfgfillword` command. For more information, see the *Fabric OS Command Reference Manual* supporting FOS V6.4 and the FOS release notes. |
| 8 | When connecting the 5820 to multiple SANs, unique discovery and data VLANs are required for each SAN. |

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

For current information on supported 5820X firmware, operating systems, and options, see the SPOCK website at **http://www.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access.

## C-series and Cisco FCoE Converged Network switches

This section describes support for C-series and Cisco FCoE CN switches, switch models, and FCoE fabric rules. For more information on configuring FCoE CN switches, see **FCoE switch configuration quick-setup instructions** on page 415.

### Cisco Nexus 5548UP FCoE Converged Network Switch

The Cisco Nexus 5548UP Converged Network Switch is a 1RU, 10-GbE and FCoE switch offering up to 960 Gb/s throughput. It has up to 48 ports: 32 fixed unified ports and 1 16-port expansion slot. The unified ports can be configured as 1- and 10-Gb/s SFP+ Ethernet and FCoE ports, or as Fibre Channel ports. These two configurations are mutually exclusive, but either type can be configured for any of the ports.

### Cisco Nexus 5596UP FCoE Converged Network Switch

The Cisco Nexus 5596UP Converged Network Switch is a 2RU switch. This 10-GbE and FCoE switch offers up to 1,920 Gb/s throughput and up to 96 ports. The switch has 48 unified 1- and 10-GbE and FCoE ports and 3 16-port expansion slots. The unified ports can be configured as either 1- or 10-GbE and FCoE, or 1/2/4/8 Gb/s native Fibre Channel ports. The use of GbE or Fibre Channel on a port is mutually exclusive, but either type can be configured for any of the ports.

### Expansion Module Options for the Cisco Nexus 5500 Series switches

The Cisco Nexus 5500 platform is equipped with expansion modules that can be used to increase the number of 10-GbE and FCoE ports, connect to Fibre Channel SANs with 8, 4, 2, or 1 Gb/s Fibre Channel switch ports, or both.



https://t.me/learningnets

Hewlett Packard Enterprise support for the Cisco Nexus 5548UP and 5596UP includes the following types of expansion modules:

- 8-port FC (8/4/2/1G) + 8-port Eth/FCoE Module — Fibre Channel plus Ethernet module that provides eight 1- or 10-GbE and FCoE ports using the SFP+ interface, and eight ports of 8/4/2/1 Gb/s native Fibre Channel connectivity using the SFP+/SFP interface.
- 16-port Unified Port Expansion Module — Unified port module that provides up to 16 1- or 10-GbE and FCoE ports using the SFP+ interface or up to 16 ports of 8/4/2/1 Gb/s native Fibre Channel connectivity using the SFP+ and SFP interfaces. The use of GbE or Fibre Channel on a port is mutually exclusive, but either type can be configured for any of the ports.

## Expansion Modules for the Cisco Nexus 6004EF switch

12-port 40 Gigabit Ethernet and FCoE expansion module. Each 40 Gigabit Ethernet expansion module provides 12 ports of 40 Gigabit Ethernet and FCoE ports using a QSFP interface. With all eight expansion modules installed, the Cisco Nexus 6004EF delivers 96 ports of QSFP or 384 ports of 10 Gigabit Ethernet (SFP+) using the breakout cables.

20-port unified expansion module. The unified port module provides up to twenty 1 and 10 Gigabit Ethernet and FCoE ports using the SFP+ interface or up to 20 ports of 8/4/2-Gbps native Fibre Channel connectivity using the SFP+ and SFP interfaces; the use of 1 and 10 Gigabit Ethernet or 8/4/2-Gbps Fibre Channel on a port is mutually exclusive but can be selected for any of the 20 physical ports per module

| Switch | Maximum number of Fibre Channel ports | Maximum number of Ethernet ports |
|---|---|---|
| Cisco Nexus 6001P | 0 | 48 fixed 1/10 Gigabit Ethernet SFP+ and 4 fixed 40 Gigabit Ethernet QSFP+ports |
| Cisco Nexus 6004EF | Up to 160 through the use of 8 expansion modules | Up to 96 40 Gigabit Ethernet and FCoE ports through the use of 8 expansion modules |

| Expansion | Nexus 6004EF |
|---|---|
| 12-port 40 Gigabit Ethernet and FCoE expansion module | yes |
| 20-port unified expansion module | yes |

| Switch | Redundant/ hot-swappable power | Redundant/ hot-swappable cooling | Redundant control processor | Nondisruptive code activation | Expansion module support | Protocol support |
|---|---|---|---|---|---|---|
| Nexus 6001P | Yes/Yes | Yes/Yes | No | Yes | No | FCOE CEE |
| Nexus 6004EF | Yes/Yes | Yes/Yes | No | Yes | Yes | FC FCoE CEE |



https://t.me/learningnets

## Models

For the latest information on supported C-series FCoE CN switches and firmware versions:

1. Go to **http://www.hpe.com/support/c-seriesSwitchesDownloads/**.
2. From the list of C-series switches, select **HP StoreFabric C-series Fabric Switches**.
3. From the list of StoreFabric C-series Fabric Switches, select **Cisco Nexus 5000 Converged Network Switches**. The webpage for Cisco Nexus 5000 Converged Network Switches displays.
4. Under **Download Options**, click the **drivers, software & firmware** link.
5. Select your switch.

### Table 16: Nexus switches for NX-OS

| Switch | Maximum number of Fibre Channel ports | Maximum number of Ethernet ports |
|---|---|---|
| Cisco Nexus 5548UP | 48 with expansion module | 48 with expansion module |
| Cisco Nexus 5596UP | 96 with expansion modules | 96 with expansion modules |

### Table 17: Nexus expansion module support matrix

| Expansion module | Switch | |
|---|---|---|
| | Nexus 5548UP | Nexus 5596UP |
| Nexus 5500 8-port FC (8/4/2/1G) + 8-port Eth/FCoE Module | Yes | Yes |
| Nexus 5500 16-port Unified Port Expansion 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)



https://t.me/learningnets

- Layer 2 multipathing (hardware capable)
- Convergence and compatibility with existing FC fabrics
- High availability (see C-series Nexus 5000 series and Cisco Nexus 5500 series switch high-availability feature comparison table)

**Table 18: C-series Cisco Nexus 5500 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 5548UP | Yes/Yes | Yes/Yes | No | No | Yes | FCFCoECEE compliant DCBX |
| Nexus 5596UP | Yes/Yes | Yes/Yes | No | No | Yes | |

### 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** on page 145.

### Operating systems and storage products

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

# Boot from SAN support

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.

For information about configuring boot from SAN, see the *HPE Boot from SAN Configuration Guide* available at **www.hpe.com/info/BootFromSan**. Under **Features**, click **Boot from SAN documentation**.

# Multipathing software support

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



https://t.me/learningnets

# Fabric infrastructure rules

Fabric infrastructure rules are presented in these chapters:


- **HPE FlexFabric switches and storage support** on page 103
- **B-series switches and fabric rules** on page 105
- **C-series switches and fabric rules** on page 137
- **H-series switches and fabric rules** on page 151
- **SAN fabric connectivity and switch interoperability rules** on page 163




https://t.me/learningnets

# HPE FlexFabric switches and storage support

This chapter describes the HPE FlexFabric switch products that HPE supports in storage configurations. It describes the following topics:


- **HPE FlexFabric converged switches** on page 103
- **Features and storage support** on page 104


## HPE FlexFabric converged switches

FlexFabric Switches:

- Support multiple storage networking protocols — FCoE, Fibre Channel, iSCSI
- Can be used as leaf or spine switches. When configured in a leaf-spine fabric topology, a leaf or Top-Of-Rack (ToR) switch typically connects to servers and storage; a spine switch typically connects to other switches in the network.

  **NOTE:** In converged-Ethernet network terminology, single-layer switch topologies are often referred as a leaf or Top-Of-Rack network topologies. Multi-hop switch topologies are often referred as a leaf-spine network topology.

- Support multi-switch networks with up to 30 switches and 7 hops (FCoE) or 3 hops (Fibre Channel).
- Convergence — Many FlexFabric switches have 10GbE, 40GbE or converged ports (10GbE/4/8Gb Fibre Channel). The 10GbE and 40GbE ports can be used as Ethernet data ports, IRF ports or FCoE ports. The converged ports can be used as Ethernet data ports, IRF ports, or 4 or 8Gb native Fibre Channel ports.
- Intelligent Resilient Framework (IRF) support — An IRF domain simplifies management by combining multiple physical switches into one virtual switch. An IRF fabric appears as one node or virtual switch and is accessible as a single IP address on the network. You can use this IP address to log in at any member device to manage all the members of the IRF domain.
- Support plug-and-play compatibility.
- Have two hot-swappable power supplies or one replaceable power supply depending on model.
- Scalability — Leaf-spine switches can be interconnected in multi-hop configurations to increase the size of the network. The 40GbE ports on many of the switches can be split into (4) individual 10GbE ports using splitter cables, increasing the total number of available 10GbE ports.

### Model numbering

FlexFabric switches use a numbering scheme based on switch families or switch series, such as 5900 series or 12900 series.

The FlexFabric switch series numbers are described below:

- Spine or Core Switches — 5950, 5940, 5930, 7900, 12900 Series
- Leaf or Edge Switches — 5700, 5900, 5930, 5950, 5940 series
- Blade Switches — 6120, 6125 and 6127



https://t.me/learningnets

## Model naming

The naming convention typically begins with the series number, followed by a letter designation, then a number to indicate port count, and an interface-type description.

## Switch models

For the latest information about FlexFabric switch models, refer to the Hewlett Packard Enterprise Networking website at **http://www.hpe.com/networking**.

# Features and storage support

For more detailed information about FlexFabric switch features, storage support, including 3PAR Remote Copy over long distance, and configuration/setup, see the *HPE FlexFabric Switch Series Storage Solutions Configuration Guide*, available at **http://www.hpe.com/storage/spock** listed under **Design Guides/Other Configuration Guides**.

The following FlexFabric switches are currently supported for use with HPE storage systems:

- 5700
- 5900AF
- 5900CP
- 5930 32QSFP+
- 5930 modular 2 slot and 4 slot
- 5950 modular 4 slot and 5950 fixed port switch
- 5940 modular 2 slot and 4 slot
- 7900
- 12900

Hewlett Packard Enterprise supports the use of FlexFabric converged switches with a number of HPE storage products, including 3PAR StoreServ, StoreVirtual, XP, EVA, and MSA storage systems. For the latest information on supported storage systems, FlexFabric switch models and firmware versions, see the SPOCK website at **http://www.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access.



https://t.me/learningnets

# 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 Router Blade, and fabric rules for building B-series fabrics. It describes the following topics:


- **B-series Fibre Channel switches, HPE StoreFabric SN4000B SAN Extension Switch, HPE StoreFabric SN8600B 32Gb SAN Extension Blade, 1606 Extension SAN Switches, and MP Router Blade** on page 105
- **Fibre Channel switch fabric rules** on page 118


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** on page 163.

## B-series Fibre Channel switches, HPE StoreFabric SN4000B SAN Extension Switch, HPE StoreFabric SN8600B 32Gb SAN Extension Blade, 1606 Extension SAN Switches, and MP Router Blade

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, 12, 16, 24, 32, 36, 40, 48, 64, 80, 96, 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 configuration of a 10 Gb/s port on a 32 Gb or 16 Gb switch or director blade).
- Support Fibre Channel routing, see
  - **1606 Extension SAN Switch and DC Dir Switch MP Extension Blade fabric rules** on page 127
  - **HPE StoreFabric SN8600B 32Gb SAN Extension Blade and SN4000B SAN Extension Switch fabric rules** on page 125
- Support FCIP through the 1606 Extension SAN Switch, DC Dir Switch MP Extension Blade, StoreFabric SN4000B SAN Extension Switch, MP Router Blade, and HPE SN8600B 32Gb SAN Extension Blade, see **SAN extension** on page 274.
- Support data at rest encryption through the encryption switch and encryption blade.
- Support in-flight encryption and/or compression on 32 Gb and 16 Gb switch or director blade ISL ports.

Switch blade support is as follows:



https://t.me/learningnets

- HPE StoreFabric SN8600B 8-slot Power Pack+ Director Switch and HPE StoreFabric SN8600B 4-slot Power Pack+ Director Switch support the 32Gb 48-port Fibre blade and 32Gb SN Extension Blade.
- SN8000B 8-Slot SAN Backbone Director and SN8000B 4-Slot SAN Director support the 64-port 8 Gb/s Fibre Channel blade, 32, 48, and 64 port 16 Gb/s Fibre Channel blades, 32 and 48-port enhanced 8 Gb/s Fibre Channel blades, encryption blade, 10/24 FCoE 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, MP Router Blade, encryption blade, 10/24 FCoE blade, and the DC Dir Switch MP Extension Blade.

The B-series switches offer:

- High availability
- Scalability
- Cost efficiency

## Model naming

The B-series Fibre Channel switches are typically called 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 B-series embedded switches for c-Class BladeSystem server environment are called the Brocade 16Gb SAN Switch for BladeSystem c-Class and the Brocade 8Gb SAN Switch for BladeSystem c-Class. The B-series Fibre Channel routers is called the Multi-protocol Router Blade (MP Router Blade). The 8 Gb SAN extension products are called the 1606 Extension SAN Switch and DC SAN Director Multi-protocol Extension Blade. The 16 Gb SAN extension product is called the StoreFabric SN4000B SAN Extension Switch. The 32 Gb SAN extension blade is called the HPE StoreFabric SN8600B 32Gb SAN Extension Blade.

The current HPE model naming convention has the B-series Storage Networking switches named SN**xxxx**B where SN stands for Storage Networking switch and B for B-series. The **xxxx** numbering is 8000 (Director Class), and 6000, 4000, and 3000 (Fabric Class).

The 32 Gb/s, 16 Gb/s, 8 Gb/s, and 4 Gb/s switch ports autonegotiate the signaling speed. For example, when you connect a 16 Gb/s or 8 Gb/s port directly to a 4 Gb/s port, both ports run at 4 Gb/s in each direction. If the ports are not directly connected, the fabric switch that connects the ports determines the signaling speed.

## Switch models

Table B-series Fibre Channel switches and B-series Fibre Channel switches and routers describe the B-series Fibre Channel switches and routers currently available from HPE (32 Gb/s, 16 Gb/s and 8 Gb/s B-series Fibre Channel switches).

Hewlett Packard Enterprise supports all B-series Fibre Channel switches in a fabric if you:

- Use the recommended firmware versions. All like switch models must use the same version when configured in a single-fabric or multi-fabric SAN. 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** on page 118.



https://t.me/learningnets

For the latest information on supported B-series Fibre Channel switches and firmware versions, see the SPOCK website at **http://www.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access.

After logging in, click **Switches** under **Other Hardware** in the last navigation panel of the window to access the Fibre Channel Switch Streams. Click on the **B-Series FC Switch Connectivity Stream** to open the document.

To download the latest firmware from the Hewlett Packard Enterprise Support Center website:

1. Go to **http://www.hpe.com/**.
2. Select **Support** from the drop-down menu in the top right corner of the home page.
3. Under Product Support, click **HPE Support Center**.
4. Log into the HPE Support Center with your HPE Passport information.
5. Enter your B-series switch (i.e. SN6600B) into the search box, and you will be presented with a list of models associated with this switch. Click on the link for your model.
6. Click **Drivers & Software**.
7. Select **Firmware (Entitlement required) Version: vX.X.X**
8. To read the Firmware Release Notes, click on the **Release Notes** link.

### Table 19: B-series Fibre Channel switches

| HPE switch name | Fabric Management | Number of Ports |
|---|---|---|
| SN8600B 8-slot Power Pack+ Director Switch | Network Advisor | 48 to 384 ports |
| SN8600B 4-slot Power Pack+ Director Switch | | 48 to 192 ports |
| SN6600B 32Gb Fibre Channel Switch | | 24 to 64 ports |
| SN6650B 32 Gb Fibre Channel Switch | | 48 to 128 ports |
| SN3600B 32Gb Fibre Channel Switch | | 8 to 24 ports |
| SN8000B 8-Slot SAN Backbone Director Switch | Network Advisor | 32 to 512 |
| SN8000B 4-Slot SAN Director Switch | | 32 to 256 |
| SN6000B 16Gb Fibre Channel Switch | | 24 to 48 |

*Table Continued*



https://t.me/learningnets

| HPE switch name | Fabric Management | Number of Ports |
|---|---|---|
| SN3000B 16Gb Fibre Channel Switch | | 12 to 24 |
| StoreFabric SN6500B 16Gb Fibre Channel Switch | | 48 to 96 |
| Brocade 16Gb SAN Switch for BladeSystem c-Class | | 1 to 16 internal<br>1 to 12 external |
| Brocade 16Gb Fibre Channel SAN Switch Module for HPE Synergy | | 1 to 12 internal<br>1 to 24 external |
| StorageWorks Encryption SAN Switch | | 32 |
| StorageWorks Encryption FC Blade | Network Advisor | 16 |
| Brocade 8Gb SAN Switch for BladeSystem c-Class | | 1 to 16 internal<br>1 to 8 external |
| StorageWorks 8/8 SAN Switch | | 8 to 24 |
| StorageWorks 8/24 SAN Switch | | 16 to 24 |

**Table 20: B-series Fibre Channel switches and routers**

| HPE switch name | Fabric Management | Number of ports |
|---|---|---|
| StoreFabric SN4000B SAN Extension Switch | Network Advisor | 24 at 16 Gb/s Fibre Channel<br>16 at 1-GbE or 10-GbE FCIP<br>2 at 40-GbE FCIP |
| StorageWorks 1606 Extension SAN Switch | | 16 at 8 Gb/s Fibre Channel<br>6 at 1-GbE FCIP |
| SN8600B 32Gb SAN Extension Blade | | 16 at 32 Gb/s Fibre Channel<br>16 at 1/10 GbE FCIP<br>2 at 40GbE FCIP |
| HPE Multi Protocol Extension Blade | | 12 at 8 Gb/s Fibre Channel<br>10 at 1-GbE FCIP<br>2 at 10-GbE FCIP |
| HPE StoreFabric Enhanced Multiprotocol Extension Blade[1] | | |


https://t.me/learningnets

[1] IPsec support for the XGEO port has been added to this blade enabling creation of IP security enabled FCIP tunnels. FOS FOS 7.1.0a or later is required to take advantage of the enhanced IPsec capability. This feature is not enabled when running earlier versions of FOS firmware.

### Table 21: B-series legacy Fibre Channel switches and routers

| Legacy HPE switch name | Fabric Management | Number of ports |
|---|---|---|
| HPE Multi-protocol Router Blade for B-Series | Network Advisor | 16 at 1, 2, or 4 Gb/s Fibre Channel<br>2 at 1-GbE FCIP |
| StorageWorks 400 Multi-Protocol Router | | |
| StorageWorks DC SAN Backbone Director | | 16 to 512 |
| StorageWorks 8/40 SAN Switch | | 24 to 40 |
| StorageWorks 8/80 SAN Switch | | 48 to 80 |
| StorageWorks DC04 SAN Director | | 16 to 256 |

## 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 16Gb Fibre Channel SAN Switch Module for HPE Synergy, Brocade 16Gb SAN Switch for BladeSystem c-Class, Brocade 8Gb SAN Switch for BladeSystem c-Class, StorageWorks SAN Switch 8/24, 8/40 (with 40 ports enabled), SN3000B, SN3600B, SN6000B, SN6600B FC switches.

- **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 10 km 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** on page 163.
- **HPE Network Advisor** — Centralizes fabric management through a host-based application. This is available free for dual-fabric management (Network Advisor Professional), or for purchase for multiple-fabric management and increased device count (Network Advisor Professional Plus and Network Advisor Enterprise).
- **ISL Trunking** — Combines multiple links between switches to form a single, logical ISL with a total bandwidth of 256 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, 10-GbE, or 40-GbE interfaces to provide a high-bandwidth FCIP tunnel and



https://t.me/learningnets

failover resiliency. 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 SN8600B 32Gb SAN Extension Blade, StoreFabric SN4000B SAN Extension Switch, 1606 Extension SAN Switch, and the DC Dir Switch MP Extension Blade. For more information, see **HPE StoreFabric SN4000B SAN Extension Switch** on page 306 and **B-series 1606 Extension SAN Switch and DC Dir Switch MP Extension Blade** on page 310.

- **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 on the 1606 Extension SAN Switch and DC Dir Switch MP Extension Blade or SN8600B 32Gb SAN Extension Blade and StoreFabric SN4000B.

- **NPIV** — See **SAN fabric connectivity and switch interoperability rules** on page 163.

  The following B-series Fibre Channel switches are NPIV compliant:

  - **32Gb/s:** SN8600B 8-slot Power Pack+ Director Switch, SN8600B 8-slot Power Pack+ Director Switch, SN6600B 32Gb Fibre Channel Switch and SN3600B 32Gb Fibre Channel Switch.
  - **16 Gb/s:** SN8000B 8-Slot SAN Backbone Director, SN8000B 4-Slot SAN Director, SN6000B 16Gb Fibre Channel Switch, StoreFabric SN4000B SAN Extension Switch; Brocade 16Gb SAN Switch for BladeSystem c-Class, Brocade 16Gb SAN Switch Module for HPE Synergy
  - **8 Gb/s:** Brocade 8Gb SAN Switch for BladeSystem c-Class; SAN Switch 8/8, 8/24, 8/40, and 8/80; DC SAN Backbone Director; DC04 SAN Director; Encryption SAN Switch; Encryption FC Blade

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

- **Fabric Vision** — Fabric Vision technology provides a breakthrough hardware and software solution that helps simplify monitoring, maximize network availability, and dramatically reduce costs. Featuring innovative monitoring, management, and diagnostic capabilities, Fabric Vision technology enables administrators to avoid problems before they impact operations, helping their organizations meet SLAs. Fabric Vision technology features for storage extension management include:

  - Monitoring and Alerting Policy Suite (MAPS): Provides a pre-built, policy-based threshold monitoring and alerting tool that proactively monitors storage extension network health based on a comprehensive set of metrics at tunnel, circuit, and QoS layers. Administrators can configure multiple fabrics at one time using pre-defined or customized rules and policies for specific ports or switch elements.
  - Fabric Performance Impact (FPI) Monitoring: Uses pre-defined thresholds and alerts in conjunction with MAPS to automatically detect and alert administrators to severe levels of latency and identifies slow drain devices that might impact the network. This feature uses advanced monitoring



https://t.me/learningnets

capabilities and intuitive MAPS dashboard reporting to indicate various latency severity levels, pinpointing exactly which devices are causing or impacted by a bottlenecked port.

- Dashboards: Provides integrated dashboards that display an overall SAN health view, along with details on out-of-range conditions, to help administrators easily identify trends and quickly pinpoint issues occurring on a switch or in a fabric.

- **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 SN8000B 8-Slot SAN Backbone Director Switch, the SN8000B 4-Slot SAN Director Switch, the SN6000B 16Gb FC Switch, and the SN6500B 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 DC SAN Director Multiprotocol Extension Blade and the StoreFabric SN4000B SAN Extension Switch
  - Multi-Gigabit circuits which allows more than 1 Gb/s to be configured on a single FCIP Circuit
  - 10 GigE Adaptive Rate Limiting (ARL) on DC SAN Director Multiprotocol Extension Blade and the StoreFabric SN4000B SAN Extension Switch
  - 40GbE support on the StoreFabric SN4000B SAN Extension Switch
  - Non-disruptive firmware upgrade on the StoreFabric SN4000B SAN Extension Switch

- **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 the StoreFabric SN4000B SAN Extension Switch, the DC SAN Director Multiprotocol Extension Blade, and the 1606 SAN Extension Switch.
- **XISL support on HPE DC SAN Director Multiprotocol Extension Blade** — Provides the ability to use VE ports as XISLs on DC SAN Director Multiprotocol Extension Blade only in both 1 Gb/s and 10 Gb/s modes.
- **InBand Management on HPE StoreFabric SN4000B SAN Extension Switch, HPE DC SAN Director Multiprotocol Extension Blade, and HPE 1606 SAN Extension Switch** — Allows you to



https://t.me/learningnets

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 (HPE StorageWorks Encryption SAN Switch/HPE 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 DC SAN Director 10/24 FCoE Blade in StorageWorks DC SAN Backbone Director Switch/HPE StorageWorks DC04 SAN Director Switch and in slot 1 of the SN8000B 8-Slot SAN Backbone Director with FOS v7.3.x
- **ICLs (HPE SN8000B 8-Slot and HPE SN8000B Directors)** — Provides an additional 1 Tb/s of bandwidth and does not consume usable ports with an additional full duplex connection. ICL connections operate as hardware trunked ISLs. A single license will enable all the ICL ports on the SN8000 4-Slot Director. However, a single license enables only half the available ICL ports on the SN8000B 8-Slot Director, and requires an additional license to enable the remaining ICL ports.
- **Enterprise ICL** — Enables the connection of more than four SN8000B SAN Director platforms within a single fabric using optical ICLs to deliver massive scalability. This license augments but does not replace existing licensed ICL support (ICL licenses are still required to enable ICL ports) for larger ICL-based topologies of 5 to 10 SN8000B SAN Director chassis per fabric. This license is required to be installed on every SN8000B SAN Director chassis with ICLs in a fabric when an ICL-based topology exceeds the four chassis limit. This license is recognized/displayed only when operating with Fabric OS 7.0.1 and later. There is a separate license for each of the SN8000B SAN Directors, the SN8000B 8-Slot Director, and the SN8000B 4-Slot Director.

Features of Fabric OS 8.0.1 follow (see the *HPE B-series Fabrics OS 8.x Release Notes* and appropriate Fabric OS administrators guide for configuration restrictions and more information):

- **ICLs (HPE SN8600B 8-Slot and HPE SN8600B 4-Slot Directors)** — Provides an additional 4 Tb/s of bandwidth on the SN8600B 8-Slot Director and 2 Tb/s of bandwidth on the SN8600B 4-Slot Director and does not consume usable ports with an additional full duplex connection. For the SN8600B 8-Slot SAN Director, each license has a capacity of 32 links, which activates 16 QSFP ports across both core blades. For the SN8600B 4-Slot SAN Director, each license has a capacity of 16 links, which activates 8 QSFP ports across both core blades. Each ICL license has the capacity of enabling half of the ICL bandwidth. Supports connectivity between SN8600B and SN8000B directors.
- **Fabric Vision Enhancements**
  - **Fabric Performance Impact (FPI) monitoring** — FOS v8.0.1 enhances FPI monitoring to include the FPI monitoring rules and thresholds in the MAPS default base policy. Switches running FOS v8.0.1 without a Fabric Vision license are monitored by MAPS for congestion and slow drain device conditions.
- **Monitoring and Alerting Policy Suite (MAPS) enhancements**
  - **Monitor IO Insight Metrics** — Adds the support in MAPS to monitor the IO Insight metrics for initiator-target or initiator-target-LUN flow monitors that are created with the Flow Vision feature on


https://t.me/learningnets

Brocade Gen 6 platforms. This monitoring allows administrators to be alerted if application or storage device level IO performance and latency measures drop below predetermined thresholds such as those mandated by service level agreements.

- **Monitor 32Gbps Optics** — Adds monitoring of supported 32 Gbps SWL SFP, 32 Gbps LWL SFP, and 4x32 Gbps SWL QSFP optics on Gen 6 platforms to the default MAPS policy.
- **Monitor Security Certificate Expiration** — Adds monitoring the validity period of digital security certificates present on Brocade switches. This provides administrators expiry notification before any security certificates actually expire.
- **Monitor Gigabit Ethernet Ports** — Adds the ability to monitor link layer error statistics for Gigabit Ethernet ports on extension platforms.
- **Monitor FCIP Tunnel and Tunnel QoS** — Extends MAPS monitoring of FCIP tunnels and the tunnels at each QoS level to the SX6 blade. Bandwidth utilization and state change are monitored for FCIP tunnels, while utilization and packet loss are monitored for FCIP tunnels for each QoS level. The pre-defined thresholds and rules for SX6 blades are the same as those for the Brocade 7840 platform in FOS v7.4.
- **Monitor IP Extension Circuit and Tunnel** — Adds MAPS monitoring for IP circuits and tunnels on the Brocade 7840 extension platform and SX6 blades. IP Extension monitoring provides advanced notification of WAN performance and latency problems and helps with troubleshooting IP storage traffic that has been extended through WANs.
- **Monitor System Air Flow** — Adds MAPS monitoring for system airflow directions to provide alerts if the air flow directions of the fan or power supply FRUs in the hardware platforms are installed incorrectly. This avoids potential disruptive shutdown of hardware platforms resulting from mismatched air flow directions.
- **Monitor Initiator and Target Device Ratio** — Adds the ability to monitor the ratio between initiator and target devices zoned together across entire zoning configuration. Initiator-to-target or target-to-initiator ratios higher than best practices have the potential of causing slow IO response time that may impact fabric performance. This monitoring can alert administrators to this condition, enabling them to take corrective action.

- **Flow Vision enhancements**

  - **Flow Monitor support for IO Insight** — Extends the Flow Monitor feature to support the IO Insight capability built in to B-series Gen 6 platforms. Flow Monitor provides administrators the capability to non-disruptively and non-invasively obtain key SCSI IO performance and latency metrics for initiator-target or initiator-target-LUN flows.
  - **Flow Monitor on VE port** — Supports Flow Monitor static and learning flow support on VE ports on the Brocade 7840 and the SX6 blade. The Flow Monitor support on VE ports is equivalent to those provided on E_Ports available in FOS v7.4.
  - **SRR Frame Type** — Supports a new frame type in Flow Vision, the Sequence Retransmission Request (SRR) frame, which is used in FC-Tape backup for error recovery. This frame type support allows administrators to see that the FC-Tape devices are in error recovery mode, which increases application response time.
  - **Fabric Flow Dashboard Enhancements** — Enhances the Flow Dashboard with MAPS history data for ports and back-end ports and identification of ports that have been fenced, decommissioned, or quarantined by MAPS actions.

- **FCIP and IP Extension**

  - **Extension Non-disruptive Firmware Download on the SN8600B 32Gb SAN Extension Blade** — Supports non-disruptive firmware download for the FCIP extension traffic on the SN8600B 32Gb



https://t.me/learningnets

SAN Extension Blade in an StoreFabric SN8600B 8-slot Director Switch or StoreFabric SN8600B 4-slot Director Switch. Non-disruptive firmware download is performed in parallel on the maximum number of SN8600B 32Gb SAN Extension Blade supported in the director.

Table B-series Fibre Channel switch high-availability feature comparison provides a comparison of the high-availability features for B-series Fibre Channel switches.

### Table 22: 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 |
|---|---|---|---|---|---|---|---|
| SN8600B 8-slot Director Switch | Yes/Yes | Yes/Yes | Yes | Yes | Yes | Yes | Yes |
| SN8600B 4-slot Director Switch | Yes/Yes | Yes/Yes | Yes | Yes | Yes | Yes | Yes |
| StoreFabric SN6600B 32 Gb Fibre Channel Switch | Yes/Yes | Yes/Yes | No | No | Yes | Yes | No |
| StoreFabric SN6650B 32 Gb Fibre Channel Switch | Yes/Yes | Yes/Yes | No | No | Yes | Yes | No |
| SN3600B 32Gb Fibre Channel Switch | Yes/Yes | Yes/Yes | No | No | Yes | Yes | No |
| SN8000B 8-Slot SAN Backbone Director Switch | Yes/Yes | Yes/Yes | Yes | Yes | Yes | Yes | Yes |
| SN8000B 4-Slot SAN Director Switch | Yes/Yes | Yes/Yes | Yes | Yes | Yes | Yes | Yes |
| SN6000B 16Gb Fibre Channel Switch | Yes/Yes | Yes/Yes | No | No | Yes | Yes | No |
| SN3000B 16Gb Fibre Channel Switch | Yes/Yes | Yes/Yes | No | No | Yes | Yes | No |
| StoreFabric SN6500B 16Gb Fibre Channel Switch | Yes/Yes | Yes/Yes | No | No | Yes | Yes | No |
| StoreFabric SN4000B SAN Extension Switch | Yes/Yes | Yes/Yes | No | No | Yes | Yes | No |
| Encryption SAN Switch | Yes/Yes | Yes/Yes | No | No | Yes | N/A | No |

*Table Continued*



https://t.me/learningnets

| 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 |
|---|---|---|---|---|---|---|---|
| Brocade 16Gb SAN Switch Module for HPE Synergy | N/A | N/A | No | No | Yes | Yes | No |
| Brocade 16Gb SAN Switch for BladeSystem c-Class | N/A | N/A | No | No | Yes | Yes | No |
| Brocade 8Gb SAN Switch for BladeSystem c-Class | N/A | N/A | No | No | Yes | Yes | No |
| 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 |
| StorageWorks 1606 Extension SAN Switch | Yes/Yes | Yes/Yes | No | No | Yes | No | No |



https://t.me/learningnets

## Usage

**Table 23: 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 |
|---|---|---|---|---|---|
| SN8600B 8-Slot SAN Director Switch | Excellent | | | | |
| SN8600B 4-Slot SAN Director Switch | | | | | |
| SN8000B 8-Slot SAN Backbone Director Switch | | | | | |
| SN8000B 4-Slot SAN Director Switch | | | | | |
| DC SAN Backbone Director | | | | | |
| DC04 SAN Director | | | | | |
| StoreFabric SN6500B 16Gb FC Switch | Excellent | | Very good | Good | Not recommended |
| SAN Switch 8/80 | | | | | |
| SN6600B 32Gb FC Switch | Excellent | Very good | Good | Not recommended | |
| SN6650B 32Gb FC Switch | | | | | |
| SN3600B 32Gb FC Switch | | | | | |
| SN6000B 16Gb FC Switch | | | | | |
| SN3000B 16Gb FC Switch | | | | | |
| SAN Switch 8/24, 8/40 | | | | | |

*Table Continued*



https://t.me/learningnets

| Model | 1-96 user ports | 97-224 user ports | 225-500 user ports | 501-728 user ports | 729-1280 user ports |
|---|---|---|---|---|---|
| StoreFabric SN4000B SAN Extension Switch,<br>SAN Switch 8/8, 1606 Extension SAN Switch | Good | | | Not recommended | |
| Brocade 16Gb SAN Switch Module for HPE Synergy<br>Brocade 16Gb SAN Switch for BladeSystem c-Class<br>Brocade 8Gb SAN Switch for BladeSystem c-Class | Good | Not recommended | | | |



https://t.me/learningnets

**Table 24: 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 |
|---|---|---|---|---|---|
| SN8600B 8-Slot SAN Director Switch | Excellent | | | | |
| SN8600B 4-Slot SAN Director Switch | | | | | |
| SN8000B 8-Slot SAN Backbone Director Switch | | | | | |
| SN8000B 4-Slot SAN Director Switch | | | | | |
| DC SAN Backbone Director | | | | | |
| DC04 SAN Director | | | | | |
| SN6600B 32Gb FC Switch | Excellent | | | | |
| SN6650B 32 Gb FC Switch | | | | | |
| StoreFabric SN6500B 16Gb FC Switch | | | | | |
| SN6000B 16Gb FC Switch | | | | | |
| SAN Switch 8/40, 8/80 | | | | | |
| SN3600B 32Gb Fibre Channel Switch | Excellent | Very good | | Good | |
| Brocade 16Gb SAN Switch Module for HPE Synergy | | | | | |
| SN3000B 16Gb FC Switch | | | | | |
| SAN Switch 8/8, 8/24 | | | | | |
| Brocade 16Gb SAN Switch for BladeSystem c-Class | | | | | |
| Brocade 8Gb SAN Switch for BladeSystem c-Class | | | | | |
| StoreFabric SN4000B SAN Extension Switch | | | | | |
| 1606 Extension SAN Switch | | | | | |

# 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:



https://t.me/learningnets

- Use the HPE default settings for all current HPE switches.
- Use the legacy HPE default settings if the fabric contains only legacy HPE switches.
- See **Third-party switch support** on page 175.

**NOTE:** To download the HPE Default Config Files from the Hewlett Packard Enterprise Support Center website:

1. Go to the Hewlett Packard Enterprise Support Center home page at **http://www.hpe.com/support/hpesc**.
2. From the left navigation panel, under **DOWNLOAD OPTIONS**, click **Drivers, Software & Firmware**. The **Select a Product page** appears.
3. To select an appropriate product, under **All HP products**, click **Storage**.
4. Under **Storage**, click **Storage Networking**.
5. Under **Storage Networking**, click **StoreFabric B-series Switches**.
6. Under **StoreFabric B-series Switches**, click appropriate switch product.
7. Select the appropriate switch model from the list.
8. Click **Cross operating system (BIOS, Firmware, Diagnostics, etc.)**.
9. From the Software table, click on the **HPE Default Config Files for B-Series FC Switches** or the **Download** button to obtain the Default Config Files.

This section describes the following topics:


- **Operating systems and storage products** on page 120
- **B-series Fibre Channel switch and fabric rules** on page 120
- **B-series Encryption Switch fabric rules** on page 124
- **HPE StoreFabric SN8600B 32Gb SAN Extension Blade and SN4000B SAN Extension Switch fabric rules** on page 125
- **1606 Extension SAN Switch and DC Dir Switch MP Extension Blade fabric rules** on page 127
- **Zoning limits and enforcement** on page 133
- **Zoning guidelines for B-series Fibre Channel switches** on page 134




https://t.me/learningnets

## Operating systems and storage products

### Table 25: B-series Fibre Channel switch operating system and storage system support

| Operating systems | Storage products |
|---|---|
| • HP-UX<br>• OpenVMS<br>• Tru64 UNIX<br>• Apple Mac OS X<br>• Citrix Xen<br>• IBM AIX<br>• Linux<br>• Microsoft Windows<br>• Oracle Solaris<br>• VMware ESX | • MSA 2040/2042/1040/2050/2052/1050<br>• P4300 FC<br>• HPE StoreVirtual 3200 FC/4000 FC<br>• P6550/P6500/P6350/P6300<br>• P9500<br>• XP24000/20000/12000/10000<br>• 3PAR StoreServ 10000/7000/8000/9000/20000<br>• HPE P9000 (XP) Continuous Access/P6000 Continuous Access<br>See **HPE P6000 Continuous Access SAN integration** on page 247 |

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

See **Heterogeneous server rules** on page 179 and the following storage system rules chapters:

- **MSA storage system rules** on page 227
- **HPE StoreVirtual 3200 and 4000 Storage System Rules** on page 234
- **P6000 storage system rules** on page 238
- **P9000/XP storage system rules** on page 252
- **HPE 3PAR StoreServ storage rules** on page 260

## 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 P9000 (XP) Continuous Access and P6000 Continuous Access (formerly HPE StorageWorks Continuous Access EVA) configurations. However, additional rules apply to P6000 Continuous Access implementations. For more information, see the *HPE P6000 Continuous Access Implementation Guide*.



https://t.me/learningnets

## Table 26: B-series Fibre Channel switch fabric rules

| Rule number | Description |
|---|---|
| 1 | **FOS v7.0.x and later:** A maximum of 56 switches and 6000 virtual or physical connections in a single fabric with for fabrics that contain only the following platforms: DC SAN Backbone Director Switch, DC04 SAN Director Switch, SN8000B 8-Slot SAN Backbone Director Switch, SN8000B 4-Slot SAN Director Switch, SN6500B 16Gb FC Switch, SN6000B 16Gb FC Switch, SN3000B 16Gb FC Switch, 8/80 SAN Switch, 8/40 SAN Switch, 8/8 SAN Switch, 1606 Extension SAN Switch, StoreFabric SN4000B SAN Extension Switch, or the Encryption SAN Switch.<br><br>A maximum of 56 switches and 4,096 virtual or physical connections in a single fabric that contains platforms not listed above.<br><br>**FOS v8.x:** A maximum of 56 switches and 6000 virtual or physical connections in a single fabric with for fabrics that contain only the following platforms: SN8600B 8-Slot SAN Director Switch, SN8600B 4-Slot SAN Director Switch, SN8000B 8-Slot SAN Backbone Director Switch, SN8000B 4-Slot SAN Director Switch, SN6600B FC Switch, SN3600B 32Gb Fibre Channel Switch, SN6500B 16Gb FC Switch, SN6000B 16Gb FC Switch, SN3000B Switch, or StoreFabric SN4000B SAN Extension Switch. |
| 2 | 8 Gb/s switches and ARB(FF)/ARB(FF) mode setting — Hewlett Packard Enterprise 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 `portcfgfillword` command. For more information, see the *Fabric OS Command Reference Manual supporting Fabric OS v6.4* and the FOS release notes.<br><br>**NOTE:** SN8600B 8-Slot SAN Director Switch, SN8600B 4-Slot SAN Director Switch, SN6600B 32Gb FC Switch, SN8000B 8-Slot SAN Backbone Director Switch, SN8000B 4-Slot SAN Director Switch, HPE StoreFabric SN6500B 16Gb FC Switch, HPE SN6000B 16Gb FC Switch, or the HPE SN3000B 16Gb FC Switch automatically selects the proper fill word. |
| 3 | B-series switches and routers supported routing modes with Continuous Access products:<br><br>• P9000 (XP) Continuous Access — Supported with port-based routing (aptpolicy = 1) or exchange-based routing (aptpolicy = 3), see HPE P9000 (XP) Continuous Access FCIP gateway support table.<br>• P6000 Continuous Access — Supported with port-based routing (all XCS versions) or exchange-based routing (XCS 09534000 or later), see Network requirements for long-distance IP gateways with XCS 11x, XCS 10x, or XCS 09x table.<br><br>Routing policy — Hewlett Packard Enterprise recommends that you use the same routing policy in all switches and routers within a fabric. |

*Table Continued*



https://t.me/learningnets

| Rule number | Description |
|---|---|
| 4 | Access Gateway (AG) requirements are:<br>• AG mode for SN3600B 32 Gb Fibre Channel Switch required firmware 8.1.0a (or later).<br>• AG mode for SN6600B 32Gb FC Switch requires firmware 8.0.1b (or later).<br>• AG mode for Brocade 16Gb SAN Switch Module for HPE Synergy requires firmware 8.0.1b (or later)<br>• AG mode for Brocade 16Gb/16 SAN Switch for BladeSystem c-Class requires firmware 7.1.0_blv or later.<br>• AG mode for SN3000B 16Gb FC Switch requires firmware 7.0.1 (or later).<br>• AG mode for SN6000B 16Gb FC Switch requires firmware 7.0.0a (or later).<br>• AG mode for Brocade 8Gb SAN Switch for BladeSystem c-Class and StorageWorks SAN Switch 8/8 or 8/24 (with 24 ports enabled) requires firmware 6.1**x** (or later).<br>• AG mode for the SAN Switch 8/40 requires firmware 6.2**x** (or later).<br>• All switches that communicate with Access Gateway switches require firmware 5.1.**x** (or later). see **BladeSystem with Brocade Access Gateway mode** on page 184. |
| 5 | All B-series Fibre Channel switches and routers must use the HPE default setting of `dlsSet` for non-Continuous Access configurations. P6000 Continuous Access configurations running XCS version 09534000 (or later) have two protocol options — Fibre Channel protocol or HPE SCSI Fibre Channel protocol. The settings for DLS and IOD functionality on the B-series switches should be configured as follows:<br>Fibre Channel protocol:<br>• `dlsreset` (DLS disabled)<br>• `iodset` (IOD enabled)<br>SCSI Fibre Channel protocol:<br>• `dlsset` (DLS enabled)<br>• `iodreset` (IOD disabled)<br>**NOTE:** `dlsset` is the default policy when `aptpolicy` 3 (enable exchange-based routing) is enabled on B-series switches.<br>For more information, see the *HPE P6000 Continuous Access Implementation Guide*.<br>For P9000 (XP) Continuous Access, see HPE P9000 (XP) Continuous Access FCIP gateway support table. |

*Table Continued*



https://t.me/learningnets

| Rule number | Description |
|---|---|
| 6 | 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.<br>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.<br>Routing policy — Hewlett Packard Enterprise recommends the same routing policy be used in all switches within a fabric. |
| 7 | 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* section in the appropriate Fabric OS Administrator Guide. |
| 8 | Maximum of seven hops (eight switches) between any two communicating 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).<br>**NOTE:** The 32 Gb/s and 16 Gb/s Switches and port blades do not support FL ports.<br>**NOTE:** VEX_Ports and EX_Ports cannot be connected to the same edge fabric. |
| 10 | For HPE StoreVirtual 3200 FC or StoreVirtual 4000 FC, unique domain IDs are required in the entire Fibre Channel SAN to which the StoreVirtual storage system is connected, including multiple independent fabrics. Overlapping domain IDs (even in non-merged/ independent fabrics) are not supported. |

**NOTE:** Not all topologies can support the maximum port or switch count, see **Topology maximums** on page 35.

**Switch database size**

### Table 27: Database size rules for B-series Fibre Channel switches

| Rule number | Description |
|---|---|
| 1 | **Advanced Zoning** — The maximum zoning database size is as follows:<br>For fabrics using only the DC SAN Director, DCO4 SAN Director, SN8600B Director platforms, and SN8000B Director platforms, the maximum database size is 2 MB. If any other platform in included in the fabric, then the maximum database size is 1 MB.<br>**NOTE:** Use the `cfgsize` command to determine the size of the zoning database. |

**ISL maximums**



https://t.me/learningnets

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 28: 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 HPE mSFP-LC fiber optic cable connected directly to the LC port of an existing LC-based switch or device to an HPE 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 | • Option 1: Use two mSFP-LC fiber optic cables connected together with an LC-LC coupler.<br>• Option 2 (Brocade): Use an mSFP-mSFP fiber optic cable.<br>**NOTE:** The mSFP-mSFP cable is not available from Hewlett Packard Enterprise.<br>See the DC SAN Director QuickSpec and Brocade *Best Practices Guide: High Density Cable Management Solutions*. |

## B-series Encryption Switch fabric rules

This section describes the fabric rules for the B-series Encryption Switches (StorageWorks Encryption SAN Switch and StorageWorks Encryption FC Blade) and 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.

For more information about the security uses of these switches, see **B-series Encryption Switch and Encryption FC Blade security** on page 401. Hewlett Packard Enterprise strongly recommends that you read the *Brocade Fabric OS Encryption Administrator's Guide* before configuring a fabric using the encryption switch or encryption blade.



https://t.me/learningnets

### Table 29: B-series Encryption Switch fabric rules

| Rule number | Description |
|---|---|
| 1 | Encryption blade can be installed in a SN8000B 8-Slot SAN Backbone Director Switch, 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 physical targets per encryption switch or blade. |
| 6 | There is a maximum of 1,024[1]<br>physical initiators per encryption switch or blade. |

[1] This value is half the value of the actual limit supported in a high-availability cluster.

## HPE StoreFabric SN8600B 32Gb SAN Extension Blade and SN4000B SAN Extension Switch fabric rules

This section describes the fabric rules for the StoreFabric SN8600B 32Gb SAN Extension Blade and SN4000B SAN Extension Switch, 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** on page 120. For information about configuring the StoreFabric SN4000B SAN Extension Switch for FCIP, see **HPE StoreFabric SN4000B SAN Extension Switch** on page 306.

### Table 30: HPE StoreFabric SN8600B 32Gb SAN Extension Blade and SN4000B SAN Extension Switch fabric rules

| Rule number | Description |
|---|---|
| 1 | All configurations must use the default setting for R_A_TOV (10,000 milliseconds) and E_D_TOV (2,000 milliseconds). |
| 2 | Devices connected to edge fabrics can be routed through an LSAN to a backbone fabric. |
| 3 | Direct connection of devices to the StoreFabric SN8600B 32Gb SAN Extension Blade or SN4000B SAN Extension Switch Fibre Channel ports (F_Port) is supported. |
| 4 | Fibre Channel server boot through the StoreFabric SN8600B 32Gb SAN Extension Blade or SN4000B SAN Extension Switch is not supported. |
| 5 | Each StoreFabric SN4000B SAN Extension Switch can attach to one backbone fabric only; multiple routers can attach to the same backbone fabric. |

*Table Continued*



https://t.me/learningnets

| Rule number | Description |
|---|---|
| 6 | A backbone fabric supports a maximum of 512 devices. |
| 7 | You can configure a Fibre Channel port as an F_Port, E_Port, or EX_Port (Fibre Channel routing), and configure a GbE port as a VE_Port (FCIP).<br><br>**NOTE:** VEX_Ports are not supported on the StoreFabric SN4000B SAN Extension Switch or SN8600B 32Gb Extensions Blade. |
| 8 | FCIP is supported between the StoreFabric SN4000B SAN Extension Switches and SN8600B 32Gb Extension Blade. FCIP is not supported in between the StoreFabric SN4000B SAN Extension Switch or SN8600B 32Gb Extensions Blade and the 1606 Extension SAN Switch, DC Dir Switch MP Extension Blade. |
| 9 | FCIP tunnels (virtual E_Ports) pass traffic between the SAN extension devices. The StoreFabric SN4000B SAN Extension Switch and SN8600B 32Gb Extensions Blade each supports:<br><br>• 16 1/10-GbE Ports for FCIP<br>• 2 40-GbE Ports for FCIP.<br><br>Up to 20 VE_Ports per StoreFabric SN4000B SAN Extension Switch are supported for FCIP tunnel configurations. |
| 10 | IPsec is supported on the StoreFabric SN4000B SAN Extension Switch and SN8600B 32GB Extension Blade. |
| 11 | FCIP FastWrite and Fibre Channel FastWrite support:<br><br>FCIP is supported with FCIP FastWrite SCSI write acceleration with P9000 (XP) Continuous Access. FCIP FastWrite is not supported with P6000 Continuous Access.<br><br>**NOTE:** FCIP FastWrite and IPsec are mutually exclusive and cannot be configured simultaneously. |
| 12 | For redundancy, Hewlett Packard Enterprise recommends using a minimum of two EX_Ports (router IFLs) for each fabric connected to a StoreFabric SN4000B SAN Extension Switch or SN8600B 32GB Extension Blade. The number of EX_Ports required is based on the performance requirements for all devices shared by fabrics through the router.<br><br>• 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 `portperfshow` command to monitor router port performance.<br>• For very high bandwidth requirements and applications, use a ratio of 3:1 or 1:1. |
| 13 | Supports routing between B-series secured and nonsecured edge fabrics using advanced security features or ACLs. |

*Table Continued*



https://t.me/learningnets

| Rule number | Description |
|---|---|
| 14 | Supports the following features:<br>• 16 Gb/s on FC ports<br>• FCIP Trunking<br>• Adaptive Rate Limiting<br>• 1/10 GbE ports<br>• 40 GbE ports |
| 15 | The SN8600B 32Gb Extension Blade is only supported in the SN8600B 8-Slot SAN Director Switch and the SN8600B 4-Slot SAN Director Switch. The maximum number of SN8600B 32Gb Extension Blades per SN8600B SAN Director (both 4 and 8-slot) is four. |

#### Scalability rules

The following terms describe the StoreFabric SN4000B SAN Extension Switch 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 StoreFabric SN4000B SAN Extension Switch. Each fabric has one unique translate domain, regardless of how many EX_Ports are connected to the fabric.
- **Backbone fabric** — The StoreFabric SN4000B SAN Extension Switch backbone fabric consists of at least one StoreFabric SN4000B SAN Extension Switch, and includes all B-series switches connected via E_Ports to the StoreFabric SN4000B SAN Extension Switch.
- **Edge fabric** — A fabric that is attached to an EX_Port on the StoreFabric SN4000B SAN Extension Switch.

#### HPE StoreFabric SN4000B SAN Extension Switch hop count

The StoreFabric SN4000B SAN Extension Switch is counted in the same way as a Fibre Channel switch when determining the fabric hop count.

## 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** on page 120. 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** on page 310.



https://t.me/learningnets

### Table 31: 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 SN8000B 8-Slot SAN Backbone Director, SN8000B 4-Slot SAN Director, DC SAN Backbone Director and DC04 SAN Director. The maximum number of DC Dir Switch MP Extension Blades per 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 four. |
| 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.<br>**NOTE:** Edge fabric to backbone fabric routing is supported using the 1606 Extension SAN Switch, DC Dir Switch MP Extension Blade, 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 512 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).<br>**NOTE:** 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, or MP Router Blade, see **B-series 1606 Extension SAN Switch and DC Dir Switch MP Extension Blade** on page 310. |

*Table Continued*



https://t.me/learningnets

| Rule number | Description |
|---|---|
| 11 | FCIP tunnels (virtual E_Ports) pass traffic between the SAN extension devices.<br>• The 1606 Extension SAN Switch supports up to 8 FCIP tunnels.<br>• The DC Dir Switch MP Extension Blade supports up to 20 FCIP tunnels. |
| 12 | IPsec is supported on the 1606 Extension SAN Switch beginning with firmware 6.3.1a. |
| 13 | FCIP FastWrite and Fibre Channel FastWrite support:<br>• FCIP is supported with FCIP FastWrite SCSI write acceleration when using firmware 6.3.0a (or later) with P9000 (XP) Continuous Access. FCIP FastWrite is not supported with P6000 Continuous Access.<br>**NOTE:** FCIP FastWrite and IPsec are mutually exclusive and cannot be configured simultaneously.<br>• Fibre Channel FastWrite — FC FastWrite is currently not supported with the 1606 Extension SAN Switch or DC Dir Switch MP Extension Blade. |
| 14 | For redundancy, Hewlett Packard Enterprise 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.<br>• 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 `portperfshow` command to monitor router port performance.<br>• For very high bandwidth requirements and applications, use a ratio of 3:1 or 1:1. |
| 15 | Supports routing between B-series secured and nonsecured edge fabrics using advanced security features or ACLs. |

*Table Continued*



https://t.me/learningnets

| Rule number | Description |
|---|---|
| 16 | Supports the following features:<br>• 8 Gb/s on FC ports<br>• FCIP Trunking<br>• Adaptive Rate Limiting<br>• 10-GbE ports (DC Dir Switch MP Extension Blade) |
| 17 | The 1606 Extension SAN Switch and DC Dir Switch MP Extension Blade do not support the following features with firmware 6.3.0:<br>• IPsec<br>• DiffServ priorities<br>• VLAN tagging<br>• Third-party WAN optimization<br>• IPv6 addresses for FCIP tunnels<br>• VEX_Ports (DC Dir Switch MP Extension Blade) |

**Scalability rules**

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. 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 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 an SN8000B 8-Slot or 4-Slot Director, DC SAN Backbone Director or a 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. Devices communicating across fabrics through 1606 Extension SAN Switch or DC Dir Switch MP Extension Blade 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:



https://t.me/learningnets

- If you want to add multiple fabrics to a Meta SAN, and one of those fabrics contains an SN8000B 8-Slot or 4-Slot Director, a DC SAN Backbone Director, or a 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 512 local devices in a backbone fabric connected to multiple edge fabrics
- For FCIP, if there are SN8000B 8-Slot or 4-Slot Directors, 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 an SN8000B 8-Slot or 4-Slot Director, a DC SAN Backbone Director, or a DC04 SAN Director, and do not plan to use one
- For FCIP, if there are no SN8000B 8-Slot or 4-Slot Directors, 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 SN8000B 8-Slot or 4-Slot Director, DC SAN Backbone Director, or DC04 SAN Director to install a DC Dir Switch MP Extension Blade

### Table 32: Integrated Fibre Channel routing, 1606 Extension SAN Switch, and DC Director Switch MP Extension Blade scalability rules for firmware 7.x

| | | |
|---|---|---|
| **Edge fabric scalability** | | |
| Domains | Maximum number of real domains (local switches) per edge fabric | 26 |
| Devices (user ports) | Maximum number of local and remote devices per edge fabric. | 1,200[1] |
| | Maximum number of imported devices per edge fabric | 1,000 |
| **LSAN zone scalability** | | |
| Zoning | Maximum number of entries per LSAN zone<br>**NOTE:** LSAN zones support WWN-based zoning only. | 64 |
| **Meta SAN scalability** | | |

*Table Continued*



https://t.me/learningnets

| | | |
|---|---|---|
| Edge fabrics | Maximum number of edge fabrics connected in a Meta SAN | Edge fabrics containing up to 1,500 WWNs:<br>• 48 (DC and SN8000B Directors)<br>• 24 (all others)<br>Edge fabrics containing up to 2,000 WWNs (FOS v6.3 and higher):<br>• 32 (DC and SN8000B Directors)<br>• 12 (all others) |
| Meta SAN | Maximum number of total devices per Meta SAN | 10,000 (with only DC and SN8000B Directors as routers)<br>5,000 (with all other 8 Gb or 16 Gb switches as routers) |
| | Maximum number of edge fabrics per chassis:<br>MP Router Blade | 12 |
| | Maximum number of edge fabrics per chassis:<br>DC Director, 8/80, 8/40, 1606, SN6000B, SN6500B switches with Integrated Routing<br>SN8000B 8-Slot SAN Backbone Director Switch, SN8000B 4-Slot SAN Director Switch | 24 (DC Director, SN8000B 8-Slot SAN Backbone Director Switch, SN8000B 4-Slot SAN Director Switch with Integrated Routing)<br>12 (all other chassis) |
| | Maximum number of EX_Ports per chassis with Integrated Routing<br>(DC and SN8000B Directors, 8/80, 8/40, 1606, SN6000B, SN6500B switches) | 128 for DC and SN8000B Directors.<br>Up to maximum port count for all other chassis |
| Routers | Maximum number of Routers per Meta SAN | 10 |

*Table Continued*



https://t.me/learningnets

| | | |
|---|---|---|
| LSAN | Maximum number of LSAN device entries (proxy devices) per Meta SAN | 10,000 (with only DC and SN8000B as routers)<br>5,000 (with any 8/40, 8/80, 1606, SN6000B as routers) |
| | Maximum number of LSAN zones per Meta SAN | 3,000[2] (6.**x** and 7.**x**) |
| | **Backbone fabric scalability** | |
| Switches | Maximum number of local switches per backbone fabric | 12 (6.**x**)<br>12 (7.**x**) |
| WWNs | Maximum number of local WWNs per backbone fabric | 512 |
| | **Hop count scalability** | |
| Hop count | Maximum number of hops between switches (including routers) in a Meta SAN | 12 |

[1] Tested to 1,200, but can support up to 1,500 or 2,000 (with reduced number of maximum edge fabrics) on v6.3 and higher. Contact your Hewlett Packard Enterprise representative for further details.

[2] Tested to 3,000, but will support up to 5,000 on FOS v6.3 and higher (all backbone FCRs with FOS v6.0.0 and higher). Contact your Hewlett Packard Enterprise representative for further details.

## Zoning limits and enforcement

Table Zoning enforcement for B-series Fibre Channel switches and MP Router LSANs describes zoning enforcement for B-series Fibre Channel switches and Routed LSANs, see also **Zoning** on page 411.



https://t.me/learningnets

**Table 33: Zoning enforcement for B-series Fibre Channel switches and Routed LSANs**

| Switches or routers | Configuration | Enforcement | Comments |
|---|---|---|---|
| SN8600B 8-Slot SAN Director<br>SN8600B 4-Slot SAN Director<br>SN6600B 32Gb FC Switch<br>SN6650B 32Gb FC Switch<br>SN3600B 32Gb FC Switch<br>Brocade 16Gb SAN Switch Module for HPE Synergy<br>Encryption SAN Switch<br>SN8000B 8-Slot SAN Backbone Director<br>SN8000B 4-Slot SAN Director<br>SN6000B 16Gb FC Switch<br>StoreFabric SN6500B 16Gb FC Switch<br>StoreFabric SN4000B SAN Extension Switch<br>SN3000B 16Gb FC Switch<br>DC SAN Backbone Director<br>DC04 SAN Director<br>SAN Switch 8/8, 8/24, 8/40, 8/80<br>Brocade 16Gb SAN Switch for BladeSystem c-Class<br>Brocade 8Gb SAN Switch for BladeSystem c-Class<br>Brocade 4Gb SAN Switch for c-Class BladeSystem<br>1606 Extension SAN Switch<br>DC Dir Switch MP Extension Blade<br>B-series FCoE CN switch | Define zones using domain number and port number | All frames are checked by the hardware | Frame-based hardware enforcement |
| | Define zones using WWNs only | | |
| | Define zones using combination of domain/port numbers and WWNs or if a port is in multiple zones, and is defined by WWN in one zone and by domain/port in another | The only frames checked by hardware are the ELS frames (such as PLOGI and RNID) used to establish a session. | Session-based hardware enforcement |

## Zoning guidelines for B-series Fibre Channel switches

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

- Always zone using the highest Fabric OS-level switch.
- Switches with earlier Fabric OS versions do not have the capability to view all the functionality that a newer Fabric OS provides, as functionality is backwards compatible but not forwards compatible.
- Zone using the core switch versus an edge switch.
- Zone using an enterprise-class platform rather than a switch.
- An enterprise-class platform has more resources to handle zoning changes and implementations.



https://t.me/learningnets

- Frame-based hardware enforcement is in effect, on a per-zone basis, if all members of a zone are identified the same way, either using WWNs or domain, index notation, with no overlapping zones.
- Session-based hardware enforcement is in effect in the following cases, on a per-zone basis:
  - A zone does not have either all WWN or all domain, index entries.
  - Overlapping zones (in which zone members appear in two or more zones).
- 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 `portzoneshow` command to display and verify the zoning status of each port. The `portzoneshow` command displays the status of each port:
  - Frame — based hardware enforcement
  - Session — based hardware enforcement

**Zoning guidelines**

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

`Zone-1026 - port <port number> Out of CAM entries.`

Indicates that the total number of entries of S_ID CAM is above the limit while creating or adding a zone group. The maximum number of CAM entries allowed depends on the application-specific integrated circuit (ASIC). If hardware zoning enforcement is preferred, edit the zoning database to have zoned port IDs (PIDs) for that port.

`Zone-1004 - Base PID: 0x <Base PID>, Port Index: <Port Index>, Port: <Slot/Port>: enforcement changed to Session-based HARD Zoning.`

Indicates that the zoning enforcement has changed to session-based hard zoning due to one of the following conditions:

- The zone has a mix of WWN and domain, index (D,I) members.
- The Source Identifier (S_ID) list of the hardware-enforced zoning exceeded the S_ID limit.

**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 management interface, 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, Hewlett Packard Enterprise 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



https://t.me/learningnets

contains switches running Fabric OS 6.**x** and switches running Fabric OS 8.**x**, these functions must be executed from a switch running Fabric OS 8.**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 the fabric management interface 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 `tsclockserver` 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.



https://t.me/learningnets

# 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 and Cisco FCoE Converged Network switches** on page 98
- **C-series Fibre Channel switches** on page 137
- **Fibre Channel switch fabric rules** on page 145


**NOTE:** For information about using switches from the different series in the same SAN or fabric, see **SAN fabric connectivity and switch interoperability rules** on page 163.

## C-series Fibre Channel switches

C-series Fibre Channel switches have the following features:

- They 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.
- They are supported for use with FCIP through the IP Distance Gateway. For more information, see **SAN extension** on page 274.
- StoreFabric SN8500C 8-slot and 4-slot 16/32 Gb FC Director can accommodate high port density (up to 384 ports). The base units include 8-slot or 4-slot modular chassis with two hot swappable redundant Supervisor–1 Modules, six hot swappable redundant 3000W Power Supplies, and three hot swappable redundant Fabric 1 Modules, providing up to 384 ports of up to 24 Tb/s of FC system bandwidth and 384 4/8/16/32-Gb/s or 10-Gb/s full-line rate FC ports in a single chassis or up to 1152 FC ports in a single rack.
    - The open expansion slots of the SN8500C Director can be filled by StoreFabric C-series Family Modules, which include 48-port 32 Gb and 16 Gb FC Modules and 48-port FCoE Module.
    - The StoreFabric SN8500C 48-port 32 Gb and 16 Gb Modules are recommended for high-performance 32 and 16 Gb/s enterprise-level host connections, storage connections, and ISL connections.
    - The StoreFabric C-series SN8500C 48-port FCoE Module offers 384 FCoE Channel ports (autosensing 10Gb/s) in a single chassis.
    - The StoreFabric SN8500C Fabric-1 Module (up to six total per chassis) allows N+1 Fabric Modules protection to provide zero impact to Application Bandwidth in the event of a fabric card or supervisor card failure.
    - HPE StoreFabric SN6610C 32 Gb FC Switch:
        – Performance: Line-rate, non-blocking, non-oversubscribed architecture at 8/16/32 G
        – Affordable: Starts with 8-port base license. Expandable up to 32 ports
        – SAN Telemetry: Integrated line-rate visibility on all ports for storage traffic analytics
        – Optimized for All Flash Arrays: FC-NVMe qualified
        – Deployment Flexibility: Optional Fan and Power redundancy



https://t.me/learningnets

- HPE SN8000C 6-Slot Supervisor 2A Director Switch, SN8000C 9-Slot Supervisor 2A Director Switch, HPE SN8000C 13-Slot Supervisor 2A Fabric 2 Director Switch, SN8000C 13-Slot Supervisor 2A Fabric 3 Director Switch 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 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 18/4 Multiservice Module provides 18 4 Gb/s Fibre Channel ports and 4 GbE IP ports.

  - 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.
  - The SN8000C 8Gb 32-Port Advanced Fibre Channel Module uses eight 4-port port groups and delivers full line rate performance across all ports when used in conjunction with the SN8000C Fabric 3 Director Module, and is recommended for high-end 8 Gb/s storage system connections and ISL connections. When used in the SN8000C 13-Slot Supervisor 2A Fabric 3 Director Switch, or the SN8000C 13-Slot Supervisor 2A Fabric 2 Director Switch and the HPE Cisco MDS 9513 Director Switch upgraded with SN8000C Fabric 3 Director Modules, up to 24 ports can be configured to 10 Gb/s for use as ISLs. These 10 Gb/s interfaces enable reduced cabling for ISLs as they provide 50 percent higher data rate than 8 Gb/s interfaces.
  - The SN8000C 8Gb 48-Port Advanced Fibre Channel Module uses eight 6-port port groups and delivers 48-ports of line-rate 8 Gb/s with Arbitrated Local Switching. It is recommended for deploying dense virtual machine clusters. When used in the SN8000C 13-Slot Supervisor 2A Fabric 3 Director Switch, or the SN8000C 13-Slot Supervisor 2A Fabric 2 Director Switch and the Cisco MDS 9513 Director Switch upgraded with SN8000C Fabric 3 Director Modules, up to 24 ports can be configured to 10 Gb/s for use as ISLs. These 10 Gb/s interfaces enable reduced cabling for ISLs as they provide 50 percent higher data rate than 8 Gb/s interfaces.

- SN6500C switch:

  - The SN6500C 16Gb Multiservice Switch integrates both Fibre Channel and IP Storage services in a single system to allow maximum flexibility in user configurations. With 20 16 Gb/s Fibre Channel



https://t.me/learningnets

ports active by default, two 10 Gigabit Ethernet IP Storage Services ports, and eight 10 Gigabit Ethernet FCoE ports, the SN6500C 16Gb Multiservice switch is a comprehensive package, ideally suited for enterprise storage networks that require high performance SAN extension or cost-effective IP Storage connectivity for applications, such as Business Continuity using Fibre Channel over IP or iSCSI host attachment to Fibre Channel storage devices.

  - Using the eight 10 Gigabit Ethernet FCoE ports, the SN6500C 16Gb Multiservice Switch attaches to directly connected FCoE and Fibre Channel storage devices and supports multitiered unified network fabric connectivity directly over FCoE.

- SN6010C switch:

  - The SN6010C 16Gb Fabric Switch is a high-performance, flexible, cost-effective platform providing high-density, linerate16-Gbps ports for storage networking deployments in small, medium-sized, and large enterprise environments. The SN6010C switch offers outstanding value by providing high-availability, security, and ease of use at a cost-effective price in a compact onerack-unit (1RU) form factor. With the flexibility to expand from 12 to 48 ports in 12-port increments, the SN6010C offers the densities required to scale from an entry-level departmental switch to top-of-rack switch to edge connectivity in enterprise SANs. The SN6010C delivers a non-blocking architecture, with all 48 16-Gbps ports operating at line rate concurrently.

- 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 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 8Gb 12-port fabric switch and the MDS 8Gb 24-port fabric switch for 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.
- 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 core switches are named SN8000C and SN8500C. The SN8000C and SN8500C switches (Multilayer Directors) are director or core switches.

Multiprotocol and Multiservice products are designated with the **i** suffix, such as the MDS 9222i or as IP storage services modules, such as the IPS-4, IPS-8, 14/2, and 18/4. For IPS products, the number



https://t.me/learningnets

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

Hewlett Packard Enterprise supports all C-series Fibre Channel switches in a fabric if you:

- Use the recommended firmware versions.

  When using NX-OS and FabricWare in the same fabric, Hewlett Packard Enterprise highly recommends upgrading all fabric switches in the SAN to the latest supported software version. All switches in the same major NX-OS/FabricWare family must use the same software version. When updating switch firmware, you can use two successive NX-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** on page 145.

For the latest information on supported C-series Fibre Channel switches and firmware versions, see the SPOCK website at **http://www.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access.

After logging in, click **Switches** under **Other Hardware** in the last navigation panel of the window to access the Fibre Channel Switch Streams. Click on the **C-Series FC Switch Connectivity Stream** to open the document.

### Table 34: C-series Fibre Channel switches for NX-OS 6.x, 7.x

| Switch | Maximum number of Fibre Channel ports |
|---|---|
| StoreFabric SN8500C 8-Slot 16Gb FC Director | 384 |
| SN8000C 13-Slot Supervisor 2A Fabric 3 Director Switch | 528 |
| SN6000C Fabric Switch | 48 |
| SN6500C 16Gb Multiservice Switch | 40 |
| SN6010C 16 Gb Fabric Switch | 48 |
| SN8000C 6-Slot Supervisor 2A Director Switch | 192 |
| SN8000C 9-Slot Supervisor 2A Director Switch | 336 |
| SN8000C 13-Slot Supervisor 2A Fabric 2 Director Switch | 528 |
| MDS 8Gb 12-Port Fabric Switch for BladeSystem c-Class | 8 internal, 4 external |
| MDS 8Gb 24-Port Fabric Switch for BladeSystem c-Class | 16 internal, 8 external |



https://t.me/learningnets

### Table 35: C-series Fibre Channel switches for NX-OS 8.x

| Switch | Maximum number of Fibre Channel ports |
|---|---|
| StoreFabric SN8500C 8-Slot 16/32 Gb FC Director | 384 |
| SN6610C 32 GB Fabric Switch | 32 |
| SN6010C 16 Gb Fabric Switch | 48 |
| SN6500C 16 Gb Multiservice Switch | 40 |

**NOTE:** MDS 8 Gb Fabric Switches for BladeSystem c-Class are not supported with NX-OS 7.x

### Table 36: C-series Fibre Channel legacy switches

| Switch | Maximum number of Fibre Channel ports |
|---|---|
| MDS 9124 Fabric Switch | 24 |
| MDS 9222i Multiservice Fabric | 66 |

## 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 95 **nn** and 92 **nn**) for SAN extension (FCIP and iSCSI)
- NPV — The following C-series Fibre Channel switches are NPV compliant: MDS 9124, 8Gb fabric switch for BladeSystem c-Class, and SN6000C. NPV requires SAN-OS 3.3(5b) (or later), see **SAN fabric connectivity and switch interoperability rules** on page 163.
- FlexAttach virtual pWWN — The Cisco MDS 9124, 8Gb fabric switch for 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** on page 190.



https://t.me/learningnets

- 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

### Table 37: 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 |
| SN6010C Fabric | Yes/Yes | Yes/Yes | No | Yes | No | FC |
| SN6500C | Yes/Yes | Yes/Yes | No | Yes | No | Fc/FCoE/ FCIP |
| SN8500C 8-slot 16/32 Gb FC Director Switch | Yes/Yes | Yes/Yes | Yes | Yes | Yes | FC |
| SN6610C 8-slot 16/32 Gb FC Director Switch | Yes/Yes | Yes/Yes | Yes | Yes | Yes | FC |
| MDS 9124 Fabric | Yes/Yes | Yes/Yes | No | Yes | No | FC |
| MDS 9222i Multiservice Fabric | Yes/Yes | Yes/Yes | No | Yes | Yes | FC<br>FCIP<br>iSCSI |

*Table Continued*



https://t.me/learningnets

| Model | Redundant/ hot-swappable power | Redundant/ hot-swappable cooling | Redundant control processor | Nondisruptive code activation | Port module support | Protocol support |
|---|---|---|---|---|---|---|
| SN8000C 13-Slot Supervisor 2A Fabric 3 Director Switch | Yes/Yes | Yes/Yes | Yes | Yes | Yes | FC<br>FCIP<br>iSCSI |
| MDS 8Gb fabric switch for BladeSystem c-Class | N/A | N/A | No | Yes | No | FC |

## Usage

**Table 38: 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<br>(48-port stackable solution) | Not recommended | |
| SN6010C | Excellent (48-port) stackable solution | | |
| SN6610C | Excellent (32-port) stackable solution | | |
| SN6500C | Excellent<br>(40 port maximum) | | |
| MDS 9124 Fabric | Excellent<br>(24 port maximum) | | |
| MDS 9222i Multiservice Fabric | Excellent | Not recommended | |
| SN8000C 6-Slot Supervisor 2A Director Switch | Good | Excellent | |
| SN8000C 9-Slot Supervisor 2A Director Switch | | | |

*Table Continued*



https://t.me/learningnets

| Model | 1–48 total ports | 49–224 total ports | 225–512 total ports |
|---|---|---|---|
| SN8000C 13-Slot Supervisor 2A Fabric 3 Director Switch<br>SN8000C 13-Slot Supervisor 2A Fabric 2 Director Switch | | | |
| SN8500 8-slot and 4-slot 16/32 Gb Director Switch | | | |
| MDS 8Gb Fabric switch for BladeSystem c-Class | Excellent<br>(8 port maximum) | Not recommended | |

**Table 39: 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 | | |
| SN6010C | | | |
| SN6500C | | | |
| SN6610C | | | |
| MDS 9124 Fabric | | | |
| MDS 9222i Multiservice Fabric | Excellent | | Very good |
| SN8000C 6-Slot Supervisor 2A Director Switch | Excellent (select modules to optimize performance or user port count) | | |
| SN8000C 9-Slot Supervisor 2A Director Switch | | | |
| SN8000C 13-Slot Supervisor 2A Fabric 3 Director Switch | | | |
| SN8000C 13-Slot Supervisor 2A Fabric 2 Director Switch | | | |

*Table Continued*



https://t.me/learningnets

| Model | 1–48 total ports | 49–224 total ports | 225–512 total ports |
|---|---|---|---|
| SN8500C 8-slot 16/32 Gb FC Director Switch | | | |
| MDS 8Gb Fabric switch for BladeSystem c-Class | Excellent | | |

# 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

### Table 40: C-series Fibre Channel switch operating system and storage system support

| Operating systems | Storage products |
|---|---|
| • HP-UX<br>• OpenVMS<br>• Citrix Xen<br>• IBM AIX<br>• Oracle Linux<br>• Red Hat Linux<br>• SUSE SLES Linux<br>• Microsoft Windows<br>• Solaris<br>• VMware ESX | • P6550/P6500/P6350/P6300 EVA<br>• MSA2050/2052/1050 SAN 2040/2042/1040 SAN<br>• P9500<br>• XP24000/20000/12000/10000<br>• XP7<br>• 3PAR StoreServ 10000/7000/8000/9000/20000<br>• P9000 (XP) Continuous Access<br>• P6000 Continuous Access See **HPE P6000 Continuous Access SAN integration** on page 247. |

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

See **Heterogeneous server rules** on page 179 and the following storage system rules chapters:

- **MSA storage system rules** on page 227
- **HPE StoreVirtual 3200 and 4000 Storage System Rules** on page 234
- **P6000 storage system rules** on page 238
- **P9000/XP storage system rules** on page 252
- **HPE 3PAR StoreServ storage rules** on page 260



https://t.me/learningnets

## 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 P9000 (XP) Continuous Access and P6000 Continuous Access configurations. However, additional rules apply to Continuous Access implementations. For more information, see the *HPE P6000 Continuous Access Implementation Guide*.

### Table 41: 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 | SN8000C 6-Slot Supervisor 2A Fabric 2 Director Switch supports up to 192 ports over four modular slots (four 48-port modules[1]). |
| 3 | SN8000C 9-Slot Supervisor 2A Fabric 2 Director Switch supports up to 336 ports over seven modular slots (seven 48-port modules[1]).<br>**NOTE:** This requires Supervisor-2 or Supervisor-2A Modules. |
| 4 | SN8000C 13-Slot Supervisor 2A Fabric 3 Director Switch, SN8000C 13-Slot Supervisor 2A Fabric 2 Director Switch supports up to 528 ports over eleven modular slots (eleven 48-port modules[1]). |
| 5 | SN8500C 8-Slot 16/32 Gb FC Director Switch, supports up to 384 ports over eight modular slots (eight 48-port modules[1]). |
| 6 | Maximum of seven switch hops (eight switches) between any two communicating devices |
| 7 | Maximum of 80 VSANs per fabric[2] |
| 8 | Maximum of 4,000 IVR shared devices per routed fabric<br>**NOTE:** IVR is not supported on the MDS 9124, MDS 9124e, and MDS 9134. |
| 9 | Maximum of 2,000 IVR zones per routed fabric |
| 10 | 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. |
| 11 | FCIP is supported with IPsec data encryption using the 18/4 Multiservice Module, 14/2 Multiprotocol Services Module, and MDS 9222i. IPsec is supported with P6000 Continuous Access and P9000 (XP) Continuous Access. |
| 12 | Maximum of 114 FLOGIs or F Disc per port group for the 9200, 9500, and SN8000C series switches. Maximum of 89 FLOGIs or F Disc per port for the 9100 series switches. |
| 13 | 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.<br>Do not configure any switch with a domain ID of 8, which is reserved for HP-UX. |



https://t.me/learningnets

[1] For information about the 48-port module, see **C-series Fibre Channel switches** on page 137.
[2] MDS 9124 supports a maximum of 16 VSANs.

**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 42: ISL maximums

| Switch | Total number of available user ports | Number of ports allowed as ISLs |
|---|---|---|
| SN8500C 48-port 4 32Gb FC Module | 48 | 48 at 32 Gb/s |
| SN8500C 48-port 16Gb FC Module | 48 | 48 at 16 Gb/s |
| SN8500C 48-port 10 Gb FCoE module | 48 | 48 at 10 Gb/s |
| SN6000C | 48 | 48 at 8 Gb/s |
| SN6500C | 50 | 40 at 16 Gb/s |
| SN6010C | 48 | 48 at 16 Gb/s |
| SN6610C | 32 | 32 at 32 Gb/s |
| SN8000C 8Gb 32-Port Advanced Fibre Channel Module | 32 | 32 at 8 Gb/s, 24 at 10Gb/s |
| SN8000C 8Gb 48-Port Advanced Fibre Channel Module | 48 | 48 at 8 Gb/s, 24 at 10Gb/s |
| MDS 9**xxx** 48-port 8 Gb/s Host-Optimized Fibre Channel Switching Module | 48 | 48 at 1 Gb/s |
| MDS 95**xx** 48-port 8 Gb/s Performance Fibre Channel Switching Module | 48 | 48 at 2 Gb/s |
| MDS 95**xx** 24-port 8 Gb/s Fibre Channel Switching Module | 24 | 24 at 4 Gb/s |
| MDS 9**xxx** 48-port 4 Gb/s Fibre Channel Switching Module | 48 | 48 at 1 Gb/s |
| MDS 9**xxx** 24-port 4 Gb/s Fibre Channel Switching Module | 24 | 24 at 2 Gb/s |

*Table Continued*



https://t.me/learningnets

| Switch | Total number of available user ports | Number of ports allowed as ISLs |
|---|---|---|
| MDS 9**xxx**12-port 4 Gb/s Fibre Channel Switching Module | 12 | 12 at 4 Gb/s |
| MDS 9**xxx** 4-port 10 Gb/s Fibre Channel Switching Module | 4 | 4 at 10 Gb/s |
| MDS 9**xxx** 18/4 Multiservice Module | 18 | 18 at 2 Gb/s |
| MDS 9124 Fabric | 24 | 24 at 4 Gb/s |

## Zoning

**Basic Zoning and Enhanced Zoning:**

The zoning feature complies with the FC-GS-4 and FC-SW-3 standards. Both standards support the basic zoning functionality and enhanced zoning functionality.

In **basic zoning**, administrators can make simultaneous configuration changes. Upon activation, one administrator can overwrite another administrator's changes. If a zone is part of multiple zonesets, you create an instance of this zone in each zoneset. In the basic zoning the default zone policy is defined per switch. To ensure smooth fabric operation, all switches in the fabric must have the same default zone setting. To distribute the zoning database, you must reactivate the same zone set.

**Enhanced zoning** enables to perform all configurations using a single configuration session which ensures consistency within the fabric. It enforces and exchanges the default zone setting throughout the fabric. Enhanced Zoning implements changes to the zoning database and distributes it without reactivation. By default, the enhanced zoning feature is disabled in all switches in the Cisco MDS 9000 Family. It can be enabled on specified VSANs. When you begin a session, the switch acquires a fabric wide lock, distributes the active and full zoning database using the enhanced zoning data structures, distributes zoning policies and then releases the lock. All switches in the fabric then move to the enhanced zoning mode.

For additional information and commands on the Basic and Enhanced Zoning, see the *Cisco MDS 9000 Family NX-OS Fabric Configuration Guide*.

**Smart Zoning:**

Smart Zoning is the preferred zoning method for C-Series switches.

Smart Zoning is supported for NX-OS Software Release 5.2(6) or later. Smart zoning implements hard zoning of large zones with fewer hardware resources than was previously required. The traditional zoning method allows each device in a zone to communicate with every other device in the zone. The administrator is required to manage the individual zones according to the zone configuration guidelines. Smart zoning eliminates the need to create a single initiator to single target zones. For more information on configuring Smart Zoning, see the *CISCO MDS 9000 Family NX-OS Fabric Configuration Guide* for the appropriate version of NX-OS running in your SAN.



https://t.me/learningnets

## Zoning limits and enforcement

### Table 43: Zoning limits for C-series Fibre Channel switches

| Rule number | Description[1] |
|---|---|
| 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. |

[1] IVR is not supported on the MDS 9124.

### Table 44: Zoning enforcement for C-series Fibre Channel switches

| Switch | Configuration | Enforcement | Comments |
|---|---|---|---|
| MDS 8Gb Fabric switch for BladeSystem c-Class<br>SN6000C<br>SN8000C 6-Slot Supervisor 2A Director Switch<br>SN8000C 9-Slot Supervisor 2A Director Switch<br>SN8000C 13-Slot Supervisor 2A Fabric 2 Director Switch<br>SN8000C 13-Slot Supervisor 2A Fabric 3 Director Switch<br>SN8500C 8-Slot 16/32 Gb FC Director Switch<br>SN6610C<br>SN6500C<br>SN6010C<br>MDS 9222i<br>MDS 9124 | Defines zones using Domain ID and port number<br>Defines zones using WWNs only<br>Defines zones using Domain ID and port number with WWNs | Access authorization at frame level in hardware | Hard zoning |

## C-series VSAN high availability

**Figure 47: C-series Fibre Channel switch high-availability VSAN management configuration** on page 150 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** on page 40.



https://t.me/learningnets

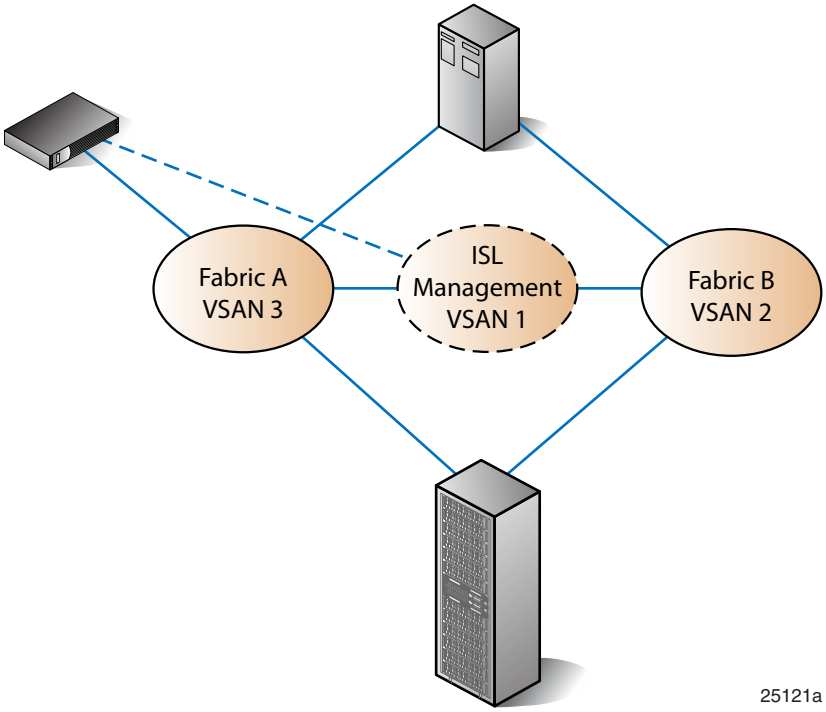



**Figure 47: C-series Fibre Channel switch high-availability VSAN management configuration**



https://t.me/learningnets

# H-series switches and fabric rules

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


- **H-series switches** on page 151
- **Fabric rules** on page 156


## 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



https://t.me/learningnets

- (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 SN6**xxx**, 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 SAN Connection Manager (SCM).

## Switch models

Hewlett Packard Enterprise 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 H-series switches table). When updating switch firmware, you can use two successive switch firmware versions temporarily.
- The fabric rules are followed, see H-series switches table.

For the latest information on supported H-series Fibre Channel switch and firmware versions, see the SPOCK website at **http://www.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access.

After logging in, click **Switches** under **Other Hardware** in the last navigation panel of the window to access the Fibre Channel Switch Streams. Click on the **H-Series FC Switch Connectivity Stream** to open the document.

To download the latest firmware from the Hewlett Packard Enterprise Support Center website:

1. Go to the Hewlett Packard Enterprise Support Center home page at **http://www.hpe.com/info/hpesc**.
2. To locate the firmware for a particular product, go to the Hewlett Packard Enterprise Support Center website: **http://www.hpe.com/support/hpesc**.


https://t.me/learningnets

3. Under **Select your HPE product**, enter the product name or number.
4. Click **Go**, and follow the next instructions.

As the product information is updated periodically, Hewlett Packard Enterprise recommends that you download release notes and additional documentation from the Hewlett Packard Enterprise Support Center website.

To download release notes and additional documentation:

1. Go to the Hewlett Packard Enterprise Support Center home page at **http://www.hpe.com/support/hpesc**.
2. From the left navigation panel, under **KNOWLEDGE BASE** click **Manuals**. The **Select a Product** page appears.
3. Under **All HP products**, click **Storage**.
4. Under **Storage**, click **Storage Networking**.
5. Under **Storage Networking**, click **H-series Switches**.
6. Under **H-series Switches**, click **H-series SAN Switches**.
7. Under **H-series SAN Switches**, click appropriate switch product.
8. Select the category of the documentation desired such as User guide, Setup and Install, Getting Started, or General Reference.

   **TIP:** Release notes are available in General Reference.

9. To download the document, click the title of the desired document from the table.

### Table 45: H-series switches

| HPE switch | Number of ports |
|---|---|
| SN6000 Fibre Channel Switch | 8/12/16/20/24 |
| 8/20q Fibre Channel Switch | 8/12/16/20 |

## 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 *HPE SN6000 Fibre Channel Switch Installation and Reference Guide*, available at the Hewlett Packard Enterprise Support Center website.
- 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)



https://t.me/learningnets

- 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 HPE, a VAR, or the customer
- Provides end-to-end management, including switch, HBA, and storage provisioning

Features of the SCM 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


https://t.me/learningnets

- 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 Fibre Channel distance rules for 32 Gb/s, 16 Gb/s and 8 Gb/s switch models (B-series, C-series, and H-series table and **H-series switch settings** on page 284)
- 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 Hewlett Packard Enterprise website **http://www.hpe.com/info/EFMS**. For information about configuring the switch using EFMS, see the *HPE 8/20q and SN6000 Fibre Channel Switch Enterprise Fabric Management Suite User Guide*.

### Table 46: H-series switches high-availability features

| Model | Redundant/hot-swappable power | Redundant/hot-swappable cooling | Nondisruptive code load/ activation | Nondisruptive port expansion/ upgrade |
|---|---|---|---|---|
| HPE SN6000 Stackable Dual Power Supply Fibre Channel Switch | Yes/Yes | Yes/Yes | Yes/Yes | Yes |
| SN6000 Stackable Single Power Supply Fibre Channel Switch | No/No | No/No | Yes/Yes | Yes |
| 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.



https://t.me/learningnets

# 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 HPE default switch settings for all configurations.
- All switch and fabric rules apply to fabrics implemented with the switch firmware versions listed in H-series switches table.

## Servers, operating systems, and storage products

### Table 47: H-series switches servers, operating systems, and storage system support

| Servers | Operating systems | Storage products |
|---|---|---|
| • HPE ProLiant ML-Series and DL-Series servers<br>• HPE Integrity servers<br>• BladeSystem c-Class ProLiant Server Blades<br>• BladeSystem c-Class Integrity Server Blades | • Microsoft Windows<br>• VMware ESX<br>• Red Hat Enterprise Linux (RHEL)<br>• SUSE Linux Enterprise Server (SLES)<br>• HP-UX<br>• Citrix Xen Enterprise<br>• Integrated Citrix XenServer<br>• Apple Mac OS X<br>• Oracle Linux<br>• Solaris | • MSA 2040/2042/1040 SAN<br>• MSA 2050/2052 SAN<br>• HPE StoreVirtual 3200 FC<br>• P6550/P6500/P6350/P6300 EVA<br>• P9500<br>• P2000 G3 FC MSA<br>• XP24000/20000, XP12000/10000<br>• 3PAR StoreServ 10000/7000/8000/20000<br>• 3PAR Remote Copy<br>• P6000 Continuous Access[1] |

[1] P6000 Continuous Access is supported with P6350/P6550 and firmware XCS version 11001000, P6300/P6500, EVA4400/6400/8400 and firmware XCS version10000000 and 10001000, with the data replication protocol option set for "HPE SCSI FC Compliant Data Replication Protocol," and with switch firmware 8.0.4.04.00 and 8.0.14.03.00.
For current storage system support and firmware, see the SPOCK website at **http://www.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access.
For more configuration information, see the *HPE P6000 Continuous Access Implementation Guide*.

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

See **Heterogeneous server rules** on page 179 and the following chapters:

- **Heterogeneous server rules** on page 179
- **MSA storage system rules** on page 227
- **HPE StoreVirtual 3200 and 4000 Storage System Rules** on page 234
- **P6000 storage system rules** on page 238



https://t.me/learningnets

- **P9000/XP storage system rules** on page 252
- **HPE 3PAR StoreServ storage rules** on page 260

To support SCM and other HPE software, the server running the SCM management software (the management station) must meet the following minimum requirements:

- Microsoft Windows Server 2008 x64/x86 with Service Pack 2
- Microsoft Windows Server 2008 R2 SP1 x64/x86

## Fabric rules for H-series switches

### Table 48: 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. |
| 4 | Within a fabric, if you assign a domain number (domain ID), it must be unique. Do not configure any switches with a domain ID of 8, which is reserved for HP-UX. |
| 5 | For HPE StoreVirtual 3200 FC or StoreVirtual 4000 FC, unique domain IDs are required in the entire Fibre Channel SAN to which the StoreVirtual storage system is connected, including multiple independent fabrics. Overlapping domain IDs (even in non-merged/ independent fabrics) are not supported. |

## 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.



https://t.me/learningnets

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. Hewlett Packard Enterprise 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:

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

### Table 49: Supported switches in an H-series switch with TR remote fabric

| HPE switch name | Firmware version |
|---|---|
| HPE StoreFabric SN6500B 16Gb 96 Fibre Channel Switch<br>HPE SN6000B 16Gb 48 port Fibre Channel Switch<br>HPE SN3000B 16Gb 24 Fibre Channel Switch<br>SN3000B 16Gb 24 Fibre Channel Switch<br>StorageWorks 8Gb DC SAN Backbone Director<br>StorageWorks 8Gb DC04 SAN Director<br>StorageWorks 8/40 SAN Switch<br>Brocade 8Gb SAN Switch for BladeSystem c-Class<br>HPE 8/8 Base, SAN Switch<br>HPE 8/24 SAN Switch<br>HPE 8/40 SAN Switch<br>HPE 8/80 SAN Switch | B-series 7.4.x or later with H-series fw 8.0.14.08.00 only |
| Brocade 8Gb SAN Switch for BladeSystem c-Class<br>8/8 Base, SAN Switch<br>8/24 SAN Switch<br>8/40 SAN Switch<br>8/80 SAN Switch | B-series 7.4.x or later |

*Table Continued*



https://t.me/learningnets

| HPE switch name | Firmware version |
|---|---|
| SN8000C 6-Slot Supervisor 2A Director Switch<br>SN8000C 9-Slot Supervisor 2A Director Switch<br>SN8000C 13-Slot Supervisor 2A Fabric 2 Director Switch<br>SN8000C 13-Slot Supervisor 2A Fabric 3 Director Switch<br>Cisco MDS 9222i Fabric Switch<br>StorageWorks SN6000C Fabric Switch (Cisco MDS 9148)<br>Cisco MDS 8Gb Fabric Switch for BladeSystem c-Class | C-series 6.2(17) or later with H-series fw 8.0.14.08.00 only |
| SN8000C 6-Slot Supervisor 2A Director Switch<br>SN8000C 9-Slot Supervisor 2A Director Switch<br>SN8000C 13-Slot Supervisor 2A Fabric 2 Director Switch<br>Cisco MDS 9222i Fabric Switch<br>Cisco MDS 9124 Fabric Switch | C-series 5.2(8h) or later |
| SN8000C 6-Slot Supervisor 2A Director Switch<br>SN8000C 9-Slot Supervisor 2A Director Switch<br>SN8000C 13-Slot Supervisor 2A Fabric 2 Director Switch<br>Cisco MDS 9222i Fabric Switch<br>Cisco MDS 9124 Fabric Switch | C-series 4.2(9) with H-series fw 8.0.14.03.00 only |

### Table 50: 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. |
| 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. |

*Table Continued*



https://t.me/learningnets

| Rule number | Description |
|---|---|
| 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:** SCM version 3.00 (or later) is required for the SN6000 Fibre Channel Switch. SCM cannot manage or discover remote switches or devices in the remote fabric when using the Transparent Router feature. The Physical Connection map displays the remote fabric as a disabled switch, indicating that SCM cannot manage the switch.

## Zoning limits and enforcement

### Table 51: 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. |



https://t.me/learningnets

### Table 52: Zoning enforcement for H-series switches

| Switch | Configuration | Enforcement | Comments |
|---|---|---|---|
| SN6000 Fibre Channel Switch | Define zones using domain number and port number | Access authorization at frame level in hardware | Hard zoning |
| 8/20q Fibre Channel Switch | Define zones using WWNs only | Access authorization at frame level in hardware | Hard zoning[1] |
| | Define zones using combination of domain/port numbers and WWNs | Name Server directory-based authentication | Soft zoning |

[1] 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.



https://t.me/learningnets

# McDATA/M-series switches

All M-series products have reached End Of Support Life (EOSL) and are no longer supported:

- McDATA 4 Gb SAN Switch for HPE p-Class BladeSystem
- HPE StorageWorks Director 2/140
- HPE StorageWorks Director 2/64
- HPE StorageWorks Edge Switch 2/32
- HPE StorageWorks Edge Switch 2/24
- HPE StorageWorks Edge Switch 2/12
- StorageWorks HA-Fabric Manager (HAFM)
- HA Fabric Manager Appliance
- HA Fabric Manager Appliance FRU
- HPE HA Fabric Manager SW UG kit

The following McDATA and Brocade products have reached End Of Support Life (EOSL) and are no longer supported:

- Sphereon 4700/M4700 Fabric Switch
- Sphereon 4400/M4400 Fabric Switch
- McDATA Intrepid 10000 Director (i10K) / Brocade Mi10K Director



https://t.me/learningnets

# 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** on page 163
- **SAN fabric switch interoperability rules** on page 174
- **Third-party switch support** on page 175
- **SAN performance considerations** on page 175


## 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 a Hewlett Packard Enterprise storage representative.

- **Fx_Port** — Switch port capable of operating as an F_Port or FL_Port.
- **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 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 **HPE FC Switches for the c-Class BladeSystem and HPE Synergy server environment** on page 181.

### Device port interfaces

The server HBA or storage system controller device node port interfaces are as follows:



https://t.me/learningnets

- **N_Port** — Connects a Fibre Channel device to a switch
- **NL_Port** — Connects a Fibre Channel arbitrated loop or switch

## Fiber optic cables

All 32 Gb/s, 16 Gb/s, 8 Gb/s, 4 Gb/s, and 2 Gb/s Fibre Channel components use industry-standard LC connectors for fiber optic cable connections; all 1 Gb/s Fibre Channel 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). Hewlett Packard Enterprise supports the use of duplex couplers, provided that you do not exceed the overall cable loss budget for that cable segment.

Hewlett Packard Enterprise recommends the use of HPE PremierFlex OM4 and OM3+ fiber optic cables for 50 micron cable installations. HPE 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 data sheet: **http://www.hpe.com/info/PremierFlex**
- Cables portfolio data sheet: **http://www.hpe.com/info/StorageOpticalTransceivers**

**NOTE:** OM3+ is a HPE 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.

Hewlett Packard Enterprise also supports industry standard fiber optic cable types OM4, OM3, OM2, and OM1, see **Fiber optic cable loss budgets** on page 165.

For cable part numbers, see the Fibre Channel switch QuickSpecs at **http://www.hpe.com/info/QuickSpecs**.

Rules for fiber optic cable connections table describes the rules for fiber optic cable connections.

### Table 53: Rules for fiber optic cable connections

| Rule number | Description |
|---|---|
| 1 | Minimum bend radius — HPE 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. Hewlett Packard Enterprise 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. |
| 2 | Hewlett Packard Enterprise does not support a mixture of 50, 62.5, and 9 micron cable in the same cable segment.<br>**NOTE:** A cable segment is defined as one or more cables connected serially between two transceivers. |

*Table Continued*



https://t.me/learningnets

| Rule number | Description |
|---|---|
| 3 | Hewlett Packard Enterprise 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). Hewlett Packard Enterprise 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:<br>• 0.5 m for 50 and 62.5 micron cable<br>• 2.0 m for 9 micron cable<br>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 µm (Belcore GR-409 compliant). The connectors must be SC or LC duplex low metal (Belcore and IEC compliant).<br>9 micron cable must be duplex, tight-buffered, single-mode 9/125 µm (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 HPE 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 PremierFlex OM3+)

  See **Fiber optic cables** on page 164 for more information on PremierFlex.

- 50 micron fiber optic cable, modal bandwidth 4700 MHz-km at 850 nm (type OM4, type PremierFlex OM4)

**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.



https://t.me/learningnets

Hewlett Packard Enterprise 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 lists the 32 Gb/s fiber optic cable loss budgets when using OM1, OM2, OM3, or OM3+ multi-mode fiber optic cable.

### Table 54: 32 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) | 20 m | 1.57 dB | 0.75 dB |
| 50/125 micron (OM3 2000 MHz-km at 850 nm, OM3+ 3000 MHz-km at 850 nm) | 70 m | 1.75 dB | |
| 50/125 micron (OM4 4700 MHz-km at 850 nm) | 100 m | 1.86 dB | |

Table 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 55: 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 | 1.63 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 | |
| 9/125 micron (singlemode) | 10 km | 6.4 dB[1] | |
| | 25km (B-series only) | 14.0 dB[2] | |



https://t.me/learningnets

[1] This number assumes use of low-loss fiber optic cables for 16 Gb/s speeds.
[2] Minimum link distance is 10 km.

Table 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 56: 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 | |
| 50/125 micron (OM4 4700 MHz-km at 850 nm) | 550 m | 3.5 dB | 0.5 dB |

Table 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 57: 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[1] | |
| | 25 km (B-series only) | 14.0 dB[2] | |

[1] This number assumes use of low-loss fiber optic cables for 8 Gb/s speeds.
[2] Minimum link distance is 10 km.

Table 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.



https://t.me/learningnets

### Table 58: 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 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 59: 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 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 60: 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 | |

*Table Continued*



https://t.me/learningnets

| Cable | Maximum distance per cable segment | Total channel insertion loss | Maximum loss per mated connector pair |
|---|---|---|---|
| 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

Unless otherwise specified, the distances specified apply to both switch-to-switch (ISL) connectivity and device-to-switch connectivity.

Table 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 61: Fibre Channel distance rules for 32 Gb/s, 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 | **OM2 fiber**<br>20 m at 32 Gb/s<br>35 m at 16 Gb/s<br>50 m at 8 Gb/s<br>150 m at 4 Gb/s<br>300 m at 2 Gb/s | **OM3, OM3+ fiber**<br>70 m at 32 Gb/s<br>100 m at 16 Gb/s<br>150 m at 8 Gb/s<br>380 m at 4 Gb/s<br>500 m at 2 Gb/s | **OM4 fiber**<br>100 m at 32 Gb/s<br>125 m at 16 Gb/s<br>190 m at 8 Gb/s<br>400 m at 4 G/s | |
| 62.5 micron[1] multi-mode fiber optic cable and short-wave SFP+ transceivers | OM1 fiber<br>21 m ISL at 8 Gb/s | | | Heterogeneous SAN servers, Fibre Channel switches, and storage systems |
| 9 micron single-mode fiber optic cable and long-wave SFPs | 10 km ISL at 32 Gb/s (B-series only)<br>25 km ISL at 16 Gb/s (B-series only)<br>10 km ISL at 8 Gb/s (H-series with additional buffer credits allocated)[2]<br>10 km ISL at 8 Gb/s (B-series and C-series only)<br>25 km ISL at 8 Gb/s (B-series only)<br>3.3 km ISL at 8 Gb/s (H-series only, base switch)<br>6.6 km ISL at 4 Gb/s (H-series only, base switch)<br>10 km ISL at 2 Gb/s (H-series only, base switch) | | | P6000 Continuous Access (B-series, C-series, and H-series only)<br>P9000 (XP) Continuous Access (B-series and C-series only) |
| Fibre Channel using WDM[3] | 40 km ISL at 8 Gb/s (B-series only)<br>100 km ISL at 4 Gb/s<br>250 km at 2 Gb/s | | | |



https://t.me/learningnets

[1] Information for 62.5 micron fiber optic cable is provided to facilitate use of installed cable. Hewlett Packard Enterprise recommends 50 micron fiber optic cable for new installations that require multi-mode fiber.

[2] 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.

[3] WDM distance is the maximum distance for the WDM link.

For information about the supported SFPs required for the distances listed in this table, see the HPE switch Streams documents on the SPOCK website at **http://www.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access.

Table 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 62: 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)<br>20 m ISL at 20 Gb/s (H-series only) | |
| 50 micron multi-mode fiber optic cable and short-wave SFP+ transceivers | **OM2 fiber**<br>82 m ISL at 10 Gb/s | **OM3, OM3+ fiber**<br>300 m ISL at 10 Gb/s |
| 62.5 micron[1] 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)<br>40 km ISL at 10 Gb/s (C-series only) | |

[1] Information for 62.5 micron fiber optic cable is provided to facilitate use of installed cable. Hewlett Packard Enterprise recommends 50 micron fiber optic cable for new installations that require multi-mode fiber.

Table describes the distance rules for 4 Gb/s Fibre Channel connections when using 4 Gb/s Fibre Channel switch models.

### Table 63: 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 | **OM2 fiber**<br>150 m at 4 Gb/s<br>300 m at 2 Gb/s<br>500 m at 1 Gb/s | **OM3, OM3+ fiber**<br>380 m at 4 Gb/s<br>500 m at 2 Gb/s | **OM4 fiber**<br>400 m at 4 G/s | Heterogeneous SAN servers, Fibre Channel switches, and storage systems<br>P6000 Continuous Access and P9000 (XP) Continuous Access<br>EBS |
| 62.5 micron multi-mode fiber optic cable and short-wave SFPs | **OM1 fiber**<br>70 m ISL at 4 Gb/s | | | |

*Table Continued*



https://t.me/learningnets

| Interface/transport | Supported distances | Supported storage products |
|---|---|---|
| 9 micron single-mode fiber optic cable and long-wave SFPs | 4 km ISL at 4 Gb/s (C-series only)<br>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)<br>35 km ISL at 2 Gb/s with 2 Gb/s SFP (B-series only) | |
| Fibre Channel using WDM | 100 km at 4 Gb/s<br>250 km at 2 Gb/s<br>500 km at 1 Gb/s | |

Table describes the distance rules for 2 Gb/s Fibre Channel connections when using 2 Gb/s Fibre Channel switch models.

### Table 64: Fibre Channel distance rules for 2 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 | **OM2 fiber**<br>300 m at 2 Gb/s<br>500 m at 1 Gb/s | **OM3, OM3+ fiber**<br>500 m at 2 Gb/s | Heterogeneous SAN servers, Fibre Channel switches, and storage systems<br>P6000 Continuous Access and P9000 (XP) Continuous Access<br>EBS |
| 62.5 micron multi-mode fiber optic cable and short-wave SFPs | **OM1 fiber**<br>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) | | |
| Total distance | 200 km (B-series, C-series) | | |
| Fibre Channel using WDM | 100 km at 2 Gb/s<br>240 km at 1 Gb/s | | |

Table describes the distance rules for 1 Gb/s Fibre Channel connections when using 1 Gb/s Fibre Channel switch models.



https://t.me/learningnets

### Table 65: Fibre Channel distance rules for 1 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 GBICs and GLMs | 500 m at 1 Gb/s | Heterogeneous SAN servers, Fibre Channel switches, and storage systems<br>P6000 Continuous Access and P9000 (XP) Continuous Access<br>EBS |
| 62.5 micron multi-mode fiber optic cable and short-wave GBICs | 200 m at 1 Gb/s | |
| 9 micron single-mode fiber optic cable and long-wave GBICs | 10 km ISL at 1 Gb/s<br>35 km ISL (C-series) | |
| 9 micron single-mode fiber optic cable and very long distance GBICs | 100 km ISL (B-series) at 1 Gb/s | |
| Total distance | 200 km (B-series, C-series) | |
| Fibre Channel using WDM | 240 km at 1 Gb/s | |

Table describes the distance rules for ATM extension Fibre Channel connections.

### Table 66: ATM extension Fibre Channel distance rules

| Interface/transport | Heterogeneous SAN host-to-disk storage systems | HPE P9000 (XP) Continuous Access | 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 | | | |

Table describes the distance rules for FCIP extension Fibre Channel connections.

### Table 67: FCIP extension IP network distance rules

| Interface/transport | HPE P6000 Continuous Access and HPE P9000 (XP) Continuous Access | EBS[1] |
|---|---|---|
| B-series 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 |
| C-series IPS-8, 14/2, 18/4 | | Not supported |

*Table Continued*



https://t.me/learningnets

| Interface/transport | HPE P6000 Continuous Access and HPE P9000 (XP) Continuous Access | EBS[1] |
|---|---|---|
| IP Distance Gateway | P6000: 100 ms IP network delay one-way or 200 ms round-trip<br>XP: 50 ms IP network delay one-way or 100 ms round-trip | |
| HPE MPX200 Multifunction Router FCIP | P6000: 100 ms IP network delay one-way or 200 ms round-trip<br>XP: Not supported | Not supported |
| **Legacy FCIP products** | | |
| IPS-4 | 100 ms IP network delay one-way or 200 ms round-trip | Not supported |
| B-series MP Router (FCIP) | | 100 ms IP network delay one-way or 200 ms round-trip |
| Third-party SAN extension devices | See **SAN extension**. | Not supported |

[1] EBS is not supported with FCIP IPsec data encryption.

Table describes the distance rules for iSCSI bridging Fibre Channel connections.

### Table 68: iSCSI bridging Fibre Channel distance rules

| Interface/transport | Heterogeneous SAN Fibre Channel distances |
|---|---|
| HPE B-series iSCSI Director Blade | See tables Fibre Channel distance rules for 16 Gb/s and 8 Gb/s switch models (B-series, C-series, and H-series) through Fibre Channel distance rules for 1 Gb/s switch models (B-series and C-series switches. |
| HPE C-series IPS-4, IPS-8, 14/2, 18/4 | |
| EVA iSCSI Connectivity Option | |
| HPE MPX200 Multifunction Router iSCSI | |

Table describes the distance rules for Fibre Channel routing connections.

### Table 69: Fibre Channel and IP network routing distance rules

| Interface/ transport | Heterogeneous SAN Fibre Channel distances | HPE P6000 Continuous Access and HPE P9000 (XP) Continuous Access | EBS |
|---|---|---|---|
| HPE B-series MP Router Blade | See tables Fibre Channel distance rules for 4 Gb/s switch models (B-series and C-series switches) through Fibre Channel distance rules for 1 Gb/s switch models (B-series and C-series switches). | See tables Fibre Channel distance rules for 4 Gb/s switch models (B-series and C-series switches) through Fibre Channel distance rules for 1 Gb/s switch models (B-series and C-series switches) and FCIP extension IP network distance rules table. | |
| HPE C-series IVR (inter-VSAN routing) | | | |



https://t.me/learningnets

# SAN fabric switch interoperability rules

Hewlett Packard Enterprise supports two heterogeneous switch SAN configurations:

- Dual interoperable, heterogeneous SAN fabrics
- Interoperable, heterogeneous switch fabrics

## 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.

Hewlett Packard Enterprise 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.

When creating a dual interoperable, heterogeneous SAN, consider the following:

- Hewlett Packard Enterprise recommends that you use the same fabric topology and configuration in both fabrics to maintain balanced SAN performance.
- HA configurations require support for common HBA, driver, multipathing software, and storage array firmware versions because servers and storage connect to both fabrics.

## Interoperable, heterogeneous switch fabrics

An interoperable, heterogeneous switch fabric can contain different series of switches.

### Table 70: Heterogeneous switches in the same fabric

| Heterogeneous switch combinations | Reference |
|---|---|
| C-series and B-series | Fabric Interoperability: Merging Fabrics Based on C-series and B-series Fibre Channel Switches Application Notes |

The document referenced in Heterogeneous switches in the same fabric table is available on the SAN Infrastructure website:

**http://www.hpe.com/info/SDGManuals**

Hewlett Packard Enterprise supports interoperable, heterogeneous switch fabrics using NPIV (N_Port ID virtualization) connectivity. This method is supported for connectivity between B-Series and C-Series or H-Series switches, or between B-Series or C-Series switches and HPE FlexFabric switches.

### B-Series, C-Series, and H-Series switches

For more information about support and implementation of heterogeneous switch fabrics using B-Series, C-Series, or H-Series switches in the same fabric, see the HPE switch Streams documents on the SPOCK website at **http://www.hpe.com/storage/spock**.



https://t.me/learningnets

You must sign up for an HP Passport to enable access.

For the B-Series stream document, see the **BladeSystem with Brocade Access Gateway mode** on page 184. For the C-Series stream document, see the "N-Port Virtualization (NPV) & FlexAttach" sections. For the H-Series stream document, see the "Transparent Router" section of the stream document.

### HPE FlexFabric switches

For more information about support and implementation of heterogeneous switch fabrics using HPE FlexFabric switches with other switch series products, see the **HPE FlexFabric Switch Series Storage Solutions Configuration Guide**. This document is available on the SPOCK website at **http://www.hpe.com/storage/spock**.

You must sign up for an HP Passport to enable access.

# Third-party switch support

Hewlett Packard Enterprise Services offers support for certain third-party switches if you purchase third-party support through Multivendor Environment Services. For more information, see the Hewlett Packard Enterprise Storage Services website:

**http://www.hpe.com/services**

Hewlett Packard Enterprise Services also offers support for multivendor SANs. Hewlett Packard Enterprise 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*



https://t.me/learningnets

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 32 Gb/s, 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 32 Gb/s, 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:

  1. 32 Gb/s ports negotiate to 16 Gb/s and 8 Gb/s
  2. 16 Gb/s ports negotiate to 8 Gb/s and 4 Gb/s
  3. 8 Gb/s ports negotiate to 4 Gb/s and 2 Gb/s
  4. 4 Gb/s ports negotiate 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, Hewlett Packard Enterprise 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.


https://t.me/learningnets

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.



https://t.me/learningnets

# Host and storage system rules

Host and storage system rules are presented in these chapters:


- **Heterogeneous server rules** on page 179
- **MSA storage system rules** on page 227
- **HPE StoreVirtual 3200 and 4000 Storage System Rules** on page 234
- **P6000 storage system rules** on page 238
- **P9000/XP storage system rules** on page 252
- **HPE 3PAR StoreServ storage rules** on page 260
- **HPE Data Availability, Protection and Retention** on page 272




https://t.me/learningnets

# Heterogeneous server rules

This chapter describes platform configuration rules for SANs with specific operating systems and heterogeneous server platforms:


- **SAN platform rules** on page 180
- **Heterogeneous storage system support** on page 180
- **HPE FC Switches for the c-Class BladeSystem and HPE Synergy server environment** on page 181
- **BladeSystem with Brocade Access Gateway mode** on page 184
- **BladeSystem with Cisco N_Port Virtualization mode** on page 187
- **NPV with FlexAttach** on page 190
- **HPE BladeSystem c3000 enclosure considerations** on page 191
- **HBA N_Port ID Virtualization** on page 192
- **NonStop servers (XP only)** on page 193
- **HP-UX SAN rules** on page 207
- **HPE OpenVMS SAN rules** on page 210
- **HPE Tru64 UNIX SAN rules** on page 212
- **Apple Mac OS X SAN rules** on page 212
- **IBM AIX SAN rules** on page 214
- **Linux SAN rules** on page 216
- **Microsoft Windows SAN rules** on page 218
- **Oracle Solaris SAN rules** on page 219
- **VMware ESX SAN rules** on page 222
- **Citrix Xen SAN rules** on page 223
- **Heterogeneous SAN storage system coexistence** on page 224
- **Server zoning rules** on page 226


The platform configuration rules in this chapter apply to SANs that comply with the following fabric guidelines:


- **B-series switches and fabric rules** on page 105
- **C-series switches and fabric rules** on page 137
- **H-series switches and fabric rules** on page 151


Before implementation, contact a Hewlett Packard Enterprise storage representative for support information for specific configurations, including the following elements:



https://t.me/learningnets

- Server model
- Storage system firmware
- SAN attachment
- HBAs and drivers
- Multipathing

# SAN platform rules

### Table 71: 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:<br>• Operating system<br>• Fabric<br>• Storage system<br>• Mixed storage system types |
| 2 | All HPE and multivendor hardware platforms and operating systems that are supported in a homogeneous SAN are also supported in a heterogeneous SAN.<br>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 EVA single-server maximum configurations table. |

For cabling options for platforms that support high-availability multipathing, see **Cabling** on page 250.

# Heterogeneous storage system support

Hewlett Packard Enterprise supports HPE storage products on shared hosts and HBAs in HPE fabric environments that also have third-party storage products. A third-party cooperative support agreement between Hewlett Packard Enterprise Services and the third party is required if HPE will provide a single support point of contact that includes the third-party storage. Hewlett Packard Enterprise provides technical support for its products and cooperates with the third party's technical support staff, as needed.

Hewlett Packard Enterprise provides best-practices recommendations for connecting devices in the SAN, see **Best practices** on page 407. These rules apply to configurations that include SAN storage products and heterogeneous third-party SAN storage products:


https://t.me/learningnets

- Use zones to isolate HPE storage ports from third-party storage ports.
- HPE storage zones are governed by Hewlett Packard Enterprise product-specific configuration guidelines. See the following HPE storage system chapters:
  - **MSA storage system rules** on page 227
  - **HPE StoreVirtual 3200 and 4000 Storage System Rules** on page 234
  - **P6000 storage system rules** on page 238
  - **P9000/XP storage system rules** on page 252
  - **HPE 3PAR StoreServ storage rules** on page 260
- 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** on page 175.

For storage system coexistence support, see **Heterogeneous SAN storage system coexistence** on page 224.

# HPE FC Switches for the c-Class BladeSystem and HPE Synergy server environment

**Table 72: Supported switches**

| Speed | Switches |
|---|---|
| 8 Gb | B-series<br>Brocade 8Gb SAN Switch for BladeSystem c-Class<br>C-series<br>Cisco MDS 8Gb Fabric Switch for BladeSystem c-Class |
| 16 Gb | B-series<br>Brocade 16Gb SAN Switch for BladeSystem c-Class<br>Brocade 16 Gb FC SAN Switch for HPE Synergy |

# HPE Virtual Connect for the c-Class BladeSystem server environment

Hewlett Packard Enterprise offers multiple Virtual Connect products for c-Class BladeSystem servers:

- **HPE Virtual Connect 20/40 F8** on page 182
- **HPE Virtual Connect FlexFabric 10 Gb/24-Port Module for c-Class BladeSystem** on page 182
- **HPE Virtual Connect Flex-10/10D Ethernet Module for c-Class BladeSystem** on page 183



https://t.me/learningnets

- **HPE Virtual Connect 8 Gb 20-Port Fibre Channel Module for c-Class BladeSystem** on page 183
- **HPE Virtual Connect 8 Gb 24-Port Fibre Channel Module for c-Class BladeSystem** on page 183

## HPE Virtual Connect 20/40 F8

HPE Virtual Connect (VC) FlexFabric-20/40 F8 Modules are the simplest, most flexible way to connect virtualized server blades to data or storage networks. VC FlexFabric-20/40 F8 modules converge traffic inside enclosures and directly connect to external LANs and SANs. Using Flex-10 and Flex-20 technology with Fibre Channel over Ethernet and accelerated iSCSI, these modules converge traffic over high-speed 10Gb/20Gb connections to servers with FlexFabric Adapters. Each redundant pair of Virtual Connect FlexFabric modules provide eight adjustable downlink connections (six Ethernet and two Fibre Channel, or six Ethernet and two iSCSI or eight Ethernet) to dual-port 10Gb/20Gb FlexFabric Adapters on each server. Up to twelve uplinks with eight Flexport and four QSFP+ interfaces, without splitter cables, are available for connection to upstream Ethernet and Fibre Channel switches. Including splitter cables up to 24 uplinks are available for connection to upstream Ethernet and Fibre Channel. VC FlexFabric-20/40 F8 modules avoid the confusion of traditional and other converged network solutions by eliminating the need for multiple Ethernet and Fibre Channel switches, extension modules, cables, and software licenses. Also, Virtual Connect wire-once connection management is built-in, enabling server adds, moves, and replacement to be done in an expedient manner.

## HPE Virtual Connect FlexFabric 10 Gb/24-Port Module for c-Class BladeSystem

HPE Virtual Connect FlexFabric 10 Gb/24-port Modules provide a simple, flexible way to connect virtualized server blades to data or storage networks, or directly to Hewlett Packard Enterprise storage systems (see *Virtual Connect Direct-attach Fibre Channel for 3PAR storage* section).

Virtual Connect FlexFabric modules converge traffic inside enclosures and connect directly to external LANs and SANs, eliminating up to 95% of network sprawl at the server edge. Using Flex-10 technology with FCoE and accelerated iSCSI, these modules converge traffic over 10 GbE connections to servers with FlexFabric adapters.

Each redundant pair of Virtual Connect FlexFabric modules provides eight adjustable downlink connections (six Ethernet and two Fibre Channel, six Ethernet and two iSCSI, or eight Ethernet connections) to dual-port 10 Gb FlexFabric adapters on servers. Up to eight uplinks are available for connection to upstream Ethernet and Fibre Channel switches.

Virtual Connect FlexFabric modules are more efficient than traditional and other converged network solutions because they do not require multiple Ethernet and Fibre Channel switches, extension modules, cables, and software licenses. Also, built-in Virtual Connect wire-once connection management enables you to add, move, or replace servers in minutes.

For more information, see the product QuickSpecs at:

**http://www.hpe.com/info/Quickspecs-VCFF10-24PortMod**

**Virtual Connect Direct-attach Fibre Channel for 3PAR storage**

The Virtual Connect FlexFabric 10 Gb/24-port Module supports direct-connection of 3PAR StoreServ Storage Fibre Channel ports. This provides the option to deploy configurations with c-Class BladeSystems and 3PAR storage without an intermediate Fibre Channel switch or fabric. The result is a significant reduction in infrastructure costs and storage provisioning time, and an increase in performance due to reduced latency.

**NOTE:** Virtual Connect Direct-attach Fibre Channel for 3PAR storage has minimum Virtual Connect firmware and 3PAR firmware requirements. For more information, see the Virtual Connect and 3PAR storage documentation at **http://www.hpe.com/info/3PAR**.



https://t.me/learningnets

## HPE Virtual Connect Flex-10/10D Ethernet Module for c-Class BladeSystem

The Virtual Connect Flex-10/10D Module simplifies server connections by separating the server enclosure from the LAN, simplifies networks by reducing cables without adding switches to manage, allows changes to servers in minutes, and tailors network connections and speeds based on application needs. HPE Flex-10 technology significantly reduces infrastructure costs by increasing the number of NICs per connection without adding extra blade I/O modules, reducing cabling uplinks to the data center network.

- Dual-hop FCoE support allows FCoE traffic to be propagated out of the enclosure to an external FCoE-capable bridge.
- Simplifies networks by reducing cables without adding switches to manage.
- Allows you to wire once, then add, move and change network connections to thousands of servers in minutes instead of days or weeks from one console without affecting LAN and SAN.
- Each module has 30 ports for a total effective full-duplex bandwidth of 600 Gb and 10 dedicated SFP+ uplink ports, which can be 1GbE or 10GbE.
- Reduces network overhead costs by wiring once and making changes dynamically without additional network administrative support. Unlike other network virtualization offerings, Virtual Connect does not require manual changes to network connections each time a server is added or moved.
- Eliminates network sprawl at the server edge and saves up to 47% on upstream ToR switch cable connections. For more information, see the product QuickSpecs available at **http://www.hpe.com/info/Quickspecs-VCFF10-10DMod**.

## HPE Virtual Connect 8 Gb 20-Port Fibre Channel Module for c-Class BladeSystem

The HPE Virtual Connect 8 Gb 20-port FC Module offers enhanced Virtual Connect capabilities, allowing up to 128 virtual machines running on the same physical server to access separate storage resources. With this module:

- Provisioned storage resource is associated directly to a specific virtual machine, even if the virtual server is re-allocated within the BladeSystem.
- Storage management of virtual machines is no longer limited by the single physical HBA on a server blade; SAN administrators can now manage virtual HBAs with the same methods and viewpoint of physical HBAs.

The Virtual Connect 8 Gb 20-port Fibre Channel Module:

- Simplifies server connections by separating the server enclosure from SAN
- Simplifies SAN fabrics by reducing cables without adding switches to the domain
- Allows you to change servers in minutes

For more information, see the product QuickSpecs at:

**http://www.hpe.com/info/Quickspecs-VCFF8-20-PortFCMod**

## HPE Virtual Connect 8 Gb 24-Port Fibre Channel Module for c-Class BladeSystem

The Virtual Connect 8 Gb 24-port FC Module offers enhanced Virtual Connect capabilities, allowing up to 24 ports of connectivity to a Fibre Channel SAN.



https://t.me/learningnets

For more information, see the product QuickSpecs at:

**http://www.hpe.com/info/Quickspecs-VCFF8-24-PortFCMod**

## HPE Virtual Connect FC connectivity guidelines

Deploy 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 73: HPE 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 and H-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 | HPE supports a maximum of 4 VC-FC modules per blade enclosure. |

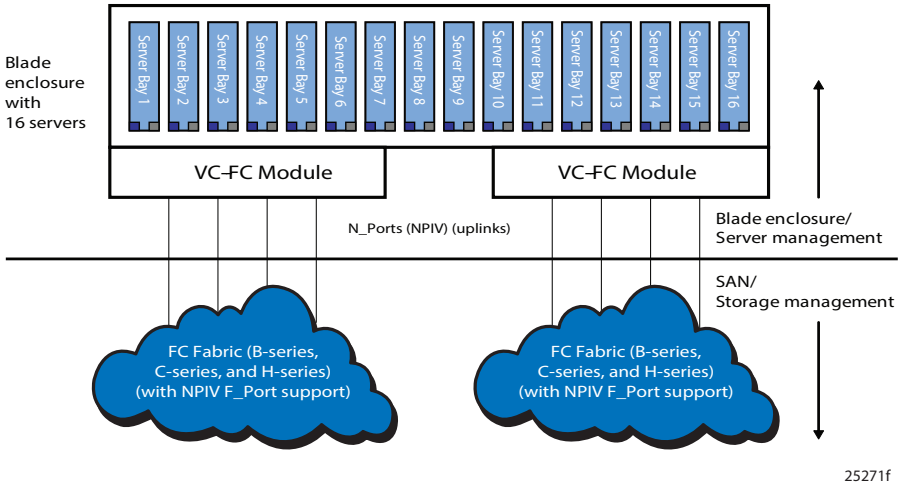



**Figure 48: HPE Virtual Connect Fibre Channel configuration**

# BladeSystem with Brocade Access Gateway mode

AG mode is a software-enabled feature available with the Brocade 8Gb SAN Switch for BladeSystem c-Class with NX-OS 5.2.8f 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.

The Brocade 8Gb SAN Switch for BladeSystem c-Class in AG mode supports a maximum of 24 ports:



https://t.me/learningnets

- 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 49: Brocade 8Gb SAN Switch for HPE c-Class BladeSystem in Access Gateway mode** on page 185 shows a view of an c-Class BladeSystem in AG mode.

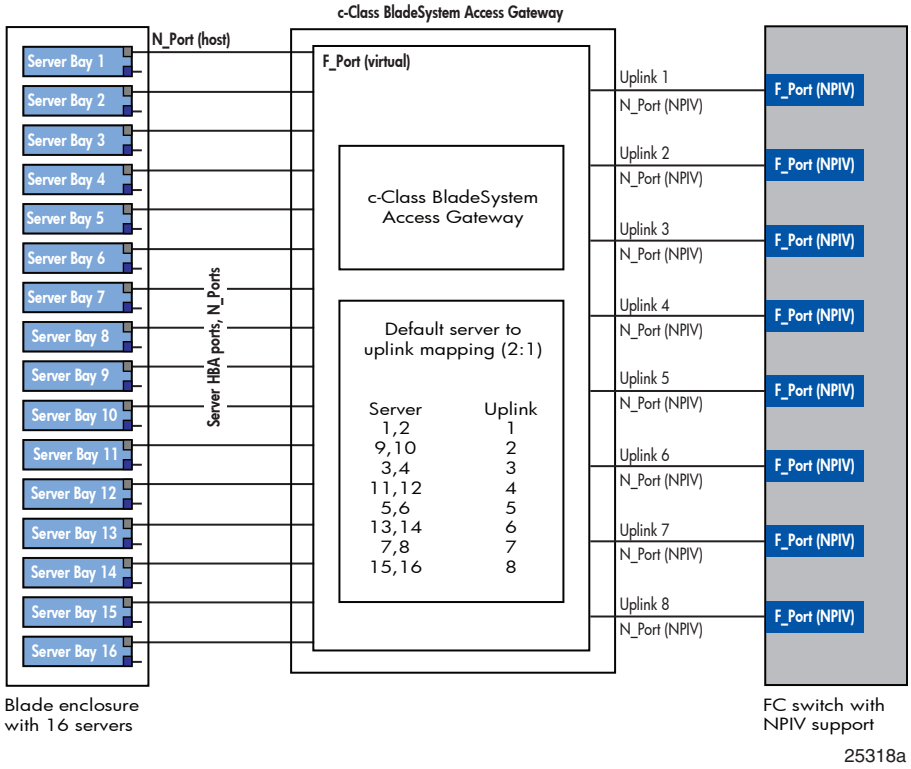



**Figure 49: Brocade 8Gb SAN Switch for HPE c-Class BladeSystem in Access Gateway mode**

**NOTE:** The uplink ports (N_Ports) in **Figure 49: Brocade 8Gb SAN Switch for HPE c-Class BladeSystem in Access Gateway mode** on page 185 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.



https://t.me/learningnets

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-series, C-series, and H-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 50: Access Gateway with dual redundant fabrics** on page 187 shows an AG with dual redundant fabrics.



https://t.me/learningnets

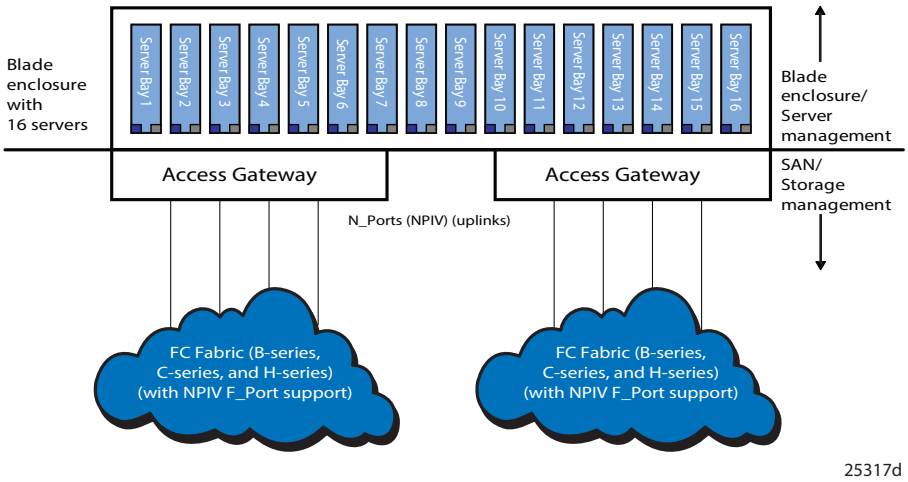



**Figure 50: Access Gateway with dual redundant fabrics**

**NOTE:**

The N_Ports in **Figure 50: Access Gateway with dual redundant fabrics** on page 187 are not host N_Ports and cannot be connected directly to storage.

Configuration highlights for **Figure 50: Access Gateway with dual redundant fabrics** on page 187 include:

- Redundant SANs, with each server connecting to one fabric through one AG module
- Ability to connect to B-series, C-series, and H-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 8Gb Fabric Switch for BladeSystem c-Class with NX-OS 5.2(8f) (or later). NPV mode does not require the purchase of additional hardware or software.

NPV is available only if the Cisco MDS 8Gb Fabric Switch for BladeSystem c-Class 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 8Gb Fabric Switch for BladeSystem c-Class is transparent to the hosts and fabric — they no longer function as standard switches.

**NOTE:** This section describes c-Class BladeSystems. NPV mode is also supported on the Cisco MDS 9124 Fabric Switches. For more information, see the *Cisco MDS 9000 Configuration Guide*.

The Cisco MDS 8Gb Fabric Switch for BladeSystem c-Class in NPV mode supports a maximum of 24 ports:



https://t.me/learningnets

- 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).

**Cisco MDS 8Gb Fabric Switch for BladeSystem c-Class** shows an c-Class BladeSystem in NPV mode.

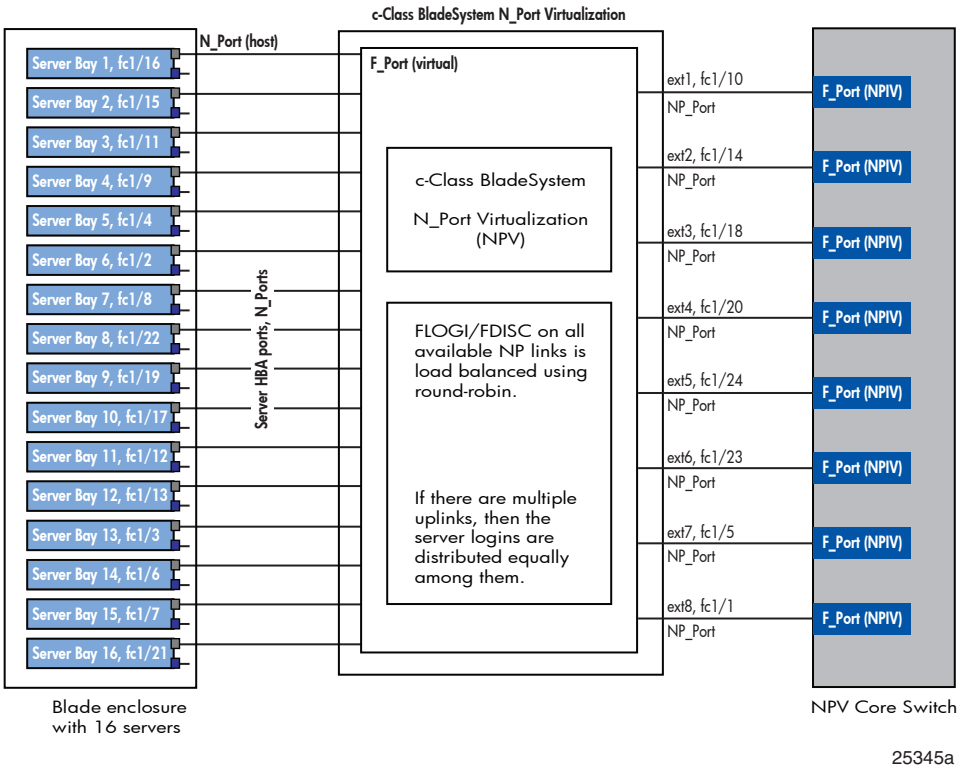



**Figure 51: Cisco MDS 8Gb Fabric Switch for BladeSystem c-Class in NPV mode**

**NOTE:** The NP_Ports in **Cisco MDS 8Gb Fabric Switch for BladeSystem c-Class in NPV mode** 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:


https://t.me/learningnets

- 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-series and C-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 52: NPV device with dual redundant fabrics** on page 189 shows an NPV device with dual redundant fabrics.

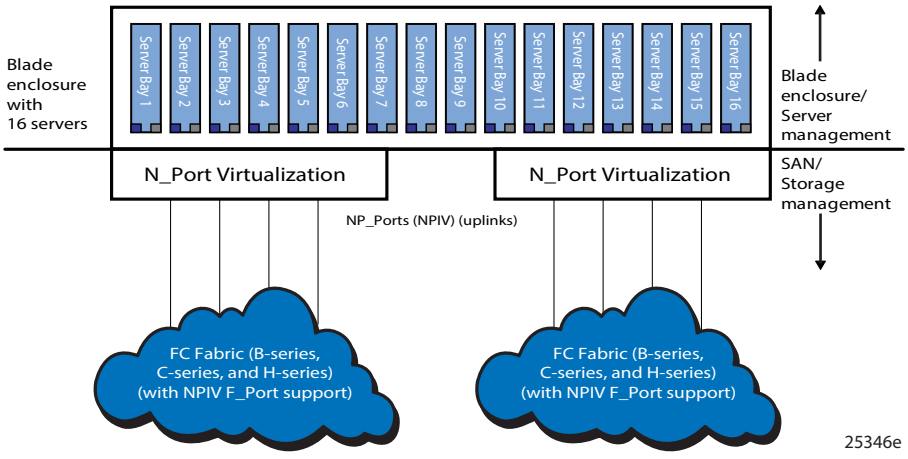



**Figure 52: NPV device with dual redundant fabrics**

**NOTE:**

The NP_Ports in **Figure 52: NPV device with dual redundant fabrics** on page 189 are not host N_Ports and cannot connect directly to storage.

The configuration shown in **Figure 52: NPV device with dual redundant fabrics** on page 189 includes:



https://t.me/learningnets

- Redundant SANs, with each server connecting to one fabric through one NPV device
- Connectivity to C-series and B-series fabrics
- Support for up to six NPV devices per HPE BladeSystem c7000 enclosure, or three NPV devices per HPE BladeSystem c3000 enclosure

# NPV with FlexAttach

The Cisco MDS 8Gb Fabric Switch for BladeSystem c-Class and MDS 9124 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. After 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 53: Cisco MDS 8Gb Fabric Switch for BladeSystem c-Class** on page 191 shows a view of an c-Class BladeSystem using NPV with FlexAttach.



https://t.me/learningnets

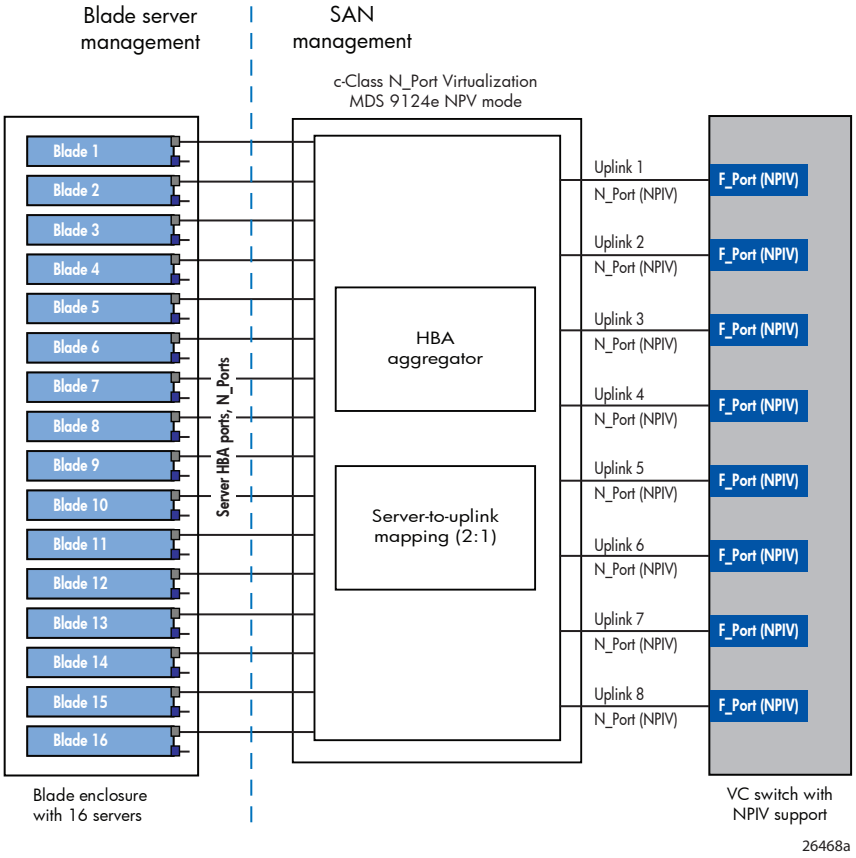



**Figure 53: Cisco MDS 8Gb Fabric Switch for BladeSystem c-Class**

**NOTE:** The names of the uplink ports (N_Ports 1 through 8) in **Cisco MDS 8Gb Fabric Switch for BladeSystem c-Class** are symbolic only. See the NPV documentation for the actual port numbers.

# HPE BladeSystem c3000 enclosure considerations

Consider the following when using the 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.



https://t.me/learningnets

# HBA N_Port ID 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 16 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, and H-series switches. For information about setting up switches for use with HBA NPIV, see the switch documentation.
- 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 54: VOS with HBA NPIV enabled** on page 193 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.



https://t.me/learningnets

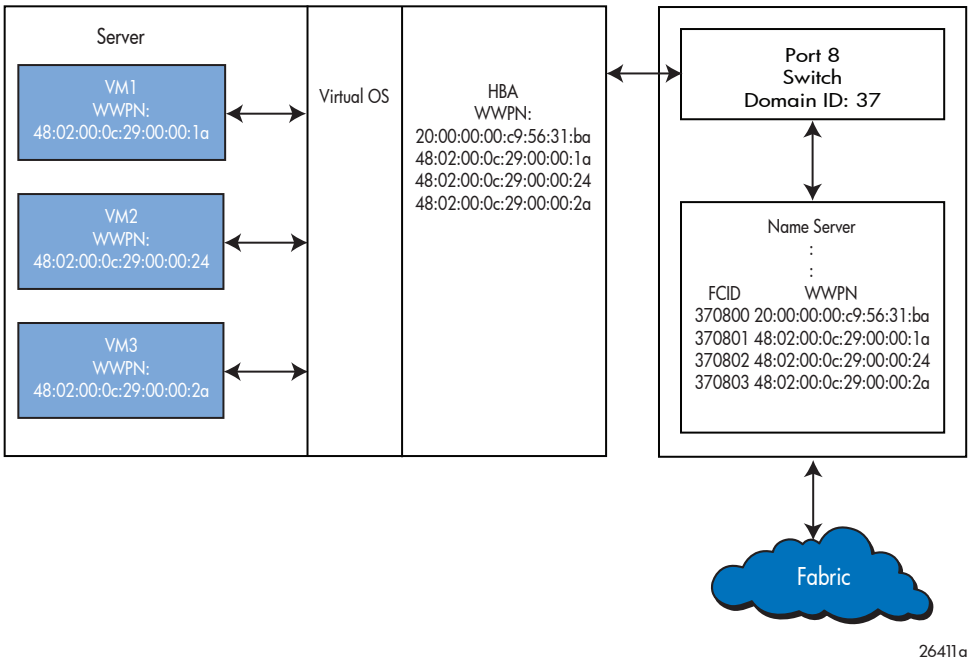



**Figure 54: VOS with HBA NPIV enabled**

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 54: VOS with HBA NPIV enabled** on page 193, 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.



https://t.me/learningnets

| NonStop servers | Storage systems |
|---|---|
| S-series servers:<br>• S760, S76000<br>• S78, S780, S7800, S78000<br>• S86000, S88000<br>NS-series servers:<br>• NS1000, NS1200<br>• NS14000, NS14200<br>• NS16000, NS16000CG, NS16200<br>NonStop Integrity servers:<br>• NS2000, NS2000T/NS2000CG<br>• NS2100<br>• NS2200, NS2200T/NS2200ST<br>• NS2300<br>• NS2400, NS2400T, NS2400ST<br>• NS3000AC<br>• NS5000T/NS5000CG<br>NonStop Integrity BladeSystem servers:<br>• NB50000c, NB50000c-cg<br>• NB54000c, NB54000c-cg<br>• NB56000c, NB56000c-cg | XP disk arrays:<br>• XP10000, XP12000 (RAID500)<br>• XP20000, XP24000 (RAID600)<br>• P9500<br>• XP7 |

There are three types of I/O interfaces used to connect NonStop servers to XP disk arrays, see Supported I/O modules for XP connectivity with specific NonStop systems table.

- **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


https://t.me/learningnets

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 HPE Enterprise Storage System (ESS) connections.

- **Storage CLIM** — 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 HPE 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 74: 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 |



https://t.me/learningnets

### Table 75: 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 76: NonStop BladeSystem with XP high-availability configurations using CLIMs (HPE Integrity NonStop NB50000c 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 77: Supported I/O modules for XP connectivity with specific NonStop systems

| NonStop servers | CLIM enclosure | IOAMEwith FCSAs installed | VIO enclosure |
|---|---|---|---|
| NonStop BladeSystems | Yes | Yes | No |
| NS5000T/NS5000CG | Yes | No | Yes |

*Table Continued*



https://t.me/learningnets

| NonStop servers | CLIM enclosure | IOAMEwith FCSAs installed | VIO enclosure |
|---|---|---|---|
| NS3000AC | Yes | No | Yes |
| NS2200, NS2400 | Yes | No | Yes |
| NS2100, NS2300 | Yes | No | Yes |
| NS2000 | Yes | No | Yes |
| NS16000, NS16200 | No | Yes | No |
| NS14000, NS14200 | No | Before December 2006<br>NS14000 only | After December 2006<br>NS14000 and NS14200 |
| NS1000, NS1200 | No | Before December 2006<br>NS1000 only | After December 2006<br>NS1000 and NS1200 |

### Table 78: 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 **Figure 55: Minimum direct host attach IOAME configuration for XP storage systems** on page 200, **Figure 58: Minimum SAN IOAME configuration for XP storage systems** on page 202, **Figure 61: SAN IOAME configuration with logical and physical redundancy for XP storage systems** on page 203, and **Figure 64: SAN IOAME configuration (two cascaded switches) with logical and physical redundancy for XP storage systems** on page 205. 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 **Figure 56: Minimum direct host attach VIO configuration for XP storage systems (NS1000, NS14000)** on page 201 and **Figure 59: Minimum SAN VIO configuration for XP storage systems (NS1000, NS14000)** on page 202. Logical and physical redundancy of the storage system, as shown in **Figure 62: SAN VIO configuration with logical and physical redundancy for XP storage systems (NS1000, NS14000)** on page 204 and **Figure 65: SAN VIO configuration (two cascaded switches) with logical and physical redundancy for XP storage systems (NS1000, NS14000)** on page 206, requires the addition of one expansion 4-port Fibre Channel PIC per VIO enclosure (total of two expansion PICs per server).<br>For optimal I/O performance, use a separate I/O path for each logical disk volumes. |

*Table Continued*



https://t.me/learningnets

| Rule number | Description |
|---|---|
| 5 | The following restrictions apply to CLIMs:<br>• 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.<br>• The XP LUN range for each port is 0 to 499.<br>• The maximum number of XP ports for each CLIM is four.<br>• The maximum number of mirrored XP volumes is 256 with two CLIMs, and 512 with four CLIMs. |
| 6 | Each LDEV requires two LUNs. |
| 7 | No boot support. |
| 8 | Contact Hewlett Packard Enterprise support for host mode requirements on storage system ports. |
| 9 | For more information on supported B-series and C-series switches, contact Hewlett Packard Enterprise 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.<br>**NOTE:** For NS1000 and NS1200 servers, expanded ports are available only to customers who have the 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. |

*Table Continued*



https://t.me/learningnets

| Rule number | Description |
|---|---|
| 14 | **Direct host attach**<br>• Hewlett Packard Enterprise 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).<br>• For high availability, the primary (P) LDEVs and mirror (M) LDEVs must be configured on separate array ACP pairs.<br>• 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.<br>• The 2 Gb FCSAs (IOAME) are supported with 1 Gb/2 Gb CHIPs for XP 10000/12000/20000/24000 and with 4 Gb CHIPs for XP10000/12000/20000/24000.<br>• The 2 Gb Fibre Channel PICs (VIO) are supported with 1 Gb/2 Gb CHIPs for XP 10000/12000/20000/24000 and with 4 Gb CHIPs for XP10000/12000/20000/24000.<br>• The 4 Gb Fibre Channel HBAs (in CLIMs) are supported with 1 Gb/2 Gb CHIPs for XP 10000/12000/20000/24000 and with 4 Gb CHIPs for XP10000/12000/20000/24000. |
| 15 | **High-availability SAN**<br>• Requires dual-redundant SAN fabrics (level 4, NSPOF high-availability SAN configuration). For information about data availability levels, see **Data availability** on page 40. Each fabric consists of either a single switch or two cascaded switches, as shown in **Figure 58: Minimum SAN IOAME configuration for XP storage systems** on page 202 through **Figure 65: SAN VIO configuration (two cascaded switches) with logical and physical redundancy for XP storage systems (NS1000, NS14000)** on page 206.<br>• A single fabric supports a maximum of three switches.<br>• 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.<br>• Configure WWN-based zoning only.<br>• Only NonStop homogeneous connections are allowed to the same zone.<br>• Heterogeneous operating systems can share the same switch or SAN if they are in different zones.<br>• Hewlett Packard Enterprise 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).<br>• For high availability, primary (P) LDEVs and mirror (M) LDEVs must be configured on separate array ACP pairs. |



https://t.me/learningnets

| Rule number | Description |
|---|---|
| | • 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.<br>• Hewlett Packard Enterprise 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.<br>• 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 XP10000/12000/20000/24000 and with 4 Gb CHIPs for XP10000/12000/20000/24000. |

**Figure 55: Minimum direct host attach IOAME configuration for XP storage systems** on page 200 shows a minimum direct host attach configuration with an IOAME.

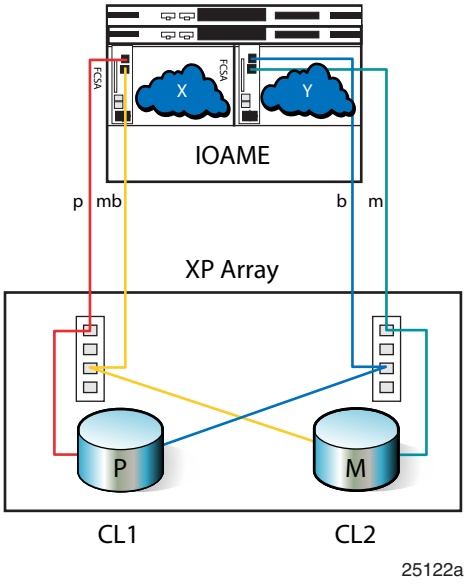



**Figure 55: Minimum direct host attach IOAME configuration for XP storage systems**

**Figure 56: Minimum direct host attach VIO configuration for XP storage systems (NS1000, NS14000)** on page 201 shows a minimum direct host attach configuration with VIO enclosures.



https://t.me/learningnets

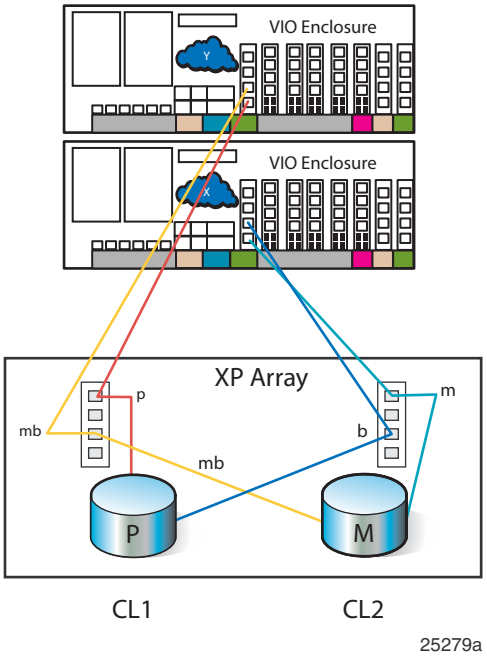



**Figure 56: Minimum direct host attach VIO configuration for XP storage systems (NS1000, NS14000)**

**Figure 57: Minimum direct host attach CLIM configuration for XP storage systems** on page 201 shows a minimum direct host attach configuration with CLIMs.

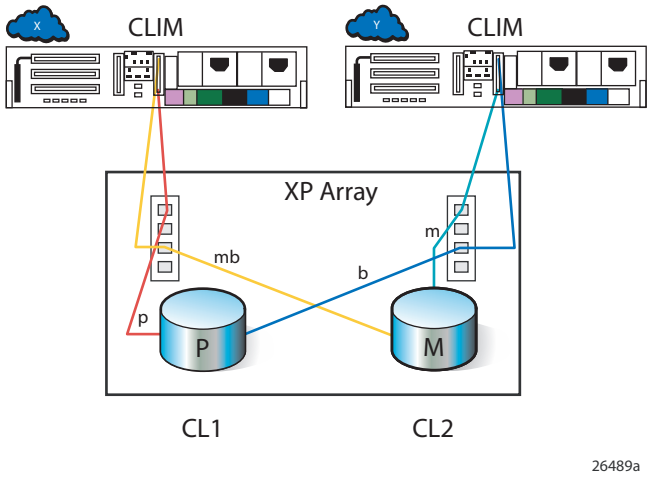



**Figure 57: Minimum direct host attach CLIM configuration for XP storage systems**

**Figure 58: Minimum SAN IOAME configuration for XP storage systems** on page 202 shows a minimum SAN configuration with an IOAME.



https://t.me/learningnets

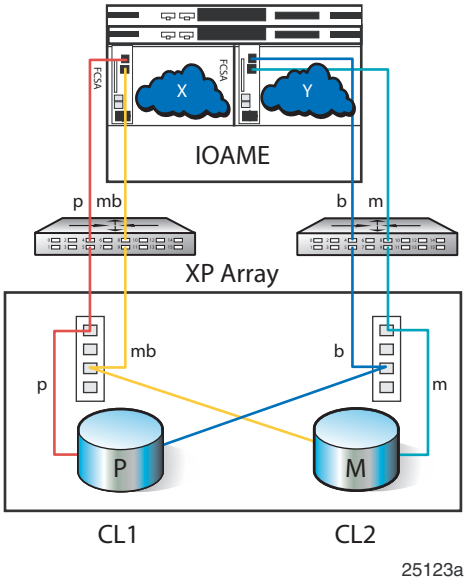



**Figure 58: Minimum SAN IOAME configuration for XP storage systems**

**Figure 59: Minimum SAN VIO configuration for XP storage systems (NS1000, NS14000)** on page 202 shows a minimum SAN configuration with VIO enclosures.

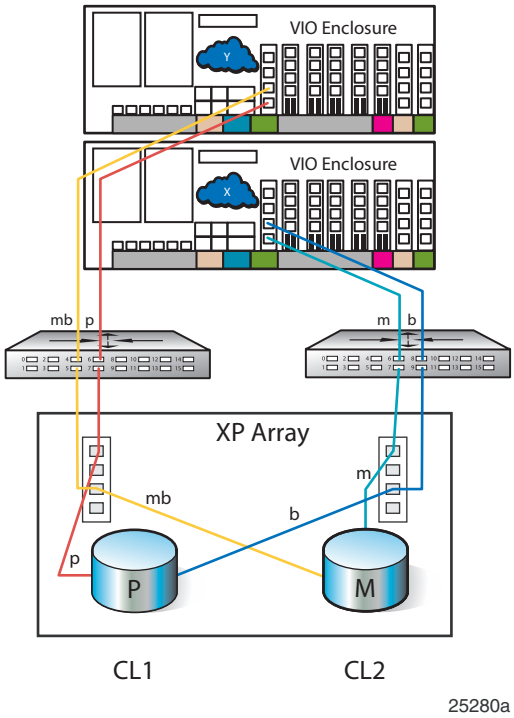



**Figure 59: Minimum SAN VIO configuration for XP storage systems (NS1000, NS14000)**

**Figure 60: Minimum SAN CLIM configuration for XP storage systems** on page 203 shows a minimum SAN configuration with CLIMs.



https://t.me/learningnets

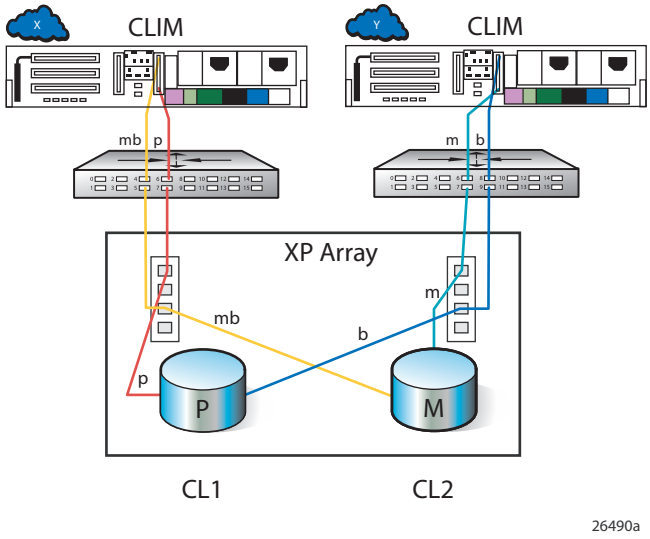



**Figure 60: Minimum SAN CLIM configuration for XP storage systems**

**Figure 61: SAN IOAME configuration with logical and physical redundancy for XP storage systems** on page 203 shows a configuration with physical IOAME redundancy.

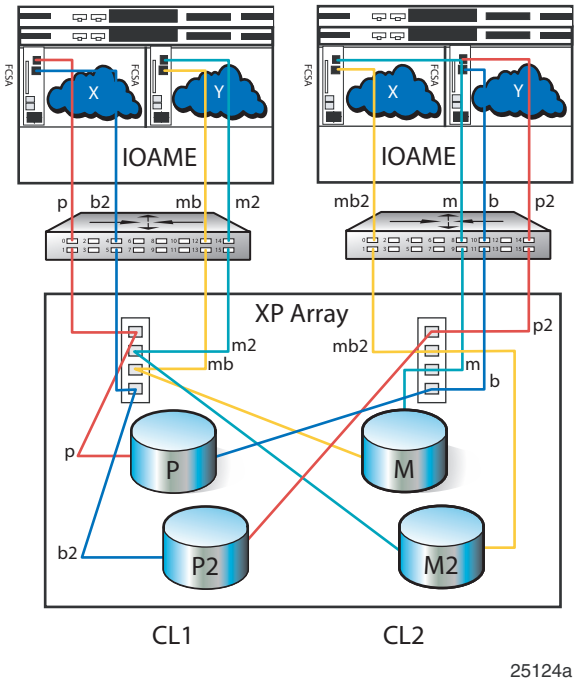



**Figure 61: SAN IOAME configuration with logical and physical redundancy for XP storage systems**

**Figure 62: SAN VIO configuration with logical and physical redundancy for XP storage systems (NS1000, NS14000)** on page 204 shows a SAN configuration with VIO Fibre Channel PIC redundancy.



https://t.me/learningnets

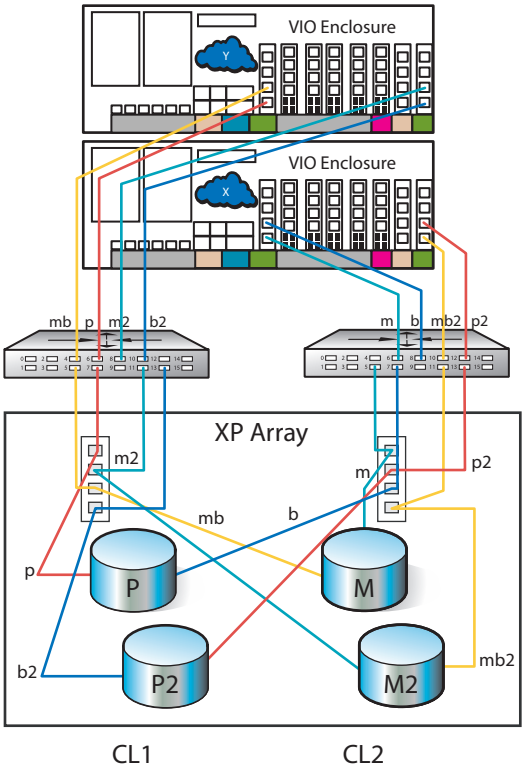



**Figure 62: SAN VIO configuration with logical and physical redundancy for XP storage systems (NS1000, NS14000)**

**Figure 63: SAN CLIM configuration with logical and physical redundancy for XP storage systems** on page 204 shows a SAN configuration with CLIM physical redundancy.

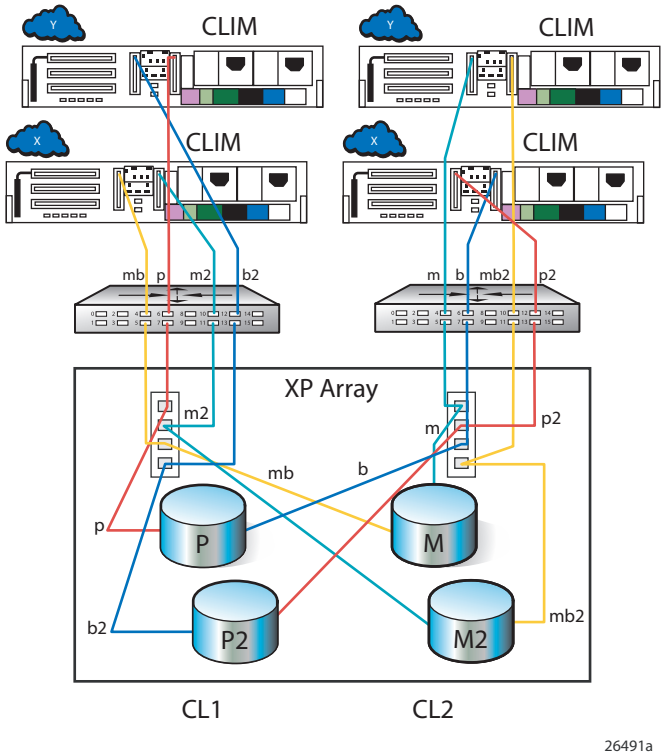



**Figure 63: SAN CLIM configuration with logical and physical redundancy for XP storage systems**



https://t.me/learningnets

**Figure 64: SAN IOAME configuration (two cascaded switches) with logical and physical redundancy for XP storage systems** on page 205 shows a configuration with physical IOAME redundancy.

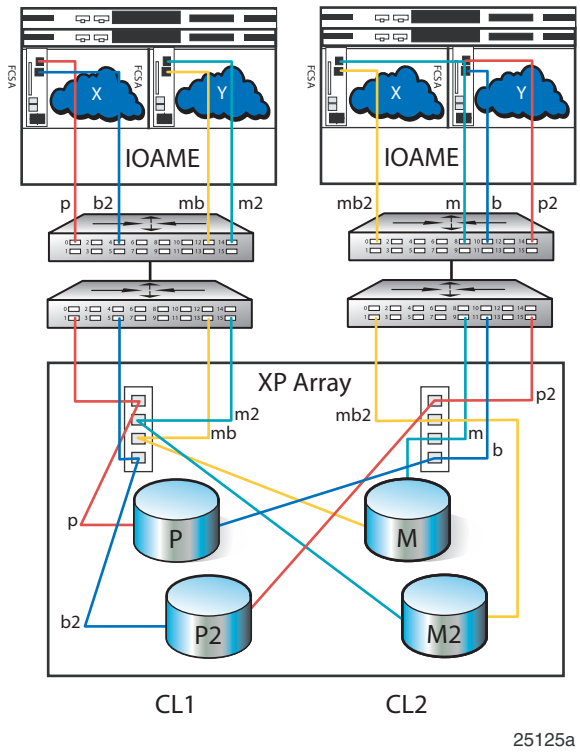



**Figure 64: SAN IOAME configuration (two cascaded switches) with logical and physical redundancy for XP storage systems**

**Figure 65: SAN VIO configuration (two cascaded switches) with logical and physical redundancy for XP storage systems (NS1000, NS14000)** on page 206 shows a SAN configuration (two cascaded switches) with VIO Fibre Channel PIC redundancy.



https://t.me/learningnets

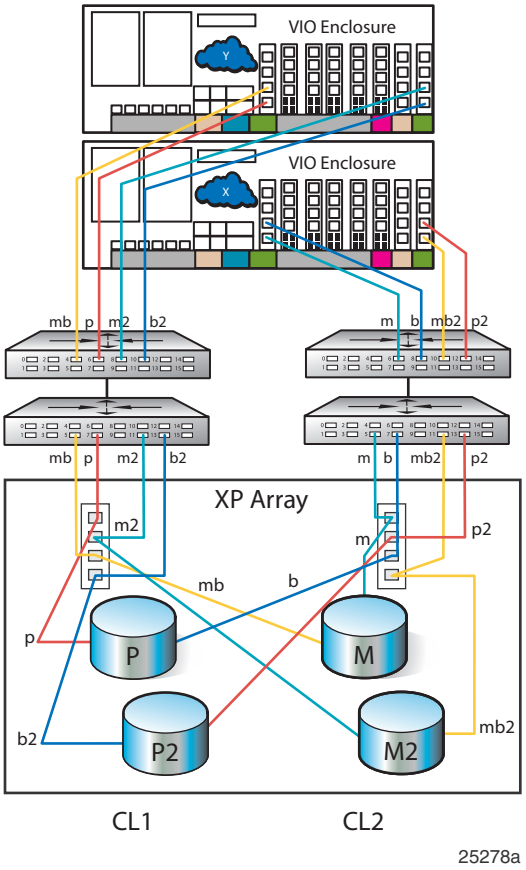



**Figure 65: SAN VIO configuration (two cascaded switches) with logical and physical redundancy for XP storage systems (NS1000, NS14000)**

**Figure 66: SAN CLIM configuration (two cascaded switches) with logical and physical redundancy for XP storage systems** on page 207 shows a SAN (two cascaded switches) configuration with CLIM physical redundancy.



https://t.me/learningnets

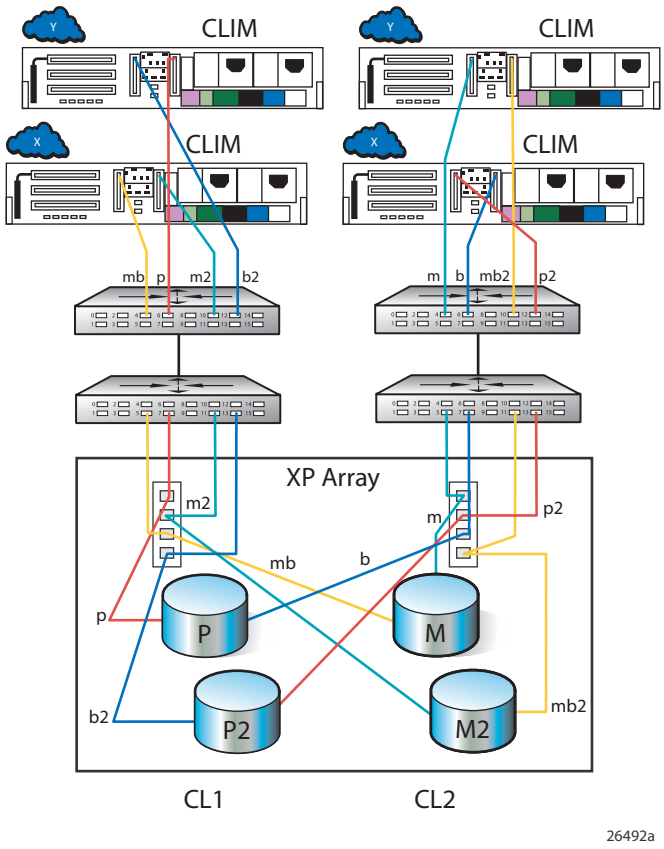



**Figure 66: SAN CLIM configuration (two cascaded switches) with logical and physical redundancy for XP storage systems**

# 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.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access.



https://t.me/learningnets

### Table 79: HP-UX SAN configuration rules

| Storage systems[1] | HP-UX SAN rules |
|---|---|
| All supported | • Supports HPE Serviceguard Clusters.<br>• Zoning is required when HP-UX is used in a heterogeneous SAN with other operating systems.<br>• Supports boot from SAN. See **P6000/EVA SAN boot support** on page 249 and **P9000/XP SAN boot support** on page 256.<br>• Supports connection to a common server for mixed storage system types, see **Common SAN storage coexistence** on page 224.<br>• Supports multipathing high-availability configuration in multiple fabrics or in a single fabric with zoned paths, see **Data availability** on page 40. |
| P2000 G3 FC | • Host name profile must be set to HP-UX<br>• Support for HP-UX 11i v2 (June 2008 or later) using PvLinks<br>• Support for HP-UX 11i v3 using native multipathing |
| MSA2050/2052 | • Host name profile must be set to HP-UX<br>• Support for HP-UX 11i v3 using native multipathing |
| MSA2040/2042/1040 | • Host name profile must be set to HP-UX<br>• Support for HP-UX 11i v3 using native multipathing |

*Table Continued*



https://t.me/learningnets

| Storage systems[1] | HP-UX SAN rules |
|---|---|
| P6300 EVA<br>P6350 EVA<br>P6500 EVA<br>P6550 EVA | • 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.<br>**NOTE:** Secure Path support for HP-UX 11i v2 is not available on the P6XX0 arrays.<br>• For P6000 Continuous Access configuration information, see **HPE P6000 Continuous Access SAN integration** on page 247. |
| 3PAR StoreServ 10000/7000; T-Class | • Supported with HP-UX 11iv2 and 11iv3<br>• Supports boot from SAN, see **HPE 3PAR SAN boot support** on page 271.<br>• Zoning by HBA is required when used in a heterogeneous SAN, including other operating systems and other storage system families or types.<br>• All hosts must have the appropriate Host Operating System type parameter set (Host Persona) and the required host settings described in the *3PAR HP-UX Implementation Guide*. |

[1] Unlisted but supported storage systems have no additional SAN configuration restrictions. For the latest support information, contact a Hewlett Packard Enterprise storage representative.

### Table 80: HP-UX storage system, HBA, and multipath software coexistence support

| | | | | P63xx/P65xx EVA | | | | | | | | XP7 | | P9500 | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | P2000 G3 | | EVA4x00/6x00/8x00[3] | | | 10000/12000/20000/24000 | | | 3PAR | | | | | |
| Notes[1] | | PL | VD | SP | PL | VD | SP | PL | VD | PL | VD | PL | VD | PL | VD |
| **P2000 G3** | **PL** | S | - | S | S | - | S | S | - | S | - | S | - | S | - |
| | **VD** | - | S | - | - | S | - | - | S | - | S | - | - | - | S |
| **P63xx/P65ixx EVA[2]** | **SP[2]** | S | - | S | S | - | S | S | - | - | - | - | - | S | - |
| **XP10000/12000 / 20000/24000** | **PL** | S | - | - | S | - | - | S | - | S | - | S | - | S | - |
| | **VD** | - | S | - | - | S | - | - | S | - | S | - | - | - | S |
| **P9500** | **PL** | S | - | S | S | - | S | S | - | S | - | S | - | S | - |
| | **VD** | - | S | - | - | S | - | - | S | - | S | - | - | - | S |
| **3PAR** | **PL** | S | - | - | S | - | - | S | - | S | - | S | - | S | - |

*Table Continued*



https://t.me/learningnets

| Notes[1] | | P2000 G3 | | P63xx/P65xx EVA<br>EVA4x00/6x00/8x00[3] | | | 10000/12000/20000/24000 | | | 3PAR | | XP7 | | P9500 | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | PL | VD | SP | PL | VD | SP | PL | VD | PL | VD | PL | VD | PL | VD |
| | VD | - | S | - | - | S | - | - | S | - | S | - | - | - | S |
| XP7 | PL | S | - | - | S | - | - | S | - | S | - | S | - | S | - |
| | VD | - | - | - | - | - | - | - | - | - | - | - | - | - | - |

[1] General rules:

- XP24000/20000/12000/10000 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 a Hewlett Packard Enterprise storage representative.
- All HP-UX 11i v3 supported storage systems are supported on the same server and HBA using native multipathing software.

[2] P63xx EVAs and P65xx EVAs do not support Secure Path.

[3] EVA4100/6100/8100 requires XCS firmware 6.110 or later.

**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

**NOTE:** Secure Path for HP-UX 11.23 is no longer provided through HPE. 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** on page 224.

# HPE OpenVMS SAN rules

This section describes the SAN rules for HPE OpenVMS.

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



https://t.me/learningnets

### Table 81: HPE OpenVMS SAN configuration rules

| Storage systems[1] | OpenVMS SAN rules |
|---|---|
| All supported | • Zoning is required when OpenVMS is used in a heterogeneous SAN with other operating systems.<br>• Supports OpenVMS Clusters.<br>• Supports active/active and active/passive failover mode. A multipathing driver is embedded in the operating system.<br>• Supports boot from SAN. For more information, see **P6000/EVA SAN boot support** on page 249 and **P9000/XP SAN boot support** on page 256.<br>• Supports multipathing high-availability configuration in multiple fabrics or in a single fabric with zoned paths. |
| P6300 EVA<br>P6350 EVA<br>P6500 EVA<br>P6550 EVA | For P6000 Continuous Access configuration information, see **HPE P6000 Continuous Access SAN integration** on page 247. |
| P2000 G3 FC<br>P2000 G3 FC/iSCSI | • Host name profile must be set to OpenVMS.<br>• OpenVMS uid must be set on all volumes mapped to a OpenVMS host.<br>• LUN ID 0 is not available for OpenVMS hosts. |

[1] Unlisted but supported storage systems have no additional SAN configuration restrictions. For the latest support information, contact a Hewlett Packard Enterprise storage representative.

### Table 82: HPE OpenVMS storage system, HBA, and multipath software coexistence support

**Legend:** D = same server and different HBA; S = same server and HBA; — = not supported

| | | P63xx/P65xx EVA4x00/6x00/8x00 | XP |
|---|---|---|---|
| | | Native multipathing driver | |
| **P2000 G3 FC**<br>**P2000 G3 FC/iSCSI** | **Native multipathing driver** | S | S | S |
| **P63xx/P65xx EVA** | | S | S | S |
| **P9500, XP24000/20000, XP12000/10000** | | S | S | S |

For more information about storage system coexistence, see **Heterogeneous SAN storage system coexistence** on page 224.



https://t.me/learningnets

## Host-based volume shadowing

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:

- P6000 Continuous Access
- P9000 (XP) Continuous Access

# HPE Tru64 UNIX SAN rules

This section describes the SAN rules for HPE Tru64 UNIX.

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

### Table 83: Tru64 UNIX SAN configuration rules

| Storage systems[1] | Tru64 UNIX SAN rules |
|---|---|
| All supported | • Zoning is required when Tru64 UNIX is used in a heterogeneous SAN with other operating systems.<br>• Supports TruCluster Server.<br>• Supports boot from SAN. For more information, see **P6000/EVA SAN boot support** on page 249 and **P9000/XP SAN boot support** on page 256.<br>• Supports multipathing high-availability configuration in multiple fabrics or in a single fabric with zoned paths. |
| | **HPE P6000 Continuous Access SAN integration** on page 247. |

[1] Unlisted but supported storage systems have no additional SAN configuration restrictions. For the latest support information, contact a Hewlett Packard Enterprise storage representative.

### Table 84: HPE Tru64 UNIX storage system, HBA, and multipath software coexistence support

| | | XP |
|---|---|---|
| | | Native multipathing driver |
| **XP24000/20000,XP12000/10000** | **Native multipathing driver** | S |

**Legend:** S = same server and HBA

For more information about storage system coexistence, see **Heterogeneous SAN storage system coexistence**.

# Apple Mac OS X SAN rules

This section describes the SAN rules for Apple Mac OS X.



https://t.me/learningnets

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

### Table 85: Apple Mac OS X SAN configuration rules

| Storage systems | Apple Mac OS X SAN rules |
|---|---|
| P6300 EVA<br>P6350 EVA<br>P6500 EVA<br>P6550 EVA | • Zoning is required when Mac OS X is used in a heterogeneous SAN with other operating systems.<br>• Supports active/active and active/passive failover mode. A multipathing driver is embedded in the host adapter driver.<br>• Supports multipathing high-availability configuration in multiple fabrics or in a single fabric with zoned paths. |
| | **P6000/EVA rules** [1]<br><br>Command View host entry operating system: Apple Mac OS X<br><br>**NOTE:** For XCS 6.100 (or later) and for XCS earlier than 09003000, select **Custom** from the Operating System Type list, and then enter **00000002024000A8** as the Custom type. |
| P6300 EVA<br>P6350 EVA<br>P6500 EVA<br>P6550 EVA | **Multipathing rules**<br><br>• Hewlett Packard Enterprise recommends zoning single Host Adapter port WWPN with single array port WWPN.<br>• Single Host Adapter port WWPN zoned with multiple array port WWPNs to the same array is supported.<br>• Single Host Adapter port WWPN zoned with multiple arrays is supported.<br>• I/O load balancing can span multiple Host Adapter ports. |
| | **P6000/EVA rules**[1]<br><br>• Celerity FC-82EN, dual channel, x8 PCI Express<br>• Celerity FC-84EN, quad channel, x8 PCI Express<br>• Command View 6.0.2 (or later) (EVA4100/6100/8100)<br>• Command View 8.0 (or later) (EVA4400)<br>• Command View host entry operating system: Custom, custom type "00000002024000A8"<br>• P6000 Command View 9.4 (or later) (P6300/P6500) |

*Table Continued*



https://t.me/learningnets

| Storage systems | Apple Mac OS X SAN rules |
|---|---|
| P6300 EVA<br>P6350 EVA<br>P6500 EVA<br>P6550 EVA | **Host adapter rules**<br>• Supported only with ATTO 8 Gb and 4 Gb PCIe and PCI-x host adapters:<br>◦ Celerity FC-41XS, single channel, 64-bit/133 MHz PCI-X<br>◦ Celerity FC-42XS, dual channel, 64-bit/133 MHz PCI-X<br>◦ Celerity FC-41EL, single channel, x4 PCI Express<br>◦ Celerity FC-41ES, single channel, x4 PCI Express<br>◦ Celerity FC-42ES, dual channel, x4 PCI Express<br>◦ Celerity FC-44ES, quad channel, x8 PCI Express<br>For more information, see **http://www.attotech.com/hp**. |
| P2000 G3 FC | • Zoning is required when Apple OS X is used in a heterogeneous SAN with other operating systems.<br>• Supported only with the following host adapters:<br>◦ Celerity FC-82EN, dual channel, x8 PCI Express<br>◦ Celerity FC-84EN, quad channel, x8 PCI Express |
| 3PAR StoreServ 10000/7000 | • Zoning by HBA is required when used in a heterogeneous SAN, including other operating systems and other storage system families or types.<br>• All hosts must have the appropriate Host Operating System type parameter set (Host Persona) and the required host settings described in the *3PAR Apple OS X Implementation Guide*, available on the Hewlett Packard Enterprise Support Center website in the **Manuals for HPE 3PAR Operating System Software** section. |

[1] Contact a Hewlett Packard Enterprise 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.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access.



https://t.me/learningnets

### Table 86: IBM AIX SAN configuration rules

| Storage systems[1] | IBM AIX SAN rules |
|---|---|
| All supported | • Supports HACMP/ES Clusters. For specific version support, contact a Hewlett Packard Enterprise storage representative.<br>• Zoning is required when IBM AIX is used in a heterogeneous SAN with other operating systems.<br>• Supports active/active and active/passive failover mode. A multipathing driver is required.<br>• Supports multipathing high-availability configuration implemented in multiple fabrics or in a single fabric with zoned paths.<br>• For specific Secure Path zoning requirements, contact a Hewlett Packard Enterprise storage representative.<br>• Supports boot from SAN for XP storage systems. For more information, see **P9000/XP SAN boot support** on page 256.<br>• For P6000 Continuous Access configuration information, see **HPE P6000 Continuous Access SAN integration**. |
| 3PAR StoreServ 10000/7000; 3PAR T-Class | • All hosts must have the appropriate Host Operating System type parameter set (Host Persona) and the required host settings described in the *3PAR AIX Implementation Guide*.<br>• Supports HACMP/ES Clusters. For specific version support, contact a Hewlett Packard Enterprise Storage representative.<br>• Zoning is required when IBM AIX is used in a heterogeneous SAN with other Operating Systems.<br>• Active/Active failover mode supported. A multipath driver is required.<br>• Supports multipathing high-availability configuration implemented in multiple fabrics or in a single fabric with zoned paths. |

[1] Unlisted but supported storage systems have no additional SAN configuration restrictions. For the latest support information, contact a Hewlett Packard Enterprise storage representative.

### Table 87: IBM AIX storage system, HBA, and multipath software coexistence support

| | | P63xx/P65xx | | P9500, XP7, XP24000/20000, XP12000/10000 | | 3PAR[2] | |
|---|---|---|---|---|---|---|---|
| Notes[1] | | MPIO | Antemeta | HDLM | MPIO | 3PAR MPIO | MPIO |
| **P63xx/P65xx** | **MPIO** | S[3] | - | - | S[7] | D | S |
| | **Antemeta** | - | S[6,7] | - | - | - | - |

*Table Continued*



https://t.me/learningnets

| Notes[1] | | P63xx/P65xx | | P9500, XP7, XP24000/20000, XP12000/10000 | | 3PAR[2] | |
|---|---|---|---|---|---|---|---|
| | | MPIO | Antemeta | HDLM | MPIO | 3PAR MPIO | MPIO |
| P9500, XP7, XP24000/20000, XP12000/10000 | HDLM | - | - | S[7] | - | - | - |
| | MPIO | S[7] | - | - | S[7] | S | S |
| 3PAR[3] | 3PAR MPIO | D | - | - | S | S | - |
| | MPIO | S | - | - | S | - | S |

1 Connection to a common server with different HBAs requires separate HBA zones for XP and P6000/EVA.

2 3PAR storage requires that you enable both DynamicTracking and FastFail parameters on each of the IBM AIX FC HBAs connected to 3PAR storage. 3PAR OS 3.1.2 is not supported with IBM AIX 5.3. For more information, see the *3PAR AIX Implementation Guide.*

3 Requires AIX 5.3 (or later).

**Legend:** S = same server and HBA; D = same server and different HBA; — = not supported

For more information about storage system coexistence, see **Heterogeneous SAN storage system coexistence**.

# Linux SAN rules

This section describes the SAN rules for Linux.

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

### Table 88: Linux SAN configuration rules

| Storage systems[1] | Linux SAN rules |
|---|---|
| All supported | Supports multipathing high-availability configuration in multiple fabrics or in a single fabric with zoned paths. |
| P2000 G3 FC<br>MSA2050/2052 SAN<br>MSA2040/2042/1040 SAN | • Zoning is required when Linux is used in a heterogeneous SAN with other operating systems.<br>• For HBA parameter settings, see MSA storage system rules.<br>• Supports boot from SAN.<br>• Supports Device Mapper. |

*Table Continued*



https://t.me/learningnets

| Storage systems[1] | Linux SAN rules |
|---|---|
| XP10000<br>XP12000<br>XP20000<br>XP24000 | • Zoning is required when Linux is used in a heterogeneous SAN with other operating systems.<br>• Provides single-path support only.<br>• Supports clusters (32-bit/64-bit).<br>• Supports active/active or active/passives failover mode.<br>• Supports boot from SAN. For more information, see **P9000/XP SAN boot support** on page 256. |
| P6300 EVA<br>P6350 EVA<br>P6500 EVA<br>P6550 EVA | • Zoning is required when Linux is used in a heterogeneous SAN with other operating systems.<br>• Supports connection of single HBA servers.<br>• Supports boot from SAN. For more information, see **P6000/EVA SAN boot support** on page 249.<br>• For P6000 Continuous Access configuration information, see **HPE P6000 Continuous Access SAN integration** on page 247. |
| 3PAR StoreServ 10000/7000/8000/20000; 3PAR T-Class | • Zoning by HBA is required when used in a heterogeneous SAN, including other operating systems and other storage system families or types.<br>• All hosts must have the appropriate Host Operating System type parameter set (Host Persona) and the required host settings described in the 3PAR Linux Implementation Guide. |

[1] Unlisted but supported storage systems have no additional SAN configuration restrictions. For the latest support information, contact a Hewlett Packard Enterprise storage representative.

## Linux multipath software coexistence support

All HPE FC storage systems supported with Linux are supported for coexistence on the same server and HBA when using Device-Mapper Multipath software. The multipathing pmeter settings for each storage system are managed by Device Mapper.



https://t.me/learningnets

**Table 89: Linux storage system, HBA, and multipath software coexistence support**

| | | P2000 G3 | P63xx/P65xx EVA EVA4x00/6x00/8x00 | P9500, XP7 XP24000/20000/ 12000/10000 | 3PAR |
|---|---|---|---|---|---|
| | | Device-Mapper Multipath | | | |
| **P2000 G3** | **Device-Mapper Multipath** | S | S | S | S |
| **P63 xx/P65xx** | | | | | |
| **XP24000/20000/12000/10000** | | S | S | S | S |
| **3PAR** | | S | S | S | S |

**Legend:** D = same server and different HBA; S = same server and HBA; - = not supported

# 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.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access.

**Table 90: Microsoft Windows storage system, HBA, and multipath software coexistence support**

| | | P2000 G3 | P63xx/P65xx EVA | | XP10000/12000/ XP20000/24000/ P9500, XP7 | | 3PAR | |
|---|---|---|---|---|---|---|---|---|
| | | MS MPIO DSM[1] | HPE MPIO FF[2] | MS MPIO DSM | HPE MPIO FF | MS MPIO DSM | 3PAR MPIO | MS MPIO DSM[3] |
| **P2000 G3** | **MS MPIO DSM** | S | S | S | S | S | - | S |
| **P63xx/P65xx EVA** | **HPE MPIO FF** | S | S | - | S | S | S | S |

*Table Continued*



https://t.me/learningnets

| | | P2000 G3 | P63xx/P65xx EVA | | XP10000/12000/ XP20000/24000/ P9500, XP7 | | 3PAR | |
|---|---|---|---|---|---|---|---|---|
| | | MS MPIO DSM[1] | HPE MPIO FF[2] | MS MPIO DSM | HPE MPIO FF | MS MPIO DSM | 3PAR MPIO | MS MPIO DSM[3] |
| | MS MPIO DSM | S | - | S | S | S | - | S |
| XP10000/12000/ XP20000/24000/P9500, XP7 | HPE MPIO FF | S | S | S | S | - | S | S |
| | MS MPIO DSM | S | S | S | - | S | - | S |
| 3PAR | HPE MPIO FF | - | S | - | S | - | S | - |
| | MS MPIO DSM | S | S | S | S | S | - | S |

1 Microsoft MPIO Device Specific Module

2 If 3PAR storage is in the same zone, you must use HPE MPIO DSM version 4.x or later if you use the HPE MPIO DSM.

3 Microsoft MPIO DSM for 3PAR storageis supported with Windows 2008 and later.

**Legend:** S = same server and HBA; — = not supported

# Oracle Solaris SAN rules

This section describes the SAN rules for Oracle Solaris.

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



https://t.me/learningnets

### Table 91: Oracle Solaris SAN configuration rules

| Storage systems[1] | Solaris SAN rules |
|---|---|
| All supported | • Zoning is required when Solaris is used in a heterogeneous SAN with other operating systems.<br>• Supports Veritas Cluster Server, Solaris Cluster, and PRIMECLUSTER. Contact a Hewlett Packard Enterprise storage representative for version and storage system support. Not all storage system types are supported on all cluster solutions. 3PAR storage does not support PRIMECLUSTER.<br>• Each cluster must be in its own zone.<br>• Supports connection of single HBA servers.<br>• Supports multipathing high-availability configuration implemented in multiple fabrics or in a single fabric with zoned paths.<br>• Requires a multipathing driver for active/active or active/passive failover with two or more paths. 3PAR storage does not support active/passive failover. |
| P2000 G3 | • Zoning is required when Oracle Solaris is used in a heterogeneous SAN with other operating systems.<br>• Single path is not supported. MPxIO is required.<br>• Retry count must be set to 10; entry in `/kernel/drv/sd.conf` must be set as follows:<br>`SD-config-list="HP MSA","retries-timeout: 10"; SD-config-list="HP P2000","retries-timeout: 10";`<br>**NOTE:** There should be six spaces between "HP" and the array model, "MSA" or "P2000." |

*Table Continued*



https://t.me/learningnets

| Storage systems[1] | | Solaris SAN rules |
|---|---|---|
| P6300 EVA | P6350 EVA<br>P6500 EVA<br>P6550 EVA<br>XP10000<br>XP12000<br>XP20000<br>XP24000 | Supports boot from SAN. For more information, see **P6000/EVA SAN boot support** and **P9000/XP SAN boot support**.<br>• Requires MPxIO, Veritas Volume Manager DMP, or HDLM multipathing driver. For version support, contact a Hewlett Packard Enterprise storage representative. For multipathing coexistence support, see Table Oracle Solaris storage system, HBA, and multipath software coexistence support.<br>• For P6000 Continuous Access configuration information, see **HPE P6000 Continuous Access SAN integration** on page 247. |
| 3PAR StoreServ 10000/7000; 3PAR T-Class | | • For information about configuring different storage system families on the same server with different HBAs, see **Common server, different HBAs** on page 225.<br>• Zoning by HBA is required when used in a heterogeneous SAN, including other operating systems and other storage system families or types.<br>• For information about Solaris support and supported Fibre Channel components, see the SPOCK website at **http://www.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access.<br>• All hosts must have the appropriate Host Operating System type parameter set (Host Persona) and the required host settings described in the *HPE 3PAR Solaris Implementation Guide*. |

[1] Unlisted but supported storage systems have no additional SAN configuration restrictions. For the latest support information, contact a Hewlett Packard Enterprise storage representative.

### Table 92: Oracle Solaris storage system, HBA, and multipath software coexistence support

| Notes[1] | | P2000 G3 | P63xx/P65xx | | P9500, XP7, XP24000/20000, XP12000/10000 | | | 3PAR | |
|---|---|---|---|---|---|---|---|---|---|
| P2000 G3 | | | | | | | | | |
| MP | | MP | MP | VD | H | MP | VD[2] | MP | VD |
| S | S | D | D | S | D | S | D | | |
| **P63xx/P65xx EVA** | **MP** | S | S | D | - | S | D | S | D |
| | **VD** | D | D | S | - | D | S | D | S |
| **P9500, XP7, XP24000/20000, XP12000/10000** | **H** | D | - | - | S | - | - | - | - |
| | **MP** | S | S | D | - | S | D | S | D |
| | **VD** | D | D | S | - | D | S | D | S |
| **3PAR** | **MP** | S | S | D | - | S | D | S | - |
| | **VD** | D | D | S | - | D | S | - | S |



https://t.me/learningnets

[1] General notes:

- Connection to a common server with different HBAs requires septe HBA zones for 3PAR, XP, 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, Oracle-branded HBAs are required for MPxIO.

[2] Veritas DMP is not supported with XP7

**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. For more information, see the *HPE 3PAR Solaris Implementation Guide*.

For more information about storage system coexistence, see **Heterogeneous SAN storage system coexistence** on page 224.

# 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.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access.

### Table 93: VMware ESX SAN configuration rules

| Storage systems[1] | ESX SAN rules |
|---|---|
| All supported | Zoning is required when ESX is used in a heterogeneous SAN with other operating systems. |
| P6300 EVA<br>P6350 EVA<br>P6500 EVA<br>P6550 EVA<br>P2000 G3 FC<br>XP10000<br>XP12000<br>XP20000<br>XP24000 | Supports Microsoft Windows Server 2003 base, with a maximum of two nodes per cluster. It does not support Windows Server 2003 SP1.<br>**NOTE:** Not all listed storage systems are supported on all versions of VMware ESX. |
| 3PAR StoreServ 10000/7000/8000/2000<br>3PAR T-Class | type parameter set (Host Persona) and the required host settings described in the *HPE 3PAR VMware ESX Implementation Guide*. |

[1] Unlisted but supported storage systems have no additional SAN configuration restrictions. For the latest support information, contact a Hewlett Packard Enterprise storage representative.

## VMware ESX HBA and multipath software coexistence support

All supported storage systems are supported on the same server using the same HBA and native multipathing software.



https://t.me/learningnets

# Citrix Xen SAN rules

This section describes the SAN rules for Xen.

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

### Table 94: Xen SAN configuration rules

| Storage systems[1] | Xen SAN rules |
|---|---|
| All supported | Supports multipathing high-availability configuration in multiple fabrics or in a single fabric with zoned paths. |
| P2000 G3 FC<br>MSA2050/2052 SAN<br>MSA2040/2042/1040 SAN | • Zoning is required when Xen is used in a heterogeneous SAN with other operating systems.<br>• For HBA parameter settings, see MSA storage system rules table.<br>• Boot from SAN is not supported.<br>• Supports Inbox Device Mapper. |
| XP10000<br>XP12000<br>XP20000<br>XP24000 | • Zoning is required when Xen is used in a heterogeneous SAN with other operating systems.<br>• Supports active/active or active/passives failover mode.<br>• Boot from SAN is not supported. |
| P6300 EVA<br>P6500 EVA | • Zoning is required when Xen is used in a heterogeneous SAN with other operating systems.<br>• Supports connection of single HBA servers.<br>• Boot from SAN is not supported.<br>• For P6000 Continuous Access configuration information, see **HPE P6000 Continuous Access SAN integration** on page 247. |
| 3PAR StoreServ 10000/7000/8000/20000<br>3PAR T-Class | • 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** on page 225.<br>• Zoning by HBA is required when used in a heterogeneous SAN, including other operating systems and other storage system families or types.<br>• All hosts must have the appropriate Host Operating System type parameter set (Host Persona) and the required host settings described in the *HPE 3PAR Citrix XenServer Implementation Guide*. |



https://t.me/learningnets

[1] Unlisted but supported storage systems have no additional SAN configuration restrictions. For the latest support information, contact a Hewlett Packard Enterprise storage representative.

# Heterogeneous SAN storage system coexistence

Hewlett Packard Enterprise 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** on page 105
- **C-series switches and fabric rules** on page 137
- **H-series switches and fabric rules** on page 151
- **Heterogeneous server rules** on page 179

## Common SAN storage coexistence

To configure different HPE 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
- P9500, XP
- 3PAR StoreServ 10000/7000; 3PAR T-Class
- Third-party storage systems
- MSA 2040
- MSA 1040
- MSA 2050
- 3PAR StoreServ 20000/8000

To configure common server access to different storage system types, see **Common server access, different storage system types** on page 224.

## Common server access, different storage system types

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 HPE storage and third-party storage is supported if you comply with the following rules:



https://t.me/learningnets

- You must use Hewlett Packard Enterprise-supported HBAs, parameter settings, drivers, and firmware to support HPE storage and third-party storage on the same HPE or multi-vendor server.
- Use different Hewlett Packard Enterprise-supported HBA vendors if the required parameter settings, drivers, and firmware are different for HPE storage and the third-party storage.
- Both HPE storage and third-party storage must use the same multipath software.

## 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 and firmware versions
- Operating systems driver versions and parameter settings
- Both HPE storage and third-party storage must use the same multipath software

See the storage system, HBA, and multipath software coexistence table specific to each operating system beginning with **HP-UX SAN rules** on page 207.

## Common server, different HBAs

For information about supported storage system types using common server access with different HBAs (Emulex and QLogic only), see the storage system, HBA, and multipath software coexistence table specific to each operating system beginning with **HP-UX SAN rules** on page 207.

**NOTE:**

HPE 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).

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 a Hewlett Packard Enterprise 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, see the 3PAR implementation guides for each of the server operating systems.



https://t.me/learningnets

# 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 — **Zoning limits and enforcement for H-series Fibre Channel switches**
- B-series — **Zoning enforcement for B-series Fibre Channel switches and MP Router LSANs**
- C-series — **Zoning limits and enforcement for C-series Fibre Channel switches**
- **Configuration parameters** on page 242 for P6000/EVA
- **Zoning** on page 248 for P6000/EVA
- **Zoning** on page 255 for XP
- **Zoning enforcement** on page 411
- **Zoning guidelines** on page 412
- **EBS zoning** on page 414
- **Zone naming** on page 414


https://t.me/learningnets

# MSA storage system rules

This chapter describes specific rules for the following entry-level storage systems:

- HPE MSA 2050/2052 SAN
- HPE MSA 1050 FC
- HPE MSA 2040/2042 SAN
- HPE MSA 1040 FC
- MSA P2000 FC/iSCSI

For the iSCSI rules for the MSA2000i, see **HPE StorageWorks MSA family of iSCSI SAN arrays** on page 349.

## HPE MSA storage system configurations

### Table 95: HPE MSA storage system configurations

| Model | Array chassis | Expansion options | Expansion maximum |
|---|---|---|---|
| MSA 2050/2052 SAN | Preconfigured dual controller LFF or SFF | MSA 2050 SFF disk enclosure<br>MSA 2050 LFF disk enclosure<br>MSA 2040 LFF disk enclosure **#GUID-D7147C7F-2016-0901-03BB-000000000E2C/fn1**<br>D2700 SFF disk enclosure**#GUID-D7147C7F-2016-0901-03BB-000000000E2C/fn1** | 8 total enclosures<br>96 LFF drives<br>192 SFF drives |
| MSA 1050 FC | Preconfigured dual controller LFF or SFF | MSA 2050 SFF disk enclosure<br>MSA 2050 LFF disk enclosure<br>MSA 2040 LFF disk enclosure**#GUID-D7147C7F-2016-0901-03BB-000000000E2C/fn1**<br>D2700 SFF disk enclosure**#GUID-D7147C7F-2016-0901-03BB-000000000E2C/fn1** | 4 total enclosures<br>48 LFF drives<br>96 SFF drives |

*Table Continued*



https://t.me/learningnets

| Model | Array chassis | Expansion options | Expansion maximum |
|---|---|---|---|
| MSA 1040 FC | Preconfigured dual controller LFF or SFF | MSA 2040 LFF disk enclosure<br>D2700 SFF disk enclosure<br>P2000 LFF drive enclosure | 4 total enclosures<br>48 LFF drives<br>99 SFF drives |
| MSA 2040/2042 SAN | MSA 2040 SFF<br>MSA 2040 LFF | MSA 2040 LFF disk enclosure<br>D2700 SFF disk enclosure<br>P2000 LFF drive enclosure | 8 total enclosures<br>96 LFF drives<br>199 SFF drives |
| P2000 FC<br>P2000 G3 FC/iSCSI | P2000 3.5-in drive bay (LFF)<br>P2000 2.5-in drive bay (SFF)<br>2012 3.5-in drive bay (LFF, upgrade only)<br>2024 2.5-in drive bay (SFF, upgrade only) | P2000 LFF drive enclosure<br>D2700 SFF disk enclosure<br>MSA2000 LFF drive enclosure<br>MSA70 SFF drive enclosure | 8 total enclosures or 149 drives<br>96 LFF drives<br>149 SFF drives |

## Heterogeneous SAN support

The MSA 2050/2052, MSA 2040/1040, MSA 1050, and P2000 G3 support heterogeneous access with any of the operating systems. Standalone servers and clustered servers are supported on the same storage system.



https://t.me/learningnets

**Table 96: MSA Fibre Channel heterogeneous SAN support**

| Storage system | Firmware | Switches | Operating systems[1] |
|---|---|---|---|
| MSA 2050/2052 | VL100R004 | H-series[2]<br>B-series<br>C-series<br>FlexFabric 5900CP | HP-UX<br>Microsoft Windows<br>Red Hat Enterprise Linux<br>SUSE Linux Enterprise Server<br>VMware ESX<br>XenServer CITRIX |
| MSA 1050 | VL100R002 | | Microsoft Windows<br>Red Hat Enterprise Linux<br>SUSE Linux Enterprise Server<br>VMware ESX |
| MSA 1040 | GL220P009 (or later) | | Microsoft Windows<br>Red Hat Enterprise Linux<br>SUSE Linux Enterprise Server<br>VMware ESX |

*Table Continued*



https://t.me/learningnets

| Storage system | Firmware | Switches | Operating systems[1] |
|---|---|---|---|
| MSA 2040/2042 | GL220P009 (or later) | | Apple OS X<br>HP-UX<br>Microsoft Windows<br>OpenVMS<br>Red Hat Enterprise Linux<br>SUSE Linux Enterprise Server<br>VMware ESX<br>XenServer CITRIX |
| P2000 G3<br>• FC<br>• FC/iSCSI | TS252P005 (or later) | | Apple OS X<br>HP-UX<br>Microsoft Windows<br>OpenVMS<br>Red Hat Enterprise Linux<br>SUSE Linux Enterprise Server<br>VMware ESX<br>XenServer CITRIX |

1 Not all operating systems listed are supported on all switch series or models.

2 The H-series switches have limited operating system support, see table H-series switches servers, operating systems, and storage system support.

For the latest information on firmware versions and MSA storage system support, see the SPOCK website at **http://www.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access.



https://t.me/learningnets

## Configuration rules

### Table 97: MSA Fibre Channel Storage System rules

| Rule number | Description |
|---|---|
| 1 | The MSA 2050/2052/1050, MSA 2040/2042/1040, and P2000 G3 FC are supported in all SAN fabric topology configurations described in this guide. |
| 2 | Unless otherwise specified, the MSA 2050/2052/1050, MSA 2040/2042/1040, and P2000 G3 FC can be configured in a SAN using the switch models listed in this guide, see:<br>• H-series switches table<br>• B-series Fibre Channel switches table through B-series legacy Fibre Channel switches and routers table |
| 3 | Servers that share access to the same storage LUN require special application software (such as cluster software) to preserve data. |
| 4 | Linux with Emulex or QLogic — See the *Installation and Reference Guide Device Mapper Multipath Enablement Kit for StorageWorks Disk Arrays*.<br>For all other supported operating systems, set the following parameters:<br>• **Brocade:** `MPIO Mode Enable = off, Path Timeout = 60`<br>• **Emulex:** `LinkTimeOut = 60, NodeTimeOut = 60`<br>• **QLogic:** `Port Down Retry Count = 60, Link Down Timeout = 60` |
| 5 | Use volume mapping to enable or disable volume access to specific hosts. |
| 6 | When connecting to a SAN, Hewlett Packard Enterprise recommends setting the topology to "Point to Point" on the MSA 2050/2052/1050, MSA 2040/2042/1040, and P2000 G3 FC. |

## Zoning

These zoning rules apply to a heterogeneous SAN with MSA 2050/2052/1050, MSA 2040/2042/1040, and P2000 G3 storage systems:

- Zoning is required for all operating systems that access the array, see **Heterogeneous SAN support** on page 228.
- 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, different storage system types** on page 224.



https://t.me/learningnets

# Maximums

**Table 98: MSA Fibre Channel connection maximum configurations**

| Storage systems | Operating systems | Drives | Hosts | LUNs | LUN size | Snapshots and clones[1] |
|---|---|---|---|---|---|---|
| MSA 2050/2052 SAN | Microsoft Windows<br>Red Hat Enterprise Linux<br>SUSE Linux Enterprise Server<br>VMware ESXi<br>HP-UX<br>Citrix XenServer | 96 3.5-in LFF<br>192 2.5-in SFF | 1024 | 1024 | 128TiB | 64 Standard (maximum 512 snapshots) |
| MSA 1050 FC | Microsoft Windows<br>Red Hat Enterprise Linux<br>SUSE Linux Enterprise Server<br>VMware ESXi | 48 3.5-in LFF<br>99 2.5-in SFF | 1024**# GUID-D7147 C7F-2 016-09 01-03 BB-00 00000 00E0F/ fn2** | 1024 | 128TiB | 64 standard (maximum 512 snapshots) |
| MSA 1040 FC | Microsoft Windows<br>Red Hat Enterprise Linux<br>SUSE Linux Enterprise Server<br>VMware ESXi | 48 3.5-in LFF<br>99 2.5-in SFF | 1024**# GUID-D7147 C7F-2 016-09 01-03 BB-00 00000 00E0F/ fn2** | 512 | Up to 64 TB depending on vdisk configuratio n | 64 standard (maximum 512 snapshots) |

*Table Continued*



https://t.me/learningnets

| Storage systems | Operating systems | Drives | Hosts | LUNs | LUN size | Snapshots and clones[1] |
|---|---|---|---|---|---|---|
| MSA 2040 SAN | Apple OS X<br>HP-UX 11.31<br>Microsoft Windows<br>Red Hat Enterprise Linux<br>SUSE Linux Enterprise Server<br>VMware ESXi | 96 3.5-in LFF<br>199 2.5-in SFF | 1024**#GUID-D7147C7F-2016-0901-03BB-000000000E0F/fn2** | 512 | Up to 64 TB depending on vdisk configuration | 64 standard (maximum 512 snapshots) |
| P2000 G3 FC<br>P2000 G3 FC/iSCSI | Apple OS X<br>HP-UX 11.23, 11.31<br>OpenVMS<br>Citrix XenServer<br>Microsoft Windows Server<br>Red Hat Enterprise Linux<br>SUSE Linux Enterprise Server<br>VMware ESX, ESXi<br>Oracle Solaris | 96 3.5-in LFF<br>149 2.5-in SFF | 64 | 512 | Up to 64 TB depending on vdisk configuration | 64 standard (maximum 512 snapshots, clones or remote snaps) |

[1] Snapshots and clones require an additional license.

Contact a Hewlett Packard Enterprise storage representative for a configuration review to determine support for configurations that exceed the stated limits.

## Management software support

The MSA 2050/2052/1050, MSA 2040/2042/1040, and P2000 G3 support target-based management interfaces, including Telnet (CLI), SSH (CLI), FTP, SFTP, SMI-S, SNMP, HTTP and HTTPS web-based interface. The web-based interface is supported with Microsoft Internet Explorer, Chrome and Mozilla Firefox.



https://t.me/learningnets

# HPE StoreVirtual 3200 and 4000 Storage System Rules

This chapter describes specific rules for management groups with at least two of the following storage systems:

- HPE StoreVirtual 3200 FC Storage
- HPE StoreVirtual 4730 FC Storage
- HPE StoreVirtual 4330 FC Storage

This chapter describes the following topics:


- **Fibre Channel on HPE StoreVirtual 3200 Storage** on page 234
- **Fibre Channel on HPE StoreVirtual 4000 Storage** on page 234
- **Campus cluster support** on page 235
- **Heterogeneous SAN support** on page 235
- **Configuration rules** on page 236
- **Configuration parameters** on page 237
- **Data migration** on page 237
- **Zoning** on page 237
- **Boot from SAN support** on page 237


## Fibre Channel on HPE StoreVirtual 3200 Storage

Fibre Channel connectivity on StoreVirtual 3200 Storage is used for host traffic only. For all management communication, an Ethernet network is required.

StoreVirtual 4000 storage clusters cannot include StoreVirtual 3200 storage pools.

## Fibre Channel on HPE StoreVirtual 4000 Storage

Before implementation, contact a Hewlett Packard Enterprise storage representative for information about support for specific configurations, including the following elements:

- LeftHand operating system version
- Operating systems and Fibre Channel HBA support
- Multipath configuration

Fibre Channel connectivity on StoreVirtual 4000 Storage is used for host traffic only. For all management and cluster communication among storage nodes and iSCSI connectivity to other initiators, a 10 GbE network is required.

To enable Fibre Channel for a management group, it must meet the following requirements:



https://t.me/learningnets

- There are at least two Fibre Channel enabled storage nodes (that is storage nodes with Fibre Channel interfaces and on which Fibre Channel is enabled) in the management group
- All storage nodes in the management group are running LeftHand operating system version 10.5 or higher

To extend Fibre Channel connectivity to existing StoreVirtual 4000 Storage clusters, Fibre Channel enabled storage nodes (see models listed at the beginning of this section) can be added to these existing clusters. However, in cases where this is not recommended (different drive type, different capacity), Fibre Channel enabled storage nodes should join the same management group, but should form a new cluster. These storage nodes with Fibre Channel connectivity can serve as gateway to volumes in other clusters in the same management group.

StoreVirtual 3200 storage pools cannot include StoreVirtual 4000 storage clusters.

# Campus cluster support

Fibre Channel on StoreVirtual 4000 Storage can be deployed in clusters that are stretched across two racks, data rooms or data centers. These installations require:

- Two stretched Fibre Channel fabrics
- A stretched 10 GbE Ethernet network

---

**NOTE:**

- 500 MB/s (4,000 Mb/s) of bandwidth per storage node pair needs to be allocated on the 10 GbE network between the two locations. Network latency among storage nodes cannot exceed 1 ms.
- Campus cluster support does not apply to StoreVirtual 3200 Storage.

---

The two Fibre Channel fabrics between the two sites can be stretched using any native and transparent fabric extension technology, such as long-range optics and DWDM. SAN extension using other intermediate protocols, like IP, is not supported in campus cluster configurations.

To set up the campus cluster functionality, create a multi-site cluster by assigning storage nodes to sites. For more information on how to configure multi-site clusters and how to work with sites, see the StoreVirtual documentation. Note that server connections, which are assigned to sites in the management group and are configured for volume access via Fibre Channel, will not honor the site preference. All multi-site cluster best practices apply.

All Fibre Channel ports of StoreVirtual 4000 Storage must be visible to servers connecting via Fibre Channel to ensure connectivity in case of port, node, or site failures. For more information, see **Zoning** on page 237.

If Fibre Channel connectivity is used to access volumes on another cluster in the same management group, the other cluster must have an identical site configuration.

# Heterogeneous SAN support

StoreVirtual 3200 Storage supports shared access with the operating system combinations listed in the configuration sets on SPOCK. StoreVirtual 4000 Storage supports shared access with the operating system combinations listed in the configuration sets on SPOCK. Currently, any combinations with Microsoft Windows Server 2008 and later, and VMware vSphere 4 and later, are supported.



https://t.me/learningnets

For more details on Fibre Channel HBA and switch support, see the HPE StoreVirtual 3200 and HPE StoreVirtual 4000 Storage Configuration Sets available on SPOCK at **http://www.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access.

# Configuration rules

### Table 99: Heterogeneous SAN storage with Fibre Channel on HPE StoreVirtual 3200 and HPE StoreVirtual 4000 Storage rules

| Rule number | Description |
|---|---|
| 1 | Fibre Channel on StoreVirtual 3200 or StoreVirtual 4000 Storage is supported in all Fibre Channel SAN fabric topology configurations described in this guide. This requires unique domain IDs in the entire Fibre Channel SAN to which the StoreVirtual 3200 or StoreVirtual 4000 storage system is connected, including multiple independent fabrics. Overlapping domain IDs (even in non-merged/independent fabrics) are not supported. |
| 2 | Fibre Channel on StoreVirtual 3200 or StoreVirtual 4000 Storage can be configured in a SAN using the Fibre Channel switch models listed in the StoreVirtual 3200 or StoreVirtual 4000 Storage sets available on SPOCK at **http://www.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access. |
| 3 | All Fibre Channel ports in a StoreVirtual 3200 or StoreVirtual 4000 Storage management group must be zoned to an HBA port of the server. The resulting zone has the HBA port of the server and all Fibre Channel port of the StoreVirtual management group on the fabric as members.<br><br>**NOTE:** Each physical Fibre Channel port on the SV3200 also has an NPIV port and both have a WWPN. When zoning by WWPN both the physical and NPIV ports need to be included. The Fibre Channel switch provides a complete list of the WWPNs. A SV3200 Fibre Channel will have a total of 8 WWPNs, 4 WWPNs on each storage controller. |
| 4 | Overlapping storage port zones are supported. |
| 5 | For information about shared access and heterogeneous platform zoning requirements, see **Heterogeneous SAN storage system coexistence** on page 224. |
| 6 | A single server's HBAs are supported for common access to other storage systems families as long as the HBAs are from the same vendor, see **Zoning** on page 237. |
| 7 | To enable or disable volume or snapshot access to specific hosts, use the **Assign/unassign volumes and snapshots** functionality in the Centralized Management Console. For more information, see the StoreVirtual 3200 Management Console Online Help or user guide or the StoreVirtual 4000 documentation. |
| 8 | Servers that share access to the same storage LUN require special application software (such as cluster software) to preserve data integrity. |

For information about configuring StoreVirtual 3200 Storage using the StoreVirtual Management Console, see the *HPE StoreVirtual 3200 Management Console User Guide* at **http://www.hpe.com/info/StoreVirtual3200Manuals**.



https://t.me/learningnets

For information about configuring StoreVirtual 4000 Storage using the Centralized Management Console, see *HPE StoreVirtual 4000 Storage User Guide*.

# Configuration parameters

For Fibre Channel on StoreVirtual 3200 or StoreVirtual 4000 Storage configuration settings, see the *HPE StoreVirtual 3200 Storage User Guide* or *HPE StoreVirtual 4000 Storage User Guide* respectively, and configuration sets on SPOCK.

# Data migration

Fibre Channel on StoreVirtual 4000 Storage does not currently support data migration using the MPX200. However, data from supported Fibre Channel source arrays can be migrated to volumes on StoreVirtual 4000 Storage via iSCSI.

This data migration is done using the Multifunction Router data migration feature. It provides for block (LUN) level data movement between source and destination storage systems.

For the latest data migration storage system, operating system, and version support, see the SPOCK website at **http://www.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access.

# Zoning

These zoning rules apply to a heterogeneous SAN with Fibre Channel on StoreVirtual 3200 or StoreVirtual 4000 Storage:

- Zoning by HBA is the HPE recommended zoning method for Fibre Channel on StoreVirtual 3200 and StoreVirtual 4000 Storage, see **Configuration rules** on page 236.
- Zoning is required for all operating systems that access StoreVirtual 3200 and StoreVirtual 4000 Storage using the Fibre Channel connection type, see **Heterogeneous server rules** on page 179.
- Fibre Channel on StoreVirtual 3200 and StoreVirtual 4000 Storage can be in multiple operating system zones.

# Boot from SAN support

For current boot from SAN support, see the SPOCK website at **http://www.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access to SPOCK.



https://t.me/learningnets

# P6000 storage system rules

This chapter describes specific rules for the following storage systems:

- P6300/P6500 EVA
- P6350/P6550 EVA

**IMPORTANT:** P6000 storage was formerly called the HPE Virtual Array product family. General references to P6000 can also refer to earlier versions of EVA products.

This chapter describes the following topics:


- **P6000/EVA storage** on page 238
- **P6000/EVA data migration** on page 243
- **HPE P6000 Continuous Access SAN integration** on page 247
- **Zoning** on page 248
- **P6000/EVA SAN boot support** on page 249
- **Storage management server integration** on page 249
- **Cabling** on page 250


## P6000/EVA storage

Before implementation, contact a Hewlett Packard Enterprise storage representative for information about support for specific configurations, including the following elements:

- Storage system firmware
- HBA firmware and driver versions



https://t.me/learningnets

## Heterogeneous SAN support

### Table 100: P6000/EVA heterogeneous SAN support

| Storage systems | Firmware[1] | Switches | Operating systems[2] |
|---|---|---|---|
| | | B-series<br>C-series<br>H-series<br>FlexFabric 5900CP | Apple Mac<br>Citrix Xen Server<br>HP-UX<br>IBM AIX<br>Microsoft Windows<br>OpenVMS<br>Red Hat Linux<br>Oracle Solaris<br>SUSE Linux<br>Tru64 UNIX<br>VMware ESX<br>**NOTE:** For the latest information on storage system support, see the SPOCK website at **http://www.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access. |
| P6300/P6500 EVA | XCS 10**x** | | |
| P6350/P6550 EVA | XCS 11**x** | | |

1 Contact a Hewlett Packard Enterprise storage representative for switch model and operating system support when using XCS 10**x**, 09**x** or 6**x**.

2 For the supported operating system versions, see **Heterogeneous server rules** on page 179. Not all operating systems listed are supported on all switch series or models.

## Configuration rules

### Table 101: 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:<br>• H-series switches table<br>• B-series Fibre Channel switches table through B-series legacy Fibre Channel switches and routers table |

*Table Continued*



https://t.me/learningnets

| Rule number | Description |
|---|---|
| 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 P6000 Continuous Access, the zoning limit is 256 HBAs. |
| 5 | For shared access and heterogeneous platform zoning requirements, see **Heterogeneous SAN storage system coexistence** on page 224. |
| 6 | For EVA4400 storage systems with the EVA4400 Embedded Switch Module, 8 Gb Brocade connectivity rules, see B-series Fibre Channel switch fabric rules table. Hewlett Packard Enterprise recommends setting the switch port speed to autonegotiate. |
| 7 | • P6300/P6500 EVA requires XCS 10001000 minimum.<br>• P6350/P6550 EVA requires XCS 11001000 minimum.<br>• All Fibre Channel host port connections to P6000/EVA controllers must be set to autonegotiate, 8 Gb/s (P63**xx**/P65**xx** EVA only), 4 Gb/s, or 2 Gb/s. Hewlett Packard Enterprise recommends setting the port speed to autonegotiate.<br>• 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.<br>**NOTE:** 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). |
| | your product. |

*Table Continued*



https://t.me/learningnets

| Rule number | Description |
|---|---|
| 9 | P6000/EVA storage systems 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.<br>• HP-UX, Linux, and Microsoft Windows are supported for direct connect.<br>• 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).<br>• Direct connect to blade servers is not supported.<br>Use the EVA OCP or WOCP to specify the proper topology setting for the controller host ports (P63**xx**/P65**xx**).<br>• **For fabric** — Set the applicable HSV controller host ports to `Fabric`.<br>• **For direct connect** — Set the applicable HSV controller host ports to `Direct Connect`.<br>For information about the OCP or WOCP, see the *HPE P6000/EVA Storage System Documentation* for your model. |
| 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:<br>• From the OCP or WOCP, set the port to `direct connect` mode.<br>• Insert a loopback connector when a P6000/EVA host port is not connected to a switch or an HBA (for direct connect). |
| 11 | Supports connection of single HBA servers. Contact a Hewlett Packard Enterprise storage representative for the white paper Connecting Single HBA Servers to the Enterprise Virtual Array without Multipathing Software.<br>Servers without native or layered multipathing software are not supported for 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. |

*Table Continued*



https://t.me/learningnets

| Rule number | Description |
|---|---|
| 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** on page 249.

## Configuration parameters

This section describes general P6000/EVA configuration parameters. Specific solutions, such as high-availability clusters, or applications, such as P6000 Continuous Access (see **HPE P6000 Continuous Access SAN integration** on page 247), can define additional configuration parameters or requirements. Solution requirements must be observed as specified by the solution configuration documentation.

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 *HPE 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 | = 0128 LUN presentations |
| LUNs #161 through #192 are presented to 4-node cluster | = 0008 LUN presentations |
| LUNs #193 through #200 are presented to single host | |



https://t.me/learningnets

When all LUNs are presented to all hosts, the number of LUNs multiplied by the number of hosts must not exceed 8,192.

There is no limit on the maximum number of EVA storage systems in a SAN.

**EVA single-server maximum configurations**

| Storage systems | Operating systems[1], [2] | EVA storage systems per HBA | LUNs per HBA target[3], [4] |
|---|---|---|---|
| P6300/P6500 EVA<br>P6350/P6550 EVA | Tru64 UNIX[5] | 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 |
| | Oracle Solaris | 16 | 255 |
| | VMware ESX 5.0 and above | 32 | 256 (per server) |
| | Windows 2008<br>Windows 2012 | 16 | 255 (per server) |
| | Xen | 16 | 255 |

1 For the supported operating system versions, see **Heterogeneous server rules** on page 179.

2 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).

3 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.

4 The maximum number of configurable LUNs per HBA target may be less based on performance requirements.

5 The P63**xx**/P65**xx** EVA is not supported with Tru64 UNIX.

# P6000/EVA data migration

The P6000/EVA family of Fibre Channel storage systems supports data migration using the `StorageWorks MPX200 Multifunction Router data migration 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 one other license-enabled function — either data



https://t.me/learningnets

migration or FCIP for standalone or concurrent operation. This section describes data migration usage and support.

- For information about FCoE, see **FCoE SAN fabrics** on page 25.
- For information about iSCSI configurations, see **MPX200 Multifunction Router with iSCSI for P6000/EVA storage** on page 359, **MPX200 Multifunction Router with iSCSI for XP storage** on page 370, and **MPX200 Multifunction Router with iSCSI for 3PAR StoreServ Storage** on page 366.
- For information about FCIP, see **MPX200 Multifunction Router with FCIP** on page 302.

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:

- **HPE Storage Works MPX200 1TB Full Chassis Data Migration License** — Includes the license to migrate 1 TB of data using an MPX200 chassis.
- **HPE Storage Works MPX200 5TB Full Chassis Data Migration License** — Includes the license to migrate 5 TB of data using an MPX200 chassis.
- **HPE 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. This license cannot be used when performing an iSCSI to iSCSI data migration. For iSCSI to iSCSI data migration, you must use 1TB or 5TB capacity licenses.

All arrays listed are supported for offline and online data migration, except where noted. For information about configuring the MPX200 for data migration, see the *HPE MPX200 Multifunction Router Data Migration Solution Guide*.



https://t.me/learningnets

### Table 102: P6000/EVA data migration source-destination storage systems

| Source storage systems | P6000/EVA destination storage systems |
|---|---|
| • All MSA (Fibre Channel) and P6000/EVA models<br>• P9500/XP24000/20000, XP12000/10000<br>• SVSP<br>• 3PAR S-Class | • P6300/P6500 EVA<br>• P6350/P6550 |
| Third-party array models:<br>• Dell EqualLogic family (iSCSI), Compellent Series 30 and 40 Controllers<br>• EMC CLARiiON AX series, CX Series, Symmetrix DMX Series, Symmetrix VMAX SE (offline only), VNX5500 (offline only), VNX5300<br>• Fujitsu ETERNUS DX400, DX440 S2 (offline only), DX8400<br>• Hitachi Data Systems V series (offline only), AMS Family, WMS, DF-500, USP (offline only), VSP families (offline only), TagmaStore Network StorageController model NSC55<br>• IBM DS3000 series, DS4000 series (offline only), DS5000 series, SVC, DS6000 series (offline only), DS8000 series, XIV Storage System family (offline only), nlStorwize V7000 Unified disk system (offline only)<br>• NEC D-Series SAN Storage arrays (offline only)<br>• NetApp FAS270, FAS250, FAS2000 Series, FAS3100 Series, FAS6000 Series (Fibre Channel and iSCSI)<br>• Xiotech Emprise 5000, Mag3D 4000 | |



https://t.me/learningnets

**Table 103: Online data migration operating system support**

| MPX200 online data migration support[1] | Online data migration destination storage system and firmware (minimum) |
|---|---|
| • Citrix XenServer 6.0<br>• HP-UX 11iv3, 11iv2, Clusters (Service Guard)<br>• IBM AIX 6.1, 5.3[2]<br>• Microsoft Windows 2012 R2, 2008 R2 and 2003, Failover Clustering and MSCS<br>• Red Hat 6, 5 U4 , 4 U8, U7, U6, U3, Clusters (RH 6)<br>• SUSE 11, 10 SP3, SP1, Clusters (SUSE 11 U1)<br>• Oracle Solaris 10, Clusters 3.3<br>• VMware ESXi 5.1, 5.0, 4.1, 4.0, 3.5, Clusters | • P2000 G3 FC (TS251P002-04)<br>• P4000 (9.0)<br>• P6350/P6550 (11001000)<br>• P6500/P6300 (10001000)<br>See **P6000/EVA data migration** on page 243. |
| | • P9500 (70-00-50-00)<br>• XP24000/20000 (60-06-10-00)<br>• XP12000/10000 (50-09-83-00)<br>See **P9000/XP data migration** on page 257. |
| | • 3PAR StoreServ 7200, 3.1.3 (MU1)<br>• 3PAR StoreServ 7400, 3.1.3 (MU1)<br>• 3PAR StoreServ 7450, 3.1.3 (MU1)<br>• 3PAR StoreServ 7200c, 3.2.1<br>• 3PAR StoreServ 7400c, 3.2.1<br>• 3PAR StoreServ 7440c, 3.2.1<br>• 3PAR StoreServ 7450c, 3.2.1<br>• 3PAR StoreServ 10400, 3.1.3 (MU1)<br>• 3PAR StoreServ 10800, 3.1.3 (MU1)<br>• 3PAR T-Class, 2.3.1 (MU42)<br>See **HPE 3PAR data migration** on page 265. |

1 For operating system updates, the minimum supported version is specified.

2 IBM AIX 5.3 not supported with 3PAR OS 3.1.2 or later.

For current data migration storage system support and up-to-date operating system version support, see the SPOCK website at **http://www.hpe.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



https://t.me/learningnets

operating system/version and data migration storage system support, see the SPOCK website at **http://www.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access.

# HPE P6000 Continuous Access SAN integration

### Table 104: HPE P6000 Continuous Access heterogeneous SAN configuration rules

| Rule number | Description |
|---|---|
| 1 | Hewlett Packard Enterprise strongly recommends that all P6000 Continuous Access deployments implement level 4 NSPOF SANs using two or more separate fabrics, see **Data availability** on page 40.<br><br>P6000 Continuous Access supports a subset of operating systems listed in this guide, which limits the types of servers that can reside in the P6000 Continuous Access management zone. For more information, see the *HPE P6000 Enterprise Virtual Array Compatibility Reference Guide* at www.hpe.com/info/P6000-ContinuousAccess. |
| 2 | The maximum number of switches and switch hops supported in a fabric with P6000 Continuous Access 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 more information, see the *HPE P6000 Continuous Access Implementation Guide* at www.hpe.com/info/P6000-ContinuousAccess and the switch and fabric rules. For switch and fabric rules, see:<br><br>• **B-series switches and fabric rules** on page 105<br>• **C-series switches and fabric rules** on page 137<br>• **H-series switches and fabric rules** on page 151<br><br>For P6000 Continuous Access with XCS version 09534000 (or later), HPE supports two storage system data replication protocols: the HPE-FC protocol and the HPE SCSI-FCP protocol. You must configure the Fibre Channel switches based on which data replication protocol you use.<br><br>For H-series switches, you must use the SCSI FC Compliant Data Replication Protocol. |
| 3 | Shared usage of P6000 Continuous Access-configured storage systems by non-Hewlett Packard Enterprise P6000 Continuous Access-configured servers (for example, a single HBA or an operating system without multipathing support) or non-Hewlett Packard Enterprise P6000 Continuous Access-supported operating systems is not supported. |
| 4 | For information about supported versions of clustering software and Secure Path, contact a Hewlett Packard Enterprise storage representative. |
| 5 | Each 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 P6000 Continuous Access solutions can exist in the same SAN, provided that no solution exceeds the 16-array limit, imposed by zoning. |

*Table Continued*



https://t.me/learningnets

| Rule number | Description |
|---|---|
| 6 | 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 *HPE P6000 Enterprise Virtual Array Compatibility Reference*). 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. |
| 7 | The P6000 Continuous Access link supports mixed heterogeneous SAN, P6000 Continuous Access, and OpenVMS host-based shadowing traffic. |
| 8 | 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 P6000 Command View can be used for initial setup of P6000 Continuous Access storage. Management of the operational P6000 Continuous Access environment is done through the HPE Replication Solutions Manager and other products installed on the Storage Management Appliances.<br>For more information, see the *HPE P6000 Continuous Access Implementation Guide*, available at www.hpe.com/info/P6000-ContinuousAccess. |
| 9 | 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 **Cabling** on page 250 and the *HPE P6000 Continuous Access Implementation Guide*. |

# 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** on page 179.
- 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, different storage system types** on page 224.

### Table 105: 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** on page 249. |
| 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:



https://t.me/learningnets

- 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.
- Hewlett Packard Enterprise recommends using a separate tape-to-HBA connection for servers when backups require more than four DLT8000 tape drives or two Ultrium (LTO) tape drives.

For more information about tape storage support, contact your Hewlett Packard Enterprise representative.

# P6000/EVA SAN boot support

For current storage system support, see the Boot from SAN website at **http://www.hpe.com/info/BootFromSan** and the SPOCK website at **http://www.hpe.com/storage/spock**. For the SPOCK website, you must sign up for an HP Passport to enable access.

For 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 Blade 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 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. After you enable security, you cannot disable it.

### Table 106: 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. P6000 Command View 10.1 (or later) is required for P6350/P6550 EVA. |
| 3 | Command View EVA 5.0 (or later) supports management of multiple fabrics from a single instance. |
| 4 | Command View EVA 5.0 (or later) is required for the EVA iSCSI Connectivity Option. |

*Table Continued*



https://t.me/learningnets

| Rule number | Description |
|---|---|
| 5 | If a storage management server resides in the fabric, it is recommended that it be configured in a separate zone from all operating systems.<br>Create a storage management zone for the storage management server and the elements it will monitor and manage. |
| 6 | It is not necessary to include the switch WWNs or server HBA WWNs in the storage management server zone.<br>Management communication to these devices from the storage management server is done out-of-band or outside the fabric via TCP/IP. |
| 7 | 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. |
| 8 | Multiple storage management servers per fabric are allowed. A single P6000/EVA can be zoned with more then one instance of P6000 Command View, independent of where P6000 Command View is loaded (SMA, GPS, management station, or Storage Server). |
| 9 | Each storage system can be managed by only one active storage management server.<br>Any standby storage management server can be powered on, but the P6000 Command View or P6000 Continuous Access user interface must not control the storage system.<br>For more information, see the *HPE P6000 Continuous Access Implementation Guide*, available at **http://www.hpe.com/info/P6000-ContinuousAccess**. |
| 10 | For fabrics with more than 1,024 HBAs, the HSV controller must be zoned to limit access to a maximum of 1,024 HBAs.<br>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.

## 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 67: Single PCI slot with dual-channel HBA and one switch** on page 251 or **Figure 68: Single PCI slot with dual-channel HBA and two switches** on page 251. 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 67: Single PCI slot with dual-channel HBA and one switch** on page 251 shows two HBA paths connected to the same Fibre Channel switch.



https://t.me/learningnets

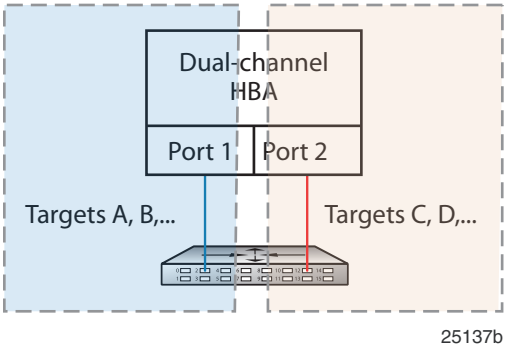



**Figure 67: Single PCI slot with dual-channel HBA and one switch**

The configuration shown in **Figure 68: Single PCI slot with dual-channel HBA and two switches** on page 251 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.

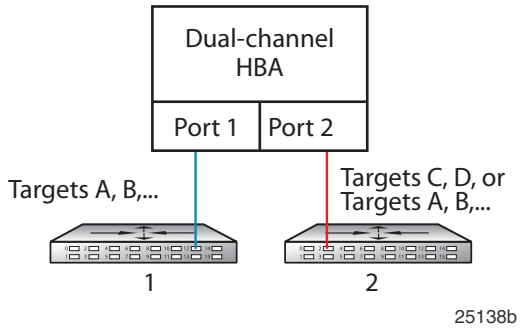



**Figure 68: Single PCI slot with dual-channel HBA and two switches**

**Figure 69: Two dual-channel HBAs (NSPOF)** on page 251 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** on page 40.

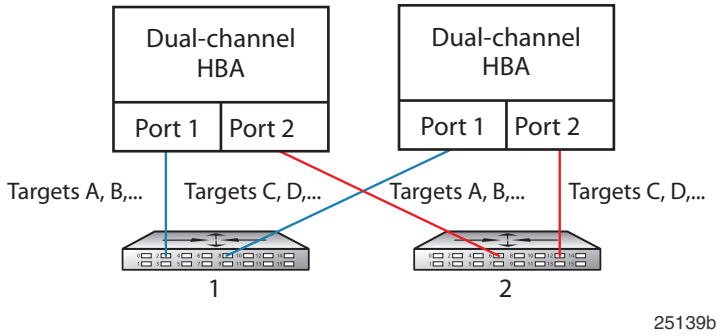



**Figure 69: Two dual-channel HBAs (NSPOF)**



https://t.me/learningnets

# P9000/XP storage system rules

This chapter describes specific rules for the following storage systems:

- XP7
- P9500
- XP24000
- XP20000

This chapter describes the following topics:


- **P9000/XP storage systems** on page 252
- **P9000/XP SAN boot support** on page 256
- **LUN Configuration and Security Manager XP support** on page 257
- **P9000/XP data migration** on page 257


## P9000/XP storage systems

Before implementation, contact a Hewlett Packard Enterprise storage representative for information about support for specific configurations, including the following elements:

- Storage system firmware
- HBA firmware and driver versions
- Multipath software



https://t.me/learningnets

## Heterogeneous SAN support

### Table 107: P9000/XP heterogeneous SAN support

| Storage systems | Firmware version[1] | Switches[2][3] | Operating systems |
|---|---|---|---|
| XP7 | 80x | B-series<br>C-series | HP-UX<br>IBM AIX<br>Microsoft Windows<br>Red Hat Linux<br>Oracle Solaris<br>SUSE Linux<br>VMware ESX<br>NonStop |
| P9500 | 70**x** | B-series[4]<br>C-series<br>H-series<br>FlexFabric 5900CP | HP-UX<br>IBM AIX<br>Microsoft Windows<br>OpenVMS<br>Red Hat Linux<br>Oracle Solaris<br>SUSE Linux<br>VMware ESX<br>NonStop<br>Citrix 5.6 |
| XP20000<br>XP24000 | 60**x** | B-series<br>C-series<br>H-series[5] | HP-UX<br>IBM AIX<br>Microsoft Windows<br>OpenVMS<br>Red Hat Linux<br>Oracle Solaris<br>SUSE Linux<br>Tru64 UNIX<br>VMware ESX |
| XP12000<br>XP10000 | 50**x** | | |

1 Contact a Hewlett Packard Enterprise storage representative for the latest firmware version support.

2 XP7 and P9500 storage systems are not supported with 2 Gb/s switches.

3 For the supported operating system versions, see **Heterogeneous server rules** on page 179. Not all operating systems listed are supported on all switch series or models.

4 B-series switches set to mode 3 are not supported with the following firmware versions: 70-06-09-00/00, 70-06-11-00/00, 70-06-13-00/00 and 70-06-15-00/00.



https://t.me/learningnets

[5] The H-series switches have limited operating system support, see H-series switches servers, operating systems, and storage system support table.

## Configuration rules

### Table 108: P9000/XP storage system rules

| Rule number | Description |
|---|---|
| 1 | P9000/XP storage systems are supported in all SAN fabric topology configurations described in this guide. |
| 2 | Unless otherwise specified, P9000/XP storage systems can be configured in a SAN using the switch models listed in this guide. See the following tables:<br>• H-series switches table<br>• B-series Fibre Channel switches table through B-series legacy Fibre Channel switches and routers table<br>XP7 and P9500 storage systems are not supported with 2 Gb/s switches.<br>XP7 storage systems are not supported with H-series switches. |
| 3 | Storage ports can be accessed from heterogeneous operating system types and multiple clusters for HA and non-HA configurations.<br>Overlapping zones are supported with disk and tape. |
| 4 | LUN Configuration and Security Manager XP 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 storage systems.<br>See **Heterogeneous SAN storage system coexistence** on page 224. |
| 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:<br>• Switch/Fabric (F_Port), Fabric = On, Connection = Point to Point<br>• Switch/Fabric (FL_Port), Fabric = On, Connection = FC-AL<br>• Direct connect (no Fibre Channel switch) (NL_Port), Fabric = Off, Connection = FC-AL<br>See *HPE P9000 (XP) Continuous Access configuration support* section for settings when using P9000 (XP) Continuous Access. |
| 7 | Multiple operating systems and multiple clusters can be supported on the same fabric with appropriate zoning. |
| 8 | Host zones must contain homogeneous operating system types. |

*Table Continued*



https://t.me/learningnets

| Rule number | Description |
|---|---|
| 9 | Overlapping storage port zones are supported if more than one operating system must share an array port. |
| 10 | 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. |
| 11 | 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 storage systems:

- Zoning is required for all operating systems that access P9000/XP storage systems. See **Heterogeneous server rules** on page 179.
- P9000/XP 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, different storage system types** on page 224.

## Tape storage

The following rules apply to tape storage support in a SAN with P9000/XP 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.
- Hewlett Packard Enterprise recommends using a separate tape-to-HBA connection for servers when backups require more than four DLT8000 tape drives or two Ultrium (LTO) tape drives.

For more information about tape storage support, contact your Hewlett Packard Enterprise representative.

**Figure 70: P9000/XP storage systems with tape storage in a shared fabric** on page 256 shows a configuration with XP storage systems in the same fabric as tape storage.



https://t.me/learningnets

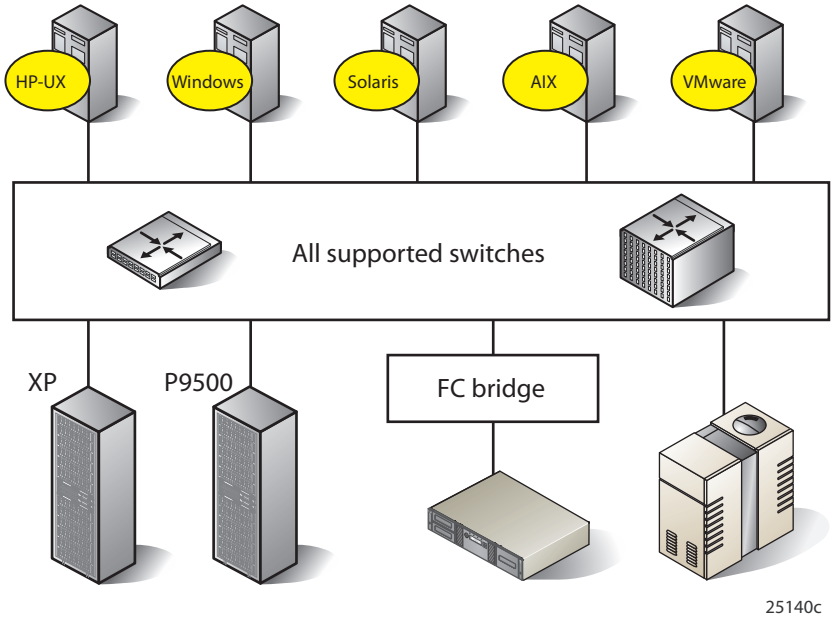



**Figure 70: P9000/XP storage systems with tape storage in a shared fabric**

# P9000/XP SAN boot support

P9000/XP LUNs can be booted from the SAN using B-series, C-series, and H-series switches.

SAN boot through the B-series MP Router Blade 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 a Hewlett Packard Enterprise storage representative.

### Table 109: P9000/XP SAN boot support by operating system

| Switches | Storage systems | HP-UX | Linux | Windows | Tru64UNIX | OpenVMS | OracleSolaris | IBMAIX | IRIX | Xen |
|---|---|---|---|---|---|---|---|---|---|---|
| B-series<br>C-series<br>H-series | XP7 | . | . | . | - | - | - | . | - | - |
| | P9500 | . | . | . | - | . | . | . | - | - |
| | XP24000 | . | . | . | . | . | . | . | - | . |
| | XP20000 | . | . | . | . | . | . | . | - | . |
| | XP12000 | . | . | . | . | . | . | . | - | . |
| | XP10000 | . | . | . | . | . | . | . | - | . |



https://t.me/learningnets

**Legend:** . = supported; — = not supported

# 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 StorageWorks MPX200 Multifunction Router data migration 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 one other license-enabled function — either data migration or FCIP for standalone or concurrent operation. This section describes data migration usage and support.

- For information about FCoE, see **FCoE SAN fabrics** on page 25.
- For information about iSCSI configurations, see **MPX200 Multifunction Router with iSCSI for P6000/EVA storage** on page 359, **MPX200 Multifunction Router with iSCSI for XP storage** on page 370, and **MPX200 Multifunction Router with iSCSI for 3PAR StoreServ Storage** on page 366.
- For information about FCIP, see **MPX200 Multifunction Router with FCIP** on page 302.

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:

- **HPE Storage Works MPX200 1TB Full Chassis Data Migration License** — Includes the license to migrate 1 TB of data using an MPX200 chassis.
- **HPE Storage Works MPX200 5TB Full Chassis Data Migration License** — Includes the license to migrate 5 TB of data using an MPX200 chassis.
- **HPE 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. This license cannot be used when performing an iSCSI to iSCSI data migration. For iSCSI to iSCSI data migration, you must use 1TB or 5TB capacity licenses.



https://t.me/learningnets

All arrays listed are supported for offline and online data migration, except where noted. For information about configuring the MPX200 for data migration, see the *HPE MPX200 Multifunction Router Data Migration Solution Guide*.

### Table 110: P9500/XP data migration source-destination storage systems

| Source storage systems | P9000/XP destination storage systems |
|---|---|
| • All MSA (Fibre Channel) and EVA models<br>• P9500/XP24000/20000, XP12000/10000<br>• SVSP<br>• 3PAR S-Class | P9500/XP24000/20000 |
| Third-party array models:<br>• Dell EqualLogic family (iSCSI), Compellent Series 30 and 40 Controllers<br>• EMC CLARiiON AX series, CX Series, Symmetrix DMX Series, Symmetrix VMAX SE (offline only), VNX5500 (offline only), VNX5300<br>• Fujitsu ETERNUS DX400, DX440 S2 (offline only), DX8400<br>• Hitachi Data Systems V series (offline only), AMS Family, WMS, DF-500, USP (offline only), VSP families (offline only), TagmaStore Network StorageController model NSC55<br>• IBM DS3000 series, DS4000 series (offline only), DS5000 series, SVC, DS6000 series (offline only), DS8000 series, XIV Storage System family (offline only), nlStorwize V7000 Unified disk system (offline only)<br>• NEC D-Series SAN Storage arrays (offline only)<br>• NetApp FAS270, FAS250, FAS2000 Series, FAS3100 Series, FAS6000 Series (Fibre Channel and iSCSI)<br>• Xiotech Emprise 5000, Mag3D 4000 | |



https://t.me/learningnets

**Table 111: Online data migration operating system support**

| MPX200 online data migration support[1] | Online data migration destination storage system and firmware (minimum) |
|---|---|
| • Citrix XenServer 6.0<br>• HP-UX 11iv3, 11iv2, Clusters (Service Guard)<br>• IBM AIX 6.1, 5.3[2]<br>• Microsoft Windows 2012 R2, 2008 R2, and 2003, Failover Clustering and MSCS<br>• Red Hat 6, 5 U4 , 4 U8, U7, U6, U3, Clusters (RH 6)<br>• SUSE 11, 10 SP3, SP1, Clusters (SUSE 11 U1)<br>• Oracle Solaris 10, Clusters 3.3<br>• VMware ESXi 5.1, 5.0 Clusters | • P2000 G3 FC (TS251P002-04)<br>• P4000 (9.0)<br>• P6350/P6550 (11001000)<br>See **P6000/EVA data migration** on page 243<br>• P9500 (70-00-50-00)<br>• XP24000/20000 (60-06-10-00)<br>• XP12000/10000 (50-09-83-00)<br>See **P9000/XP data migration** on page 257<br>• 3PAR StoreServ 7200, 3.1.3 (MU1)<br>• 3PAR StoreServ 7400, 3.1.3 (MU1)<br>• 3PAR StoreServ 7450, 3.1.3 (MU1)<br>• 3PAR StoreServ 7200c, 3.2.1<br>• 3PAR StoreServ 7400c, 3.2.1<br>• 3PAR StoreServ 7440c, 3.2.1<br>• 3PAR StoreServ 7450c, 3.2.1<br>• 3PAR StoreServ 10400, 3.1.3 (MU1)<br>• 3PAR StoreServ 10800, 3.1.3 (MU1)<br>• 3PAR T-Class, 2.3.1 (MU2)<br>See **HPE 3PAR data migration** on page 265. |

1 For operating system updates, the minimum supported version is specified.
2 IBM AIX 5.3 not supported with 3PAR OS 3.1.2 or later.

## 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.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access.



https://t.me/learningnets

# HPE 3PAR StoreServ storage rules

This chapter describes specific rules for the following storage systems:

- HPE 3PAR T400/T800 Storage
- HPE 3PAR StoreServ 10400/10800 Storage
- HPE 3PAR StoreServ 7200/7400/7450 Storage
- HPE StoreServ 8200/8400/8440/8450 AFA Storage
- HPE StoreServ 20450/20800/20850 Storage

This chapter describes the following topics:


- **HPE 3PAR StoreServ storage** on page 260
- **HPE 3PAR data migration** on page 265
- **HPE 3PAR storage management** on page 271


# HPE 3PAR StoreServ storage

Before implementation, contact a Hewlett Packard Enterprise 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

### Table 112: 3PAR heterogeneous SAN support

| Storage systems | HPE 3PAR OS software (minimum) | Switches[1] | Operating systems[2] |
|---|---|---|---|
| 8200/8400/8440/8450 AFA | 3.3.1<br>3.2.2 | B-series<br>C-series | Apple Mac OS X[45]<br>HPE Integrity Virtual Machines |
| 20450/20800/20850 | 3.2.1 | H-series | HP-UX |
| 7450 | 3.2.1 | FlexFabric 5900CP[3] | Citrix Xen |
| 7200/7400 | 3.2.1<br>3.1.3 | | IBM AIX<br>Microsoft Hyper-V<br>Microsoft Windows<br>Oracle VM Server |

*Table Continued*



https://t.me/learningnets

| Storage systems | HPE 3PAR OS software (minimum) | Switches[1] | Operating systems[2] |
|---|---|---|---|
| 10400/10800<br>T400/T800 | | | Red Hat Linux<br>Oracle Linux UEK<br>Oracle Solaris<br>SUSE Linux<br>VMware ESX |

1 For Fibre Channel switch model and firmware support, see the HPE 3PAR OS Configuration Sets available on SPOCK at **http://www.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access.

2 For the supported operating system versions, see **Heterogeneous server rules** on page 179. Not all operating systems listed are supported on all switch series or models.

3 FlexFabric 5900CP is not supported with 3PAR T400/T800 storage systems.

4 Apple Mac OS X is supported only with 10400/10800, 7200/7400, and 7450 storage systems, running a minimum 3PAR OS version of 3.2.1.

5 Apple Mac OS X is supported only with B-series and C-series switches.

## HPE 3PAR StoreServ Host Interfaces

3PAR StoreServ storage systems are available with different SAN or host interface options, including Fibre Channel, FCoE and iSCSI. The number of ports available varies by model as described in the following tables:

### Table 113: HPE 3PAR StoreServ 20000 and 8000 host ports

| Model | Max host ports[1] | 16 Gb/s FC | 10 GbE FCoE | 10 GbE iSCSI |
|---|---|---|---|---|
| 20000 | 160 | 0 to 160 | 0 to 80 | 0 to 80 |
| 8000 | 24 | 0 to 24 | 0 to 8 | 0 to 8 |

1 Fibre Channel and 10 GbE iSCSI host ports are mixable on the same 3PAR storage system

### Table 114: HPE 3PAR StoreServ 10000 and 7000 host ports

| Model | Max host ports[1] | 8 Gb/s FC | 16 Gb/s FC | 10 GbE FCoE | 10 GbE iSCSI |
|---|---|---|---|---|---|
| 10000 | 192 | 0 to 192 | 0 to 96 | 0 to 96 | 0 to 32 |
| 7000 | 24 | 0 to 24 | 0 to 8 | 0 to 8 | 0 to 8 |

1 Fibre Channel and 10 GbE iSCSI host ports are mixable on the same 3PAR V-Class storage system.



https://t.me/learningnets

### Table 115: 3PAR F-Class, T-Class host ports

| Model | Total host ports[1] | 4 Gb/s FC | 1 GbE iSCSI |
|---|---|---|---|
| T400 | 64 | 0 to 64 | 0 to 16 |
| T800 | 128 | 0 to 128 | 0 to 32 |

[1] Fibre Channel and iSCSI host ports are mixable on same 3PAR storage system.

## Configuration rules

**NOTE:** For information about 3PAR StoreServ FCoE target support, see the SPOCK 3PAR FCoE configuration sets on the SPOCK website at **http://www.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access.

### Table 116: 3PAR storage rules

| Rule number | Description |
|---|---|
| 1 | 3PAR StoreServ storage is supported in all Fibre Channel SAN fabric topology configurations described in this guide. |
| 2 | 3PAR StoreServ storage can be configured in a SAN using the Fibre Channel switch models listed in the 3PAR StoreServ configuration sets available on SPOCK at **http://www.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access.<br>For information about shared access and heterogeneous platform zoning requirements, see **Heterogeneous SAN storage system coexistence** on page 224. |
| 3 | For information about shared access and heterogeneous platform zoning requirements, see **Heterogeneous SAN storage system coexistence** on page 224. |
| 4 | A single server's HBAs and multipath software for 3PAR StoreServ storage systems are supported for common access to other storage systems families as long as the HBAs are from the same vendor, see **Zoning** on page 269. This applies to Windows 2003, Windows 2008 (incl. R2), Windows 2012, Windows 2012 R2, ESX 4.**x** and 5.**x**, HPUX 11 v2 and later, Solaris 10 and later, AIX 5.3 and later, and Linux (RHEL 5.6 and later, RHEL 6.1 and later , SLES 10 SP4, SLES 11 SP1 and later). 3PAR OS version must be 3.2.1 or later[1]. |

*Table Continued*



https://t.me/learningnets

| Rule number | Description |
|---|---|
| 5 | 3PAR StoreServ storage systems 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.<br>• MAC OS and OpenVMS do not support Direct Connect.<br>• Direct Connect to BladeSystem Servers is not supported.<br>• HP-UX, Linux, and Microsoft Windows are supported for direct connect.<br>• The 3PAR StoreServ 16Gb FC target adapter supports direct connection to host beginning with HP 3PAR OS 3.2.2.MU2.<br>Not all Operating Systems and HBAs support a Direct Connect configuration.<br>To determine if your OS and HBA combination supports Direct Connect to a 3PAR StoreServ storage system, see the SPOCK 3PAR Connectivity configuration sets on the SPOCK website at **http://www.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access.<br>After logging in into the SPOCK website:<br>**1.** On the left navigation panel, click **View by Array** link under **SAN Compatibility**.<br>The **Select Disk Array** window opens.<br>**2.** Click **Add to Selection** next to 3PAR.<br>A list of 3PAR storage systems will be displayed.<br>**3.** Select **3PAR Connectivity** in the **Filter Results by Configset Types:** drop-down menu at the top of the list.<br>**4.** Select your 3PAR StoreServ storage system and server operating system combination from the list.<br>A new window opens.<br>**5.** Expand the HBA list by clicking the **+** sign.<br>**6.** Now expand the specific HBA in the list. The various driver versions will be displayed along with the driver version that supports Direct Connect.<br>You can scroll the screen to the right to view all the columns in the table. |
| 6 | Overlapping storage port zones are supported if multiple operating systems share an array port. For more information, see the *HPE 3PAR Operating Systems Implementation Guides*. |
| 7 | Use storage system LUN presentation to enable/disable LUN access to specific hosts. For more information, see the HPE 3PAR Operating Systems Implementation Guides. |

*Table Continued*



https://t.me/learningnets

| Rule number | Description |
|---|---|
| 8 | 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 *HPE 3PAR Operating Systems Implementation Guides*. |
| 9 | Servers that share access to the same storage LUN require special application software (such as cluster software) to preserve data integrity. |
| 10 | Remote Copy and Peer Motion configuration rules still apply and require 1–to-1 zoning. |
| 11 | The "Persistent Port" feature (introduced in 3PAR OS 3.1.2) requires the active target port and partner port be connected to the same fabric. For more information, see the *HPE 3PAR Command Line Interface Reference Guide* or the *HPE 3PAR StoreServ Persistent Ports* Technical white paper document available at **http://www.hpe.com/info/StoreServ-Persistent-Ports-WP**. |

[1] For Windows Server 2012, OS 3.2.1 or later is required.

For information about configuring HPE 3PAR StoreServ storage using the 3PAR OS Management Console, see the *HPE 3PAR OS Management Console User's Guide*.

## Configuration parameters

For information about 3PAR configuration parameters, including maximum server connectivity, see the 3PAR OS configuration sets on the SPOCK website at **http://www.hpe.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.

## Virtual Connect Flat SAN (Direct-attach) Fibre Channel for 3PAR Storage

HPE supports Virtual Connect Flat SAN Fibre Channel for 3PAR storage using the HPE Virtual Connect FlexFabric 10 Gb/24-port Module and the HPE Virtual Connect FlexFabric 20-40 F8 Module. This provides connectivity between c-Class BladeSystems and 3PAR StoreServ Storage systems without a Fibre Channel switch or fabric, see **HPE Virtual Connect for the c-Class BladeSystem server environment** on page 181.

Virtual Connect Flat SAN is supported with the following 3PAR StoreServ Storage systems:

- 3PAR T400/T800
- 3PAR StoreServ 10400/10800
- 3PAR StoreServ 7200/7400/7450
- StoreServ 8200/8400/8440/8450 AFA
- StoreServ 20450/20800/20850 Storage

**NOTE:** Virtual Connect Flat SAN is supported from 3.2.2 with the 3PAR StoreServ 16 Gb target adapter.



https://t.me/learningnets

# HPE 3PAR data migration

The 3PAR StoreServ family of storage systems is supported for data migration using the following products:

- **HPE 3PAR Peer Motion** — Migrates block data between two 3PAR StoreServ systems without the use of an external appliance or host-based mirroring.
- **HPE Online Import for EVA Storage** — Migrates storage volumes from EVA storage systems to 3PAR StoreServ storage systems online and non-disruptively.
- **HPE 3PAR Online Import Utility** — 3PAR Online Import Utility manages the migration of data from a source third party storage system to a destination 3PAR StoreServ storage system.

## HPE 3PAR Online Import Utility

The 3PAR Online Import Utility currently supports the following third party source arrays:

- **EMC as a source array storage** — Migrates storage volumes from EMC storage systems to 3PAR StoreServ storage systems.
- **HDS as a source array storage** — Migrates storage volumes from HDS storage systems to 3PAR StoreServ storage systems.
- **IBM XiV as a source array storage** — Migrates storage volumes from IBM XiV storage systems to 3PARStoreServ storage systems.

Using 3PAR Online Import Utility, you can migrate third party storage virtual disks and host configuration information to a destination 3PAR StoreServ storage system.

3PAR Online Import Utility coordinates the movement of data from the source while servicing I/O requests from the hosts. During data migration, host I/O is served from source arrays. The host and virtual disk presentation implemented on the third party array is maintained on the destination 3PAR StoreServ storage system.

For the latest 3PAR Online Import Utility data migration support information, including storage systems, operating systems, cluster support, drivers, and version support, see the SPOCK website at **http://www.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access.

## 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 one other license-enabled function — either data migration or FCIP for standalone or concurrent operation. This section describes data migration usage and support.

- For information about FCoE, see **FCoE SAN fabrics** on page 25.
- For information about iSCSI configurations, see **MPX200 Multifunction Router with iSCSI for P6000/EVA storage** on page 359, **MPX200 Multifunction Router with iSCSI for XP storage** on page 370, and **MPX200 Multifunction Router with iSCSI for 3PAR StoreServ Storage** on page 366.
- For information about FCIP, see **MPX200 Multifunction Router with FCIP** on page 302.
- **HPE MPX200 Multifunction Router** — Provides online and offline block level data migration from HPE or third party storage systems to 3PAR StoreServ destination storage systems.



https://t.me/learningnets

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:

- **HPE Storage Works MPX200 1TB Full Chassis Data Migration License** — Includes the license to migrate 1 TB of data using an MPX200 chassis.
- **HPE Storage Works MPX200 5TB Full Chassis Data Migration License** — Includes the license to migrate 5 TB of data using an MPX200 chassis.
- **HPE 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. This license cannot be used when performing an iSCSI to iSCSI data migration. For iSCSI to iSCSI data migration, you must use 1TB or 5TB capacity licenses.

All arrays listed are supported for offline and online data migration, except where noted. For information about configuring the MPX200 for data migration, see the *HPE MPX200 Multifunction Router Data Migration Solution Guide*.



https://t.me/learningnets

### Table 117: MPX200/3PAR data migration source-destination storage systems

| Source storage systems | 3PAR destination storage systems |
|---|---|
| • All MSA (Fibre Channel) and EVA models<br>• P9500/XP24000/20000, XP12000/10000<br>• SVSP<br>• 3PAR S-Class | HPE 3PAR StoreServ 8200/8400/8440/8450; 3PAR StoreServ 20450/20800/20850; HPE 3PAR StoreServ 10400/10800; 3PAR StoreServ 7450; HPE 3PAR StoreServ 7200/7400; 3PAR F-Class, T-Class |
| Third-party array models:<br>• Dell EqualLogic family (iSCSI), Compellent Series 30 and 40 Controllers<br>• EMC CLARiiON AX series, CX Series, Symmetrix DMX Series, Symmetrix VMAX SE (offline only), VNX5500 (offline only), VNX5300<br>• Fujitsu ETERNUS DX400, DX440 S2 (offline only), DX8400<br>• Hitachi Data Systems V series (offline only), AMS Family, WMS, DF-500, USP (offline only), VSP families (offline only), TagmaStore Network StorageController model NSC55<br>• IBM DS3000 series, DS4000 series (offline only), DS5000 series, SVC, DS6000 series (offline only), DS8000 series, XIV Storage System family (offline only), nlStorwize V7000 Unified disk system (offline only)<br>• NEC D-Series SAN Storage arrays (offline only<br>• NetApp FAS270, FAS250, FAS2000 Series, FAS3100 Series, FAS6000 Series (Fibre Channel and iSCSI)<br>• Xiotech Emprise 5000, Mag3D 4000 | |



https://t.me/learningnets

**Table 118: MPX200/Online data migration operating system support**

| MPX200 online data migration support[1] | Online data migration destination storage system and firmware (minimum) |
|---|---|
| • Citrix XenServer 6.0<br>• HP-UX 11iv3, 11iv2, Clusters (Service Guard)<br>• IBM AIX 6.1, 5.3[2]<br>• Microsoft Windows 2012 R2, 2012, 2008 R2, 2008 and 2003; Failover Clustering and MSCS<br>• Red Hat 6, 5 U4 , 4 U8, U7, U6, U3, Clusters (RH 6)<br>• SUSE 11, 10 SP3, SP1, Clusters (SUSE 11 U1)<br>• Oracle Solaris 10, Clusters 3.3<br>• VMware ESXi 5.1, 5.0, 4.1, 4.0, 3.5, Clusters | • P2000 G3 FC (TS251P002-04)<br>• P4000 (9.0)<br>• P6350/P6550 (11001000)<br>• P6500/P6300 (10001000)<br>• EVA8100/6100/4100 (6.220)<br>• EVA8400/6400/4400 (09522000, 09534000)<br>See **P6000/EVA data migration** on page 243.<br><br>• P9500 (70-00-50-00)<br>• XP24000/20000 (60-06-10-00)<br>• XP12000/10000 (50-09-83-00)<br>See **P9000/XP data migration** on page 257.<br><br>• HPE 3PAR StoreServ 7200, 3.1.3 (MU1)<br>• HPE 3PAR StoreServ 7400, 3.1.3 (MU1)<br>• HPE 3PAR StoreServ 7450, 3.1.3 (MU1)<br>• HPE 3PAR StoreServ 7200c, 3.2.1<br>• HPE 3PAR StoreServ 7400c, 3.2.1<br>• HPE 3PAR StoreServ 7440c, 3.2.1<br>• HPE 3PAR StoreServ 7450c, 3.2.1<br>• 3PAR StoreServ 8200, 3.2.2<br>• 3PAR StoreServ 8400, 3.2.2<br>• 3PAR StoreServ 8440, 3.2.2<br>• 3PAR StoreServ 8450, 3.2.2<br>• HPE 3PAR StoreServ 10400, 3.1.3 (MU1)<br>• HPE 3PAR StoreServ 10800, 3.1.3 (MU1)<br>• 3PAR StoreServ 20450, 3.2.2<br>• 3PAR StoreServ 20800, 3.2.2<br>• 3PAR StoreServ 20850, 3.2.2<br>• HPE 3PAR F-Class, T-Class, 2.3.1 (MU2) |


https://t.me/learningnets

| MPX200 online data migration support[1] | Online data migration destination storage system and firmware (minimum) |
|---|---|
| | See **HPE 3PAR data migration** on page 265. |

1 For operating system updates, the minimum supported version is specified.

2 IBM AIX 5.3 not supported with 3PAR OS 3.1.2 or later.

For the latest data migration storage system, operating system, and version support, see the SPOCK website at **http://www.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access.

**HPE 3PAR data migration considerations**

MPX200 connectivity to 3PAR StoreServ 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 StoreServ storage system must be configured based on current 3PAR support for operating systems, HBAs, and multipath software. For more information, see the *HPE 3PAR Operating Systems Implementation Guides*.

# Zoning

These zoning rules apply to a heterogeneous SAN with 3PAR StoreServ storage:

- Zoning by HBA is the HPE recommended zoning method for 3PAR StoreServ storage, see **Configuration rules** on page 262.
- Zoning is required for all operating systems that access 3PAR StoreServ storage, see **Heterogeneous server rules** on page 179.
- 3PAR StoreServ storage systems can be in multiple operating system zones.

## 3PAR zoning best practices

3PAR zoning best practices considerations:

- Single initiator—Single target zones should be preferred (1 to 1 zones)
- WWPN zoning is recommended. For more information, see the *HPE 3PAR StoreServ Storage Best Practices Guide*, technical whitepaper and FC host zoning section. The document is available at **http://h20195.www2.hpe.com/v2/GetPDF.aspx/2F4AA4-4524ENW.pdf**.
- A host needs to be zoned with a minimum of 2 ports from the 2 nodes of the same pair

  Example: 0:2:1 and 1:2:2
- Hosts do not need to be connected to all nodes; 2 nodes are adequate
- No more than 64 initiators can be connected per front-end port
- Each 3PAR system has a maximum number of initiators supported
- 1 Initiator = 1 path from a host:



https://t.me/learningnets

- a host zoned with 2 FC ports will take 2 initiators
- a host zoned with 4 FC ports will take 4 initiators

- In order to keep the number of initiators below the maximum supported value, use the following recommendation:
  - Normal hosts (not hypervisors): 2 paths
  - Hypervisors: 4 paths maximum (not necessary if only 2 nodes)

# Cabling

This section describes best practice cabling recommendations for 3PAR StoreServ 10000 V-Class; 3PAR StoreServ 7000;3PAR F-Class, T-Class storage systems.

- Each controller node should be connected to both fabrics in a dual fabric SAN
- Odd ports should connect to fabric 1, even ports to fabric 2
- Ports of the same pair of nodes with the same ID should be connected to the same fabric
- Example:
  - Fabric 1 = 0:2:3 and 1:2:3
  - Fabric 2 = 0:2:4 and 1:2:4
- Example for 4 node F400 with 8 front-end ports:
  - Fabric 1 = 0:2:3, 1:2:3, 2:2:3, 3:2:3
  - Fabric 2 = 0:2:4, 1:2:4, 2:2:4, 3:2:4
- Example for 4 node V400 with 32 front-end ports:
  - Fabric 1 0:2:1, 0:2:3, 0:5:1, 0:5:3,1:2:1, 1:2:3, 1:5:1, 1:5:3, 2:2:1, 2:2:3, 2:5:1, 2:5:3,3:2:1, 3:2:3, 3:5:1, 3:5:3
  - Fabric 2 = 0:2:2, 0:2:4, 0:5:2, 0:5:4,1:2:2, 1:2:4, 1:5:2, 1:5:4, 2:2:2, 2:2:4, 2:5:2, 2:5:4, 3:2:2, 3:2:4, 3:5:2, 3:5:4

# Tape storage

The following rules apply to tape storage support in a SAN with 3PAR StoreServ 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.
- Hewlett Packard Enterprise recommends using a separate tape-to-HBA connection for servers when backups require more than four DLT8000 tape drives or two Ultrium (LTO) tape drives.


https://t.me/learningnets

For more information about tape storage support, contact your Hewlett Packard Enterprise representative.

# HPE 3PAR SAN boot support

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

# HPE 3PAR storage management

For information about 3PAR StoreServ storage management, see the *HPE 3PAR OS Management Console User's Guide*.



https://t.me/learningnets

# HPE Data Availability, Protection and Retention

One of the most significant benefits of a SAN is the ability to share the SAN infrastructure for both diskand 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 HPE solution is the HPE Data Availability, Protection and Retention Compatibility Matrix **http://www.hpe.com/storage/daprcompatibility**.

Inside the matrix, there is a complete list of Independent Software Vendors (ISVs), and Independent Hardware Vendors (IHVs) who have develop or enhance products that integrate with HPE technology. There is also a Solutions Whitepapers page with links to the Backup, Recovery and Archive (BURA) Solutions design guide and other whitepapers to help with the SAN configuration.



https://t.me/learningnets

# SAN extension and bridging

SAN extension and bridging are presented in these chapters:


- **SAN extension** on page 274
- **iSCSI storage** on page 341




https://t.me/learningnets

# 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** on page 274
- **Fibre Channel long-distance technology** on page 277
- **Multi-protocol long-distance technology** on page 284
- **HPE multi-protocol long-distance products** on page 291
- **HPE storage replication products** on page 315
- **Certified third-party WDM, iFCP, and SONET products** on page 340


## 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:
  - 100 m for 32 Gb/s Fibre Channel
  - 125 m for 16 Gb/s Fibre Channel
  - 190 m for 8 Gb/s Fibre Channel
  - 400 m for 4 Gb/s Fibre Channel
  - 500 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 OM4 fiber optic cable. For the supported distances for each Fibre Channel speed and interconnect type, see **SAN fabric connectivity rules** on page 163.

You can implement a SAN extension with any Fibre Channel topology.



https://t.me/learningnets

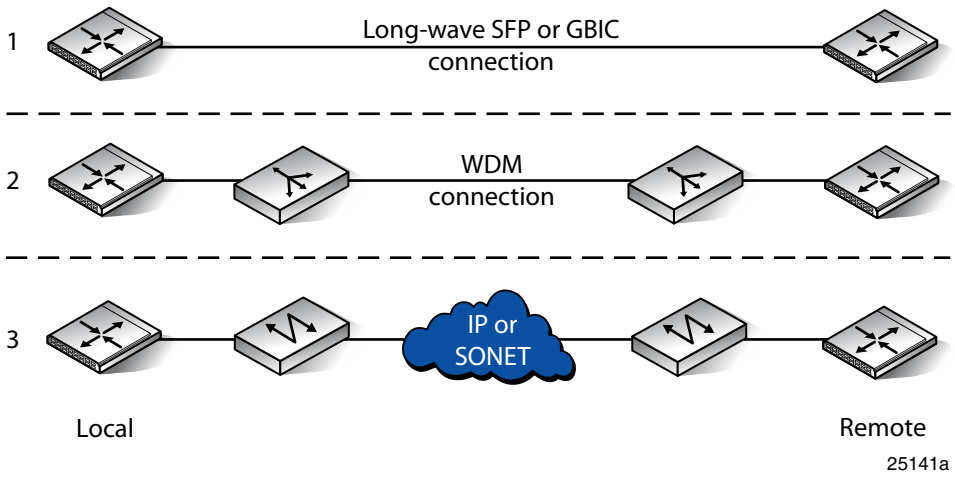



**Figure 71: SAN extension examples**

This section describes:

- **SAN extension technology** on page 275
- **SAN-iSCSI bridging technology** on page 276

## SAN extension technology

Hewlett Packard Enterprise supports the following SAN extension technologies:

- **Fibre Channel long-distance technology** on page 277, including:
  - **Fiber optic transceivers** on page 277
  - **Wavelength division multiplexing** on page 279
  - **Extended fabric settings for Fibre Channel switches** on page 283
- **Multi-protocol long-distance technology** on page 284, including:
  - **Fibre Channel over Internet Protocol** on page 285
  - **Fibre Channel over SONET** on page 289
  - **Fibre Channel over ATM** on page 291



https://t.me/learningnets

**Table 119: SAN extension technologies and HPE SAN extension products**

| SAN extension technology | HPE product |
|---|---|
| Fibre Channel using long-wave transceivers (10 km–35 km) | B-series, C-series, and H-series switches, see Long-wave transceiver distances table |
| WDM (greater than 35 km to 100–500 km) | • C-series switches with the HPE CWDM solution<br>• B-series and C-series switches with certified third-party products, see **Certified third-party WDM, iFCP, and SONET products** on page 340 |
| FCIP (greater than 10 km to 20,000 km) | B-series 1606 SAN Extension Switch, StoreFabric SN4000B SAN Extension Switch, Multi-protocol Router Blade (MP Router Blade FR4-18i), or the MP Extension Blade (FX8-24) with B-series switches<br><br>C-series SN8000C 6-Slot Supervisor 2A Director Switch, SN8000C 9-Slot Supervisor 2A Director Switch, SN8000C 13-Slot Supervisor 2A Fabric 2 Director Switch, MDS 9506, 9509, 9513, and 9222i switches using the IPS-4, IPS-8, 14/2 Multiprotocol Services Module, or 18/4 Multiservice Module<br><br>• B-series, C-series, and H-series switches with the MPX200 Multifunction Router with FCIP<br>• B-series, C-series, and H-series switches with the IP Distance Gateway (mpx110)<br><br>Certified third-party productsSee **Certified third-party WDM, iFCP, and SONET products** on page 340 |
| FC-SONET | See **Certified third-party WDM, iFCP, and SONET products** on page 340 |
| FC-ATM | See **Fibre Channel over ATM** on page 291 |

## SAN-iSCSI bridging technology

SAN-iSCSI bridging connects Fibre Channel networks and IP networks. Hewlett Packard Enterprise supports the following iSCSI to Fibre Channel bridging devices:

- C-series SN8000C 6-Slot Supervisor 2A Director Switch, SN8000C 9-Slot Supervisor 2A Director Switch, SN8000C 13-Slot Supervisor 2A Fabric 2 Director Switch, MDS 9506, 9509, 9513, and 9222i embedded IP ports
- C-series IP Storage Services Modules (IPS-4, IPS-8)
- C-series MDS 18/4 Multiservice Module
- MPX200 Multifunction Router with iSCSI


https://t.me/learningnets

For information about bridging with iSCSI, see **iSCSI storage** on page 341.

# Fibre Channel long-distance technology

This section describes the following Fibre Channel long-distance methodology and hardware for a SAN extension:

- **Fiber optic transceivers** on page 277
- **Wavelength division multiplexing** on page 279
- **Extended fabric settings for Fibre Channel switches** on page 283

## 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 32 Gb/s, 16 Gb/s, 10 Gb/s, and 8 Gb/s transceivers are known as SFP transceivers. The 4 Gb/s, and 2 Gb/s transceivers are known as SFPs. All use LC style connectors (**Figure 72: LC SFP transceiver** on page 277).

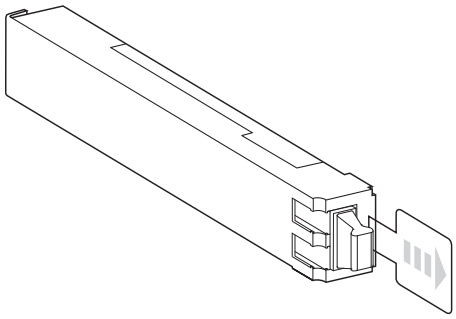



**Figure 72: LC SFP transceiver**

- The 1 Gb/s transceivers can be LC SFPs (**Figure 72: LC SFP transceiver** on page 277) or GBICs, which use SC style connectors (**Figure 73: SC GBIC transceiver** on page 277).

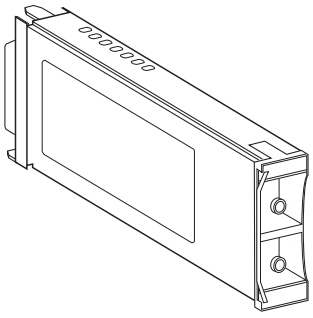



**Figure 73: SC GBIC transceiver**

Some long-wave optical transceivers can transmit up to a distance of 100 km.



https://t.me/learningnets

### Table 120: Long-wave transceiver distances

| Interconnect speed | Transceiver | Distance/rate |
|---|---|---|
| 32 Gb/s | 32 Gb/s 10 km SFP+ | 10 km ISL at 32 Gb/s (B-series) |
| 16 Gb/s | 16 Gb/s 10 km SFP+ | 10 km ISL at 16 Gb/s (B-series) |
| | 16 Gb/s 25 km SFP+ | 25 km ISL at 16 Gb/s (B-series) |
| 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)[1]<br>10 km ISL at 8 Gb/s (B-series and C-series)<br>3.3 km ISL at 8 Gb/s (H-series, base switch)<br>6.6 km ISL at 4 Gb/s (H-series, base switch)<br>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<br>(B-series, C-series) |
| | 4 Gb/s 4 km SFP | 4 km ISL at 4 Gb/s<br>(C-series) |
| | 4 Gb/s 30 km SFP (extended-reach) | 30 km ISL at 4 Gb/s<br>(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<br>(B-series) |
| | 2 Gb/s 10 km SFP | 10 km ISL at 2 Gb/s<br>(B-series, C-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) |
| | 1 Gb/s 35 km SFP | 35 km ISL at 1 Gb/s (C-series) |
| | 1 Gb/s 100 km GBIC | 100 km ISL at 1 Gb/s (B-series) |

[1] 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 the following tables:



https://t.me/learningnets

- Fibre Channel distance rules for 32 Gb/s, 16 Gb/s, and 8 Gb/s switch models (B-series, C-series, and H-series)
- Fibre Channel distance rules for 4 Gb/s switch models (B-series and C-series switches)
- Fibre Channel distance rules for 2 Gb/s switch models (B-series and C-series switches)
- Fibre Channel distance rules for 1 Gb/s switch models (B-series and C-series switches)

**NOTE:** To ensure adequate link performance, see **Extended fabric settings for Fibre Channel switches** on page 283.

# Wavelength division multiplexing

This section describes the following:


- **WDM overview** on page 279
- **WDM network implementation** on page 279
- **WDM system architectures** on page 279
- **WDM system characteristics** on page 280
- **HPE coarse wave division multiplexing** on page 281
- **Third-party WDM products** on page 282


## 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:



https://t.me/learningnets

- Passive (optical transmission protocol)
- Active signal amplification
- Active protocol handling

Most WDM products use one of these architectures or combine attributes of each.

### Table 121: WDM system architectures

| System architecture[1] | Description |
|---|---|
| Passive (optical transmission protocol) | • Transparent to transmission protocol and data-rate independent<br>• Establishes open interfaces that provide flexibility to use Fibre Channel, SONET/SDH, ATM, Frame Relay, and other protocols over the same fiber<br>• Passes the optical signal without any form of signal conditioning such as amplification or attenuation |
| Active signal amplification | • Includes line amplifiers and attenuators that connect to other devices through fiber optic links<br>• Boosts the signals that are transmitted to and received from peripheral network devices<br>• Using hardware and/or software control loops, monitors power levels to ensure that the operation does not exceed the hardware's power budgets |
| Active protocol handling | • Offers protocol-specific capabilities for Fibre Channel, enabling digital TDM and optical multiplexing to support multiple channels on each wavelength<br>• Provides network monitoring, digital retiming (to reduce timing jitter), link integrity monitoring, and distance buffering<br>• May require additional and potentially costly transmission hardware when deployed in meshed networks<br>**NOTE:** Continuous Access products using 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. |

[1] Active protocol handling and passive protocol handling require different switch port settings, see Port protocol setting based on the extension architecture table.

## WDM system characteristics

To help carriers realize the full potential of WDM, Hewlett Packard Enterprise-supported WDM systems have the following characteristics:



https://t.me/learningnets

- 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

## HPE coarse wave division multiplexing

HPE 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



https://t.me/learningnets

- 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://www.hpe.com/info/CoarseWaveDivisionMultiplexer**.

**NOTE:** The CWDM multiplexer solution is not supported on B-series or H-series switches.

For more information about B-series and C-series supported WDM devices, see **Certified third-party WDM, iFCP, and SONET products** on page 340.

## Third-party WDM products

Hewlett Packard Enterprise supports third-party WDM products that have been tested successfully with a wide range of HPE products, see **Certified third-party WDM, iFCP, and SONET products** on page 340.

**Figure 74: Basic WDM configuration using one long-distance fiber optic link** on page 282 and **Figure 75: Fully redundant WDM configuration using two long-distance fiber optic links** on page 282 show basic WDM SAN configuration options for typical Continuous Access storage systems. For specific configuration options, see the P6000 Continuous Access and P9000 (XP) Continuous Access product documentation.

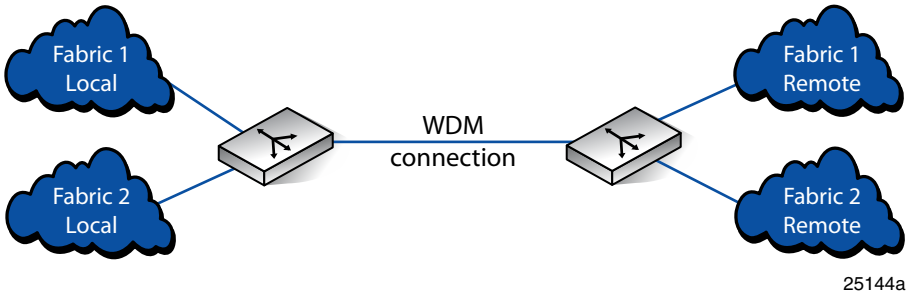



**Figure 74: Basic WDM configuration using one long-distance fiber optic link**

The configuration in **Figure 74: Basic WDM configuration using one long-distance fiber optic link** on page 282 is low cost, but has no long-distance link redundancy.

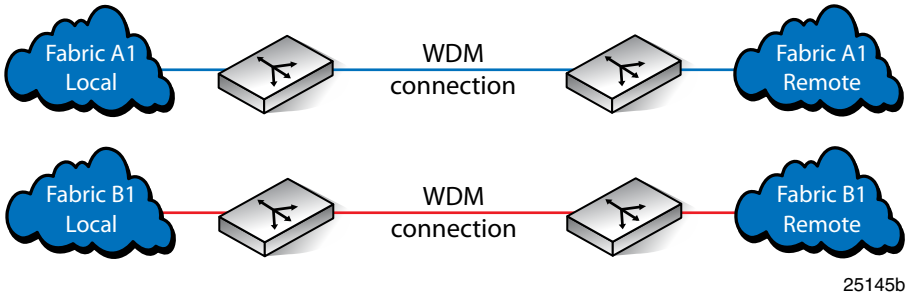



**Figure 75: Fully redundant WDM configuration using two long-distance fiber optic links**

The configuration in **Figure 75: Fully redundant WDM configuration using two long-distance fiber optic links** on page 282 is high cost, but is fully redundant.

Hewlett Packard Enterprise 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



https://t.me/learningnets

- 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** on page 169.
- Performance levels depend on the number of buffers available in the switch and the amount of application data.

For information about Hewlett Packard Enterprise-certified third-party WDM products, see **Certified third-party WDM, iFCP, and SONET products** on page 340.

# 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** on page 109.

## 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 *HPE Brocade Fabric OS Administrator's Guide*, available at **http://www.hpe.com/info/StorageWorks-SANDS-Manuals**.

**NOTE:** The C-series CWDM solution is not supported on B-series switches. For B-series supported WDM devices, see **Certified third-party WDM products** on page 340.

### Table 122: Port protocol setting based on the extension architecture

| WDM system architecture | B-series port protocol setting |
|---|---|
| Active protocol handling (see WDM system architectures table) | `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 *HPE StorageWorks Fabric OS Administrator Guide*.
- Trunking is not supported with ports configured for portCfgMode **slot/port,** 1 (R_RDY mode).



https://t.me/learningnets

### C-series switch settings

#### Table 123: 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 table. |
| Maximum number of hops | 7 |
| Maximum segment distance | For more information about distance rules for long-wave transceivers, see Fibre Channel distance rules for 4 Gb/s switch models (B-series and C-series switches) table, Fibre Channel distance rules for 2 Gb/s switch models (B-series and C-series switches) table, and Fibre Channel distance rules for 1 Gb/s switch models (B-series and C-series switches) table in **SAN fabric connectivity and switch interoperability rules** on page 163. |

Consider the following when using C-series switches with extended fabric links:

- The MDS 9222i Multiservice Fabric and 18/4 Multiservice Modules can allocate up to 4,095 BB_credits to an individual port.
- 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.

# Multi-protocol long-distance technology

This section describes the following storage replication technologies, which enable data transfer between SAN networks:



https://t.me/learningnets

- **Fibre Channel over Internet Protocol** on page 285
- **Fibre Channel over SONET** on page 289
- **Fibre Channel over ATM** on page 291


# 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** on page 285
- **FCIP link configurations** on page 285
- **FCIP network considerations** on page 286
- **FCIP bandwidth considerations** on page 287
- **FCIP gateways** on page 289
- **Third-party QoS and data encryption FCIP products** on page 289


## 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 switches connect through an E_Port, while C-series switches use plug-in modules. **Figure 76: FCIP single-link configuration** on page 286, **Figure 77: FCIP dual-link configuration** on page 286, and **Figure 78: FCIP shared-link configuration** on page 286 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 76: FCIP single-link configuration** on page 286).



https://t.me/learningnets




**Figure 76: FCIP single-link configuration**

## FCIP dual-link configuration

A dual-link configuration provides redundancy (**Figure 77: FCIP dual-link configuration** on page 286). If one link fails, the other link temporarily handles all data replication. For enhanced fault tolerance, you can use two IP providers.

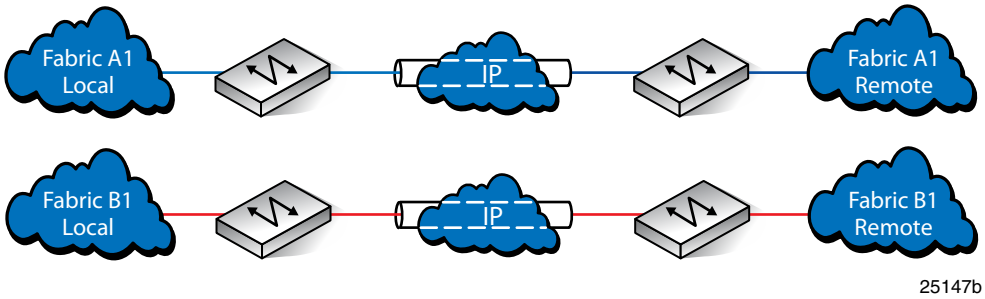



**Figure 77: FCIP dual-link configuration**

In a dual-link configuration, Hewlett Packard Enterprise 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 78: FCIP shared-link configuration** on page 286).

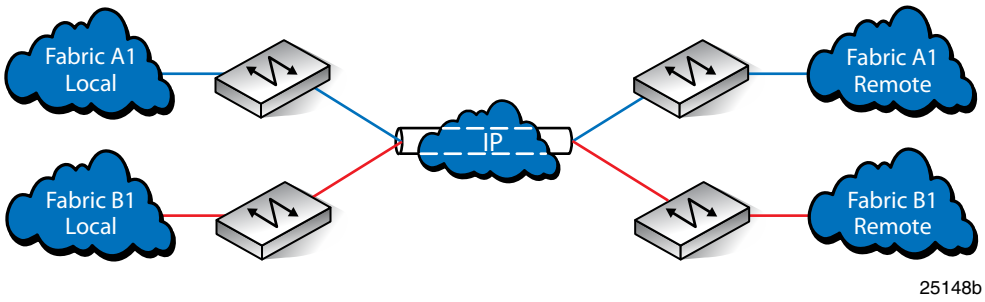



**Figure 78: FCIP shared-link configuration**

**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.

## 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.



https://t.me/learningnets

### Table 124: FCIP network consideration

| Configuration type | Use existing network? | Factors |
|---|---|---|
| Mirrored FCIP SAN | No | For peak performance, Hewlett Packard Enterprise 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 a P6000 Continuous Access environment with the EVA family of storage arrays:

**Procedure**

1. 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.
2. Increase link bandwidth while normalization is taking place.
3. 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.
4. 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:
   a. Create VPNs with QoS through your local routers for the replication circuit.
   b. Create separate physical networks for EVA storage arrays.
   c. Guarantee the bandwidth using a third-party router and/or QoS vendor.
5. Distance affects the amount of data that can be transmitted across a link. Consider these site-planning best practices:
   a. Use the shortest possible distance between remote sites.
   b. 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.



https://t.me/learningnets

**c.** Copy only data that is more expensive to re-create than to copy.

**d.** Add copy sets that will not impact normal data traffic.

**e.** Consider adding controller pairs to use available bandwidth effectively.

For additional recommendations, see the *HPE 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:

- 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 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.
- 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 P6000 Continuous Access, the RPO design space is near zero.
    - P9000 (XP) Continuous Access Asynchronous supports an RPO from near zero to many hours.
    - P6000 Continuous Access and P9000 (XP) Continuous Access all synchronous replication with an RTO equal to zero.
- 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.



https://t.me/learningnets

**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 125: HPE FCIP gateways

| Gateway | Supported switches |
|---|---|
| **B-series 1606 Extension SAN Switch and DC Dir Switch MP Extension Blade** on page 310<br>**HPE StoreFabric SN4000B SAN Extension Switch** on page 306 | B-series switches |
| **C-series MDS 9222i, IPS-4, IPS-8, 18/4 Multiservice Modules** on page 314 | C-series switches |
| **HPE IP Distance Gateway (mpx110)** on page 293 | B-series, C-series, and H-series switches |
| **MPX200 Multifunction Router with FCIP** on page 302 | B-series, C-series, and H-series switches |

### 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 126: FCIP QoS products

| Vendor name/website | Device | Purpose | Support |
|---|---|---|---|
| Allot Communications, Inc.<br>**http://www.allot.com** | NetEnforcer | Application traffic and bandwidth management system | P6000 Continuous Access |
| **http://www.packeteer.com** | PacketShaper 6500 | | |
| **http://www.riverstonenet.com** | RS 3000, RS 8000 | MPLS GbE routers | |

## WAN accelerator products

WAN accelerator products are not supported for use with Continuous Access products.

## Fibre Channel over SONET

You can connect local SANs with a SONET to create an extended SAN.



https://t.me/learningnets

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** on page 290
- **FC-SONET network considerations** on page 291
- **Third-party SONET gateways** on page 291


## 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 79: FC-SONET dual-link configuration** on page 290). 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.

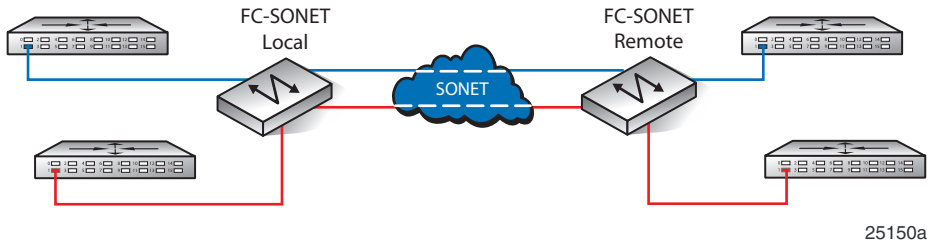



**Figure 79: FC-SONET dual-link configuration**

In a dual-link configuration, Hewlett Packard Enterprise 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 80: FC-SONET shared-link configuration** on page 290).

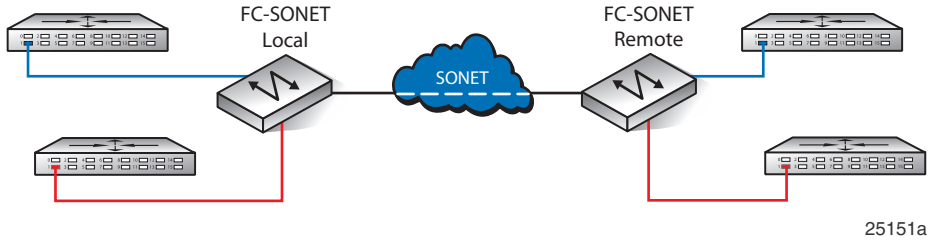



**Figure 80: FC-SONET shared-link configuration**

**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.



https://t.me/learningnets

### 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.

HPE supports the use of SONET with P9000 (XP) Continuous Access, see **Certified third-party WDM, iFCP, and SONET products** on page 340.

### Third-party SONET gateways

For a list of Hewlett Packard Enterprise-certified third-party SONET gateways, see **Certified third-party WDM, iFCP, and SONET products** on page 340.

## 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 table.

# HPE multi-protocol long-distance products

This section describes the following SAN extension products:

- **HPE IP Distance Gateway (mpx110)** on page 293
- **MPX200 Multifunction Router with FCIP** on page 302
- **B-series 1606 Extension SAN Switch and DC Dir Switch MP Extension Blade** on page 310
- **HPE StoreFabric SN4000B SAN Extension Switch** on page 306
- **C-series MDS 9222i, IPS-4, IPS-8, 18/4 Multiservice Modules** on page 314

## HPE SAN extension products summary and usage

Hewlett Packard Enterprise provides a full line of SAN extension products designed to satisfy a range of disaster recovery solutions and requirements.



https://t.me/learningnets

**Table 127: Features and usage for HPE 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 |
|---|---|---|---|---|---|---|
| IP Distance Gateway (mpx110) | B-series<br>C-series<br>H-series | E_Port<br>(fabric merge)<br>EX_Port<br>(LSAN fabric isolation)<br>E_Port<br>(VSAN fabric isolation) | P6000 Continuous Access<br>P9000 (XP) Continuous Access[3] | Software<br>(Use when RTT is ≥ 50 ms; or if guaranteed WAN bandwidth is ≤ 45 Mb/s) | Low: T3/DS3 (45 Mb/s)<br>to<br>Medium: OC-3 (155 Mb/s) | See IP Distance Gateway (mpx110) features and requirements table. |
| MPX200 Multifunction Router with FCIP (MPX200) | B-series<br>C-series<br>H-series | E_Port<br>(fabric merge)<br>EX_Port<br>(LSAN fabric isolation)[1]<br>E_Port<br>(VSAN fabric isolation)[2] | | Software[4] | Low: T3/DS3 (45 Mb/s)<br>to<br>Medium: OC-3 (155 Mb/s)<br>to<br>High: OC-6 (322 Mb/s), OC-12 (622 Mb/s) up to 1 Gb/s | See MPX200 Multifunction Router using FCIP features and requirements table. |
| B-series 1606 Extension SAN Switch<br>DC Dir Switch MP Extension Blade | B-series | VEX_Port<br>(LSAN fabric isolation) | | Hardware | Low: T3/DS3 (45 Mb/s)<br>to<br>Medium: OC-3 (155 Mb/s)<br>to<br>High: OC-6 (322 Mb/s), OC-12 (622 Mb/s) up to 1 Gb/s | See 1606 Extension SAN Switch and DC Dir Switch MP Extension Blade features and requirements table. |
| StoreFabric SN4000B SAN Extension Switch | B-series | VE_Port (EX for LSAN fabric isolation) | | Hardware | | See HPE StoreFabric SN4000B SAN Extension Switch features and requirements table |

*Table Continued*



https://t.me/learningnets

| FCIP gateway product | Supported fabrics | Inter-fabric connectivity | Supported DR software | Data compression method/use | Recommended IP bandwidths (WAN) | Network requirements |
|---|---|---|---|---|---|---|
| B-series<br>MP Router Blade | B-series | VEX_Port<br>(LSAN fabric isolation) | | Hardware | | See MP Router Blade features and requirements table. |
| C-series MDS 9222i, IPS-4, IPS-8, 18/4 | C-series | E_Port<br>(VSAN-IVR fabric isolation) | | Software (IPS-4, IPS-8)<br>Hardware (all others) | | See C-series MDS module features and requirements table. |

[1] LSAN fabric isolation is available when used with B-series switches with Integrated Routing or B-series routers with FCR.

[2] VSAN fabric isolation is available when used with C-series switches with inter-VSAN routing.

[3] P9000 (XP) Continuous Access is not supported by H-series switches.

[4] For compression usage recommendations, see the IP performance tuning section in the *HPE MPX200 Multifunction Router User Guide*.

## HPE IP Distance Gateway (mpx110)

The 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 128: IP Distance Gateway (mpx110) features and requirements

| Feature | Requirements |
|---|---|
| Fibre Channel switches for FCIP | B-series switches—See B-series Fibre Channel switches table and B-series Fibre Channel switches and routers table. |
| IP network protocols | TCP/IP IPv4, Ethernet 10 Mb/s, 100 Mb/s, 1,000 Mb/s<br>Requires dedicated IP bandwidth. See Network requirements for long-distance IP gateways with XCS 6.x table, and Network requirements for long-distance IP gateways with VCS 4.x table. |
| Storage systems | • P6000 Continuous Access<br>• P9000 (XP) Continuous Access<br>• 3PAR Remote Copy |
| FCIP-supported operating systems | For P6000 Continuous Access, see the *HPE P6000 Enterprise Virtual Array Compatibility Reference Guide* at **http://www.hpe.com/info/P6000-ContinuousAccess**. |
| Documentation | For information about using the IP Distance Gateway, see **http://www.hpe.com/info/StoreFabric**. |



https://t.me/learningnets

## IP Distance Gateway configuration examples

The IP Distance Gateway supports the configurations shown in **Figure 81: IP Distance Gateway basic FCIP configuration with one or two long-distance links** on page 294 through **Figure 91: HPE P6000 Continuous Access 3-site configuration with eight gateways** on page 300.

**Figure 81: IP Distance Gateway basic FCIP configuration with one or two long-distance links** on page 294 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.

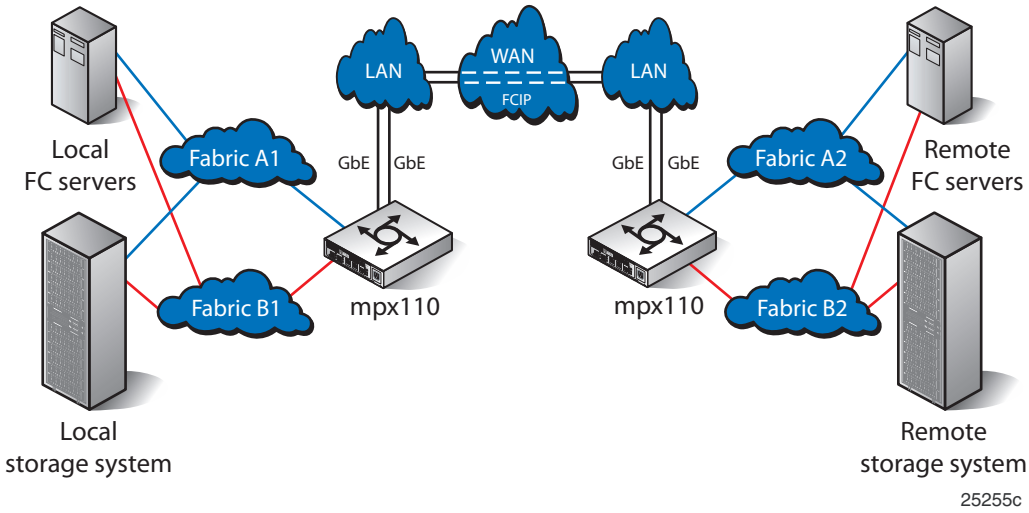



**Figure 81: IP Distance Gateway basic FCIP configuration with one or two long-distance links**

**Figure 82: IP Distance Gateway high-availability configuration with one or two long-distance links** on page 294 shows a high-availability configuration using pairs of mpx110s at the local site and remote site for path redundancy.

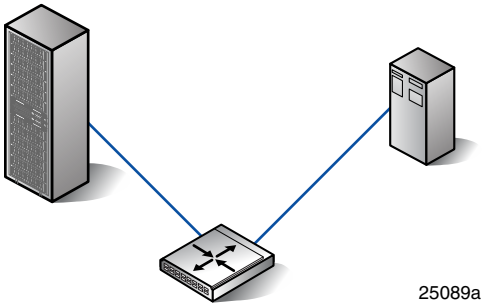



**Figure 82: IP Distance Gateway high-availability configuration with one or two long-distance links**

**Figure 83: IP Distance Gateway high-availability configuration with a redundant IP network** on page 295 shows a high-availability configuration that includes a redundant IP network.



https://t.me/learningnets

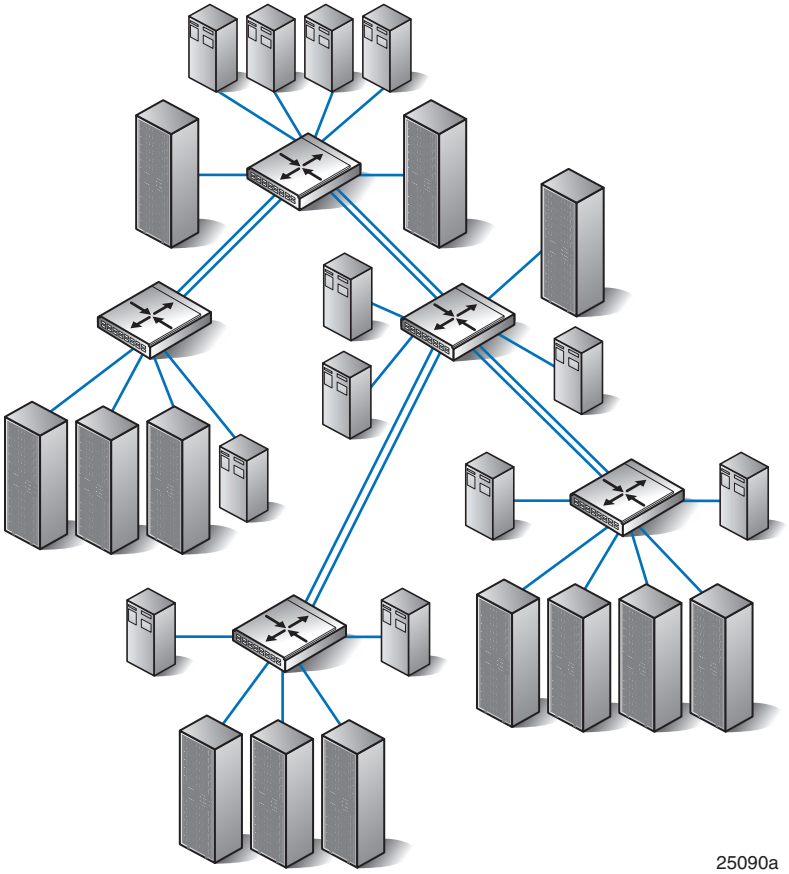



**Figure 83: IP Distance Gateway high-availability configuration with a redundant IP network**

**Figure 84: IP Distance Gateway basic FCIP configuration with single-path connectivity** on page 295 shows single-path connectivity for the servers and storage. This is the lowest-cost implementation with no redundancy.

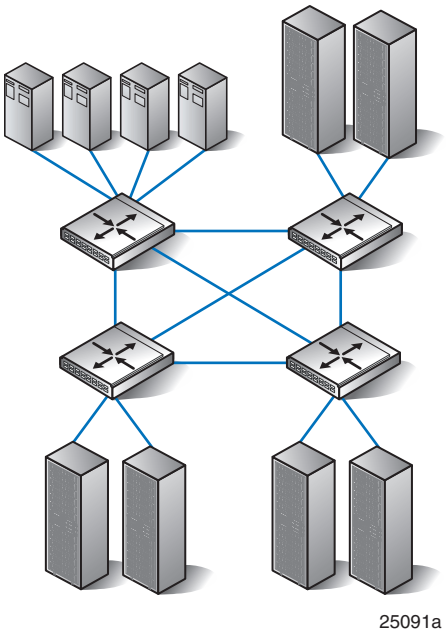



**Figure 84: IP Distance Gateway basic FCIP configuration with single-path connectivity**



https://t.me/learningnets

**Figure 85: IP Distance Gateway FCIP with B-series Integrated Routing** on page 296 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.

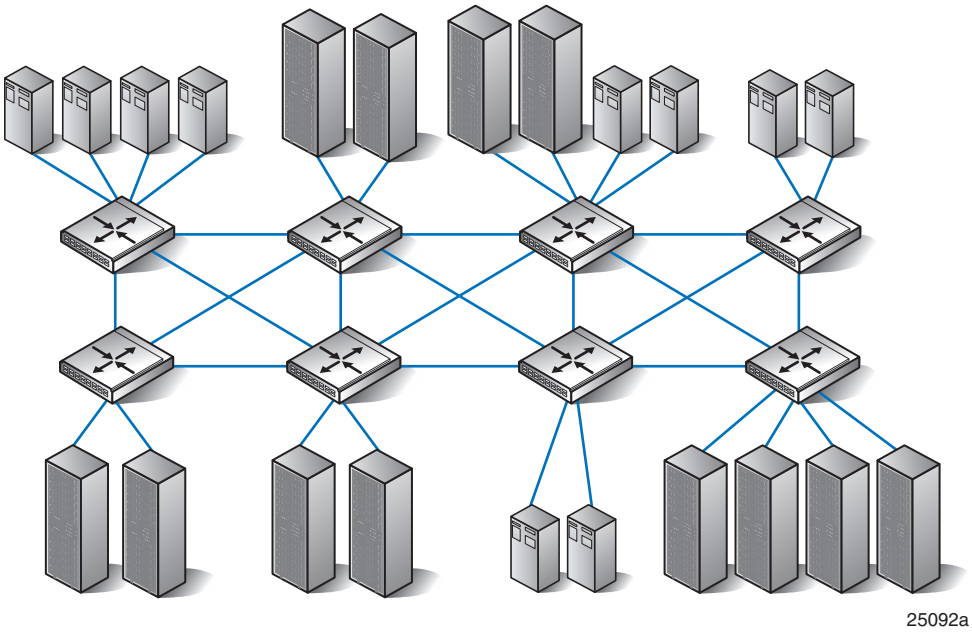



**Figure 85: IP Distance Gateway FCIP with B-series Integrated Routing**

**Figure 86: IP Distance Gateway FCIP with C-series IVR** on page 297 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.



https://t.me/learningnets

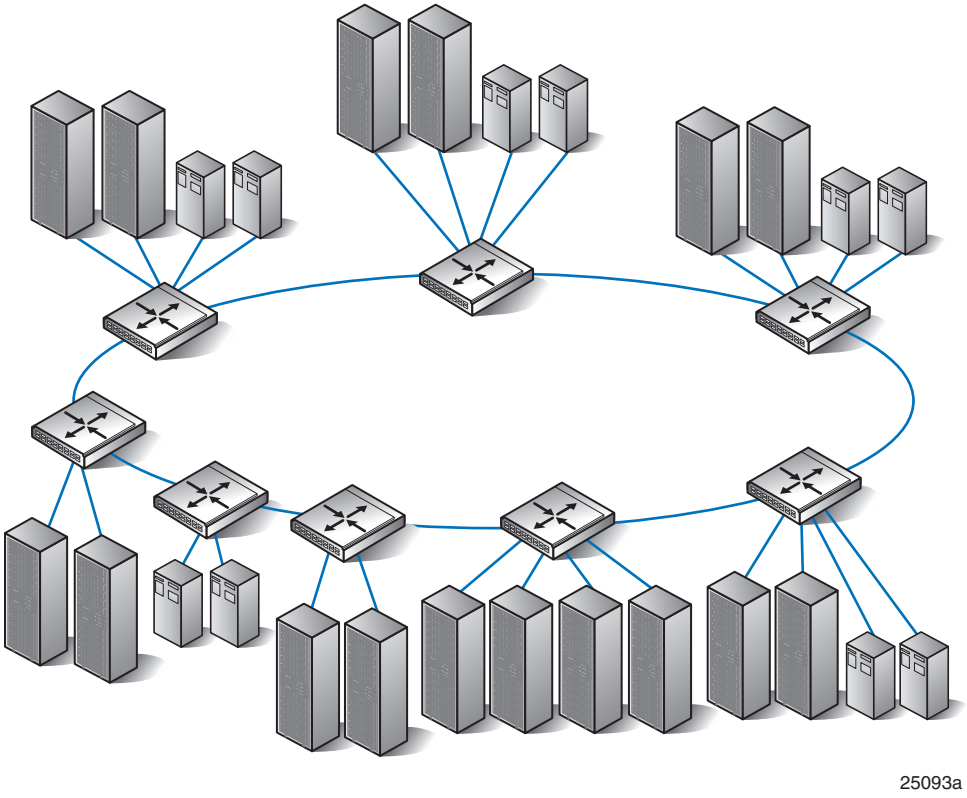



**Figure 86: IP Distance Gateway FCIP with C-series IVR**

**Figure 87: Highly redundant pairs of gateways, two long distance links** on page 297 is similar to **Figure 82: IP Distance Gateway high-availability configuration with one or two long-distance links** on page 294, but offers a higher level of redundancy using additional Fibre Channel and LAN connections.

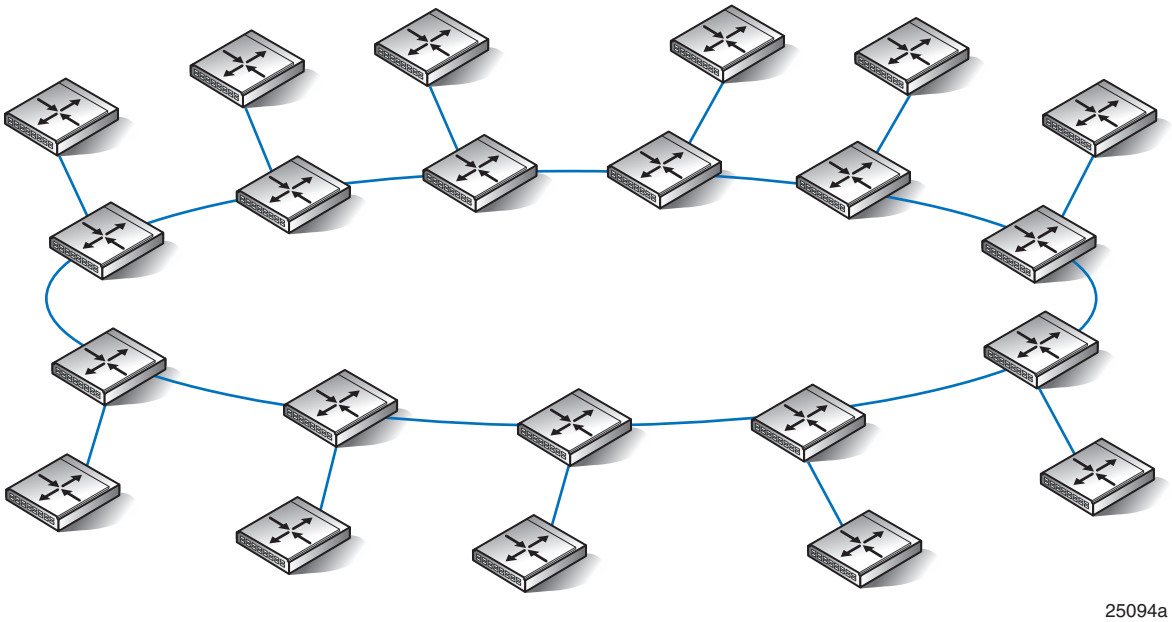



**Figure 87: Highly redundant pairs of gateways, two long distance links**

**Figure 88: Highly redundant pairs of gateways, fully redundant long-distance links** on page 298 is similar to **Figure 83: IP Distance Gateway high-availability configuration with a redundant IP**



https://t.me/learningnets

**network** on page 295, but offers a higher level of redundancy using additional Fibre Channel and LAN connections.

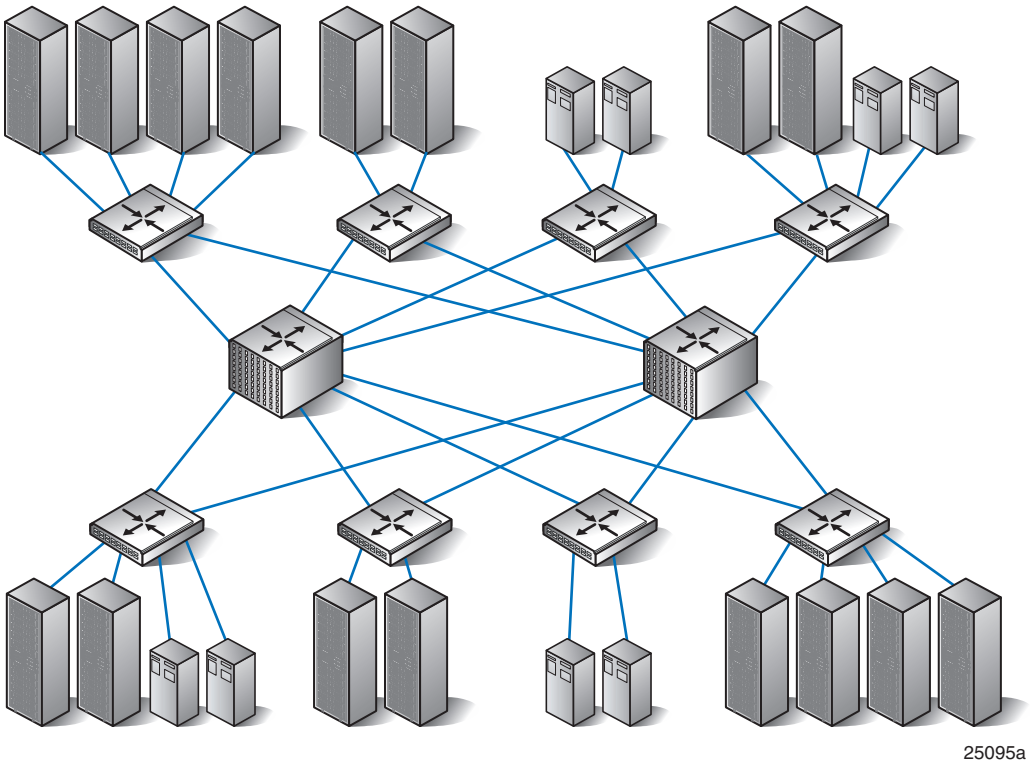



**Figure 88: Highly redundant pairs of gateways, fully redundant long-distance links**

### HPE P6000 Continuous Access 3-site configurations

This section describes P6000 Continuous Access 3-site configurations:

- **Figure 89: HPE P6000 Continuous Access 3-site configuration with four mpx110 gateways** on page 299
- **Figure 90: HPE P6000 Continuous Access 3-site configuration with six mpx110 gateways** on page 300
- **Figure 91: HPE P6000 Continuous Access 3-site configuration with eight gateways** on page 300
- **Figure 92: HPE P6000 Continuous Access 3-site configuration with six mpx110 gateways, full peer-to-peer connectivity** on page 301

The first three configurations (**Figure 89: HPE P6000 Continuous Access 3-site configuration with four mpx110 gateways** on page 299 through **Figure 91: HPE P6000 Continuous Access 3-site configuration with eight gateways** on page 300) provide a fan-in or fan-out relationship between sites. The fourth configuration (**Figure 92: HPE P6000 Continuous Access 3-site configuration with six mpx110 gateways, full peer-to-peer connectivity** on page 301) provides for a peer-to-peer relationship between all sites.

**Figure 89: HPE P6000 Continuous Access 3-site configuration with four mpx110 gateways** on page 299 shows connectivity for three sites using four mpx110 gateways, which implements the minimum-level and lowest-cost connectivity for a 3-site configuration. **Figure 90: HPE P6000 Continuous Access 3-site configuration with six mpx110 gateways** on page 300 shows additional connectivity and redundancy using six mpx110 gateways. **Figure 91: HPE P6000 Continuous Access 3-site configuration with eight gateways** on page 300 shows the highest level of 3-site connectivity using eight mpx110 gateways.



https://t.me/learningnets

The following configuration rules apply to **Figure 89: HPE P6000 Continuous Access 3-site configuration with four mpx110 gateways** on page 299 through **Figure 91: HPE P6000 Continuous Access 3-site configuration with eight gateways** on page 300 (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 92: HPE P6000 Continuous Access 3-site configuration with six mpx110 gateways, full peer-to-peer connectivity** on page 301 is similar to **Figure 90: HPE P6000 Continuous Access 3-site configuration with six mpx110 gateways** on page 300, with additional connectivity to allow for replication between Site 2 and Site 3.

The following configuration rules apply to **Figure 92: HPE P6000 Continuous Access 3-site configuration with six mpx110 gateways, full peer-to-peer connectivity** on page 301 (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 89: HPE P6000 Continuous Access 3-site configuration with four mpx110 gateways** on page 299 shows long-distance link redundancy between all three sites.

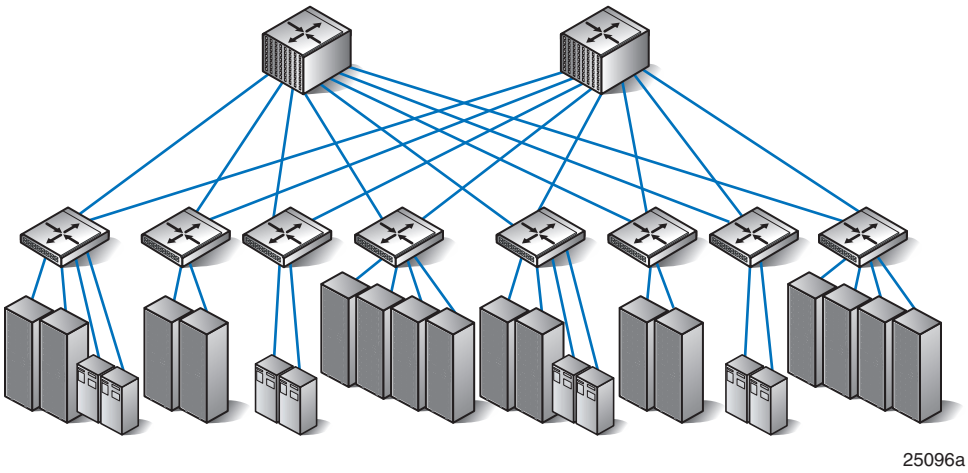



**Figure 89: HPE P6000 Continuous Access 3-site configuration with four mpx110 gateways**

**Figure 90: HPE P6000 Continuous Access 3-site configuration with six mpx110 gateways** on page 300 shows the same long-distance link redundancy as **Figure 89: HPE P6000 Continuous Access 3-site configuration with four mpx110 gateways** on page 299, with the addition of redundant mpx110 gateways at Sites 2 and 3.



https://t.me/learningnets

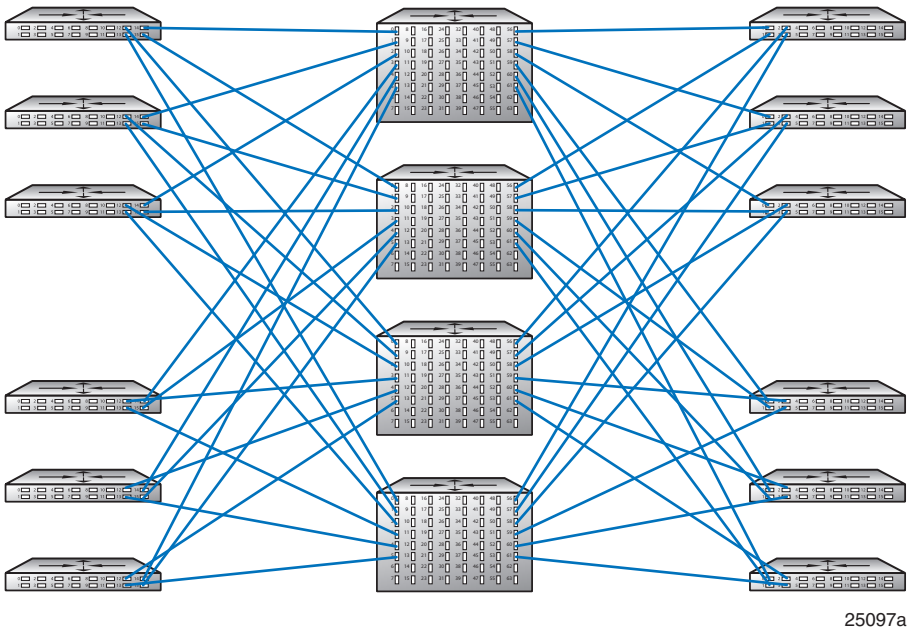



**Figure 90: HPE P6000 Continuous Access 3-site configuration with six mpx110 gateways**

**Figure 91: HPE P6000 Continuous Access 3-site configuration with eight gateways** on page 300 shows the highest level of redundancy, with a dedicated mpx110 pair for all long-distance links to all three sites.

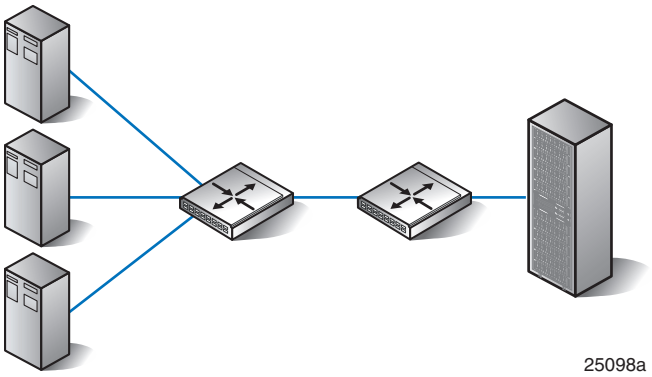



**Figure 91: HPE P6000 Continuous Access 3-site configuration with eight gateways**

**Figure 92: HPE P6000 Continuous Access 3-site configuration with six mpx110 gateways, full peer-to-peer connectivity** on page 301 shows long-distance link redundancy and full connectivity between all three sites.



https://t.me/learningnets

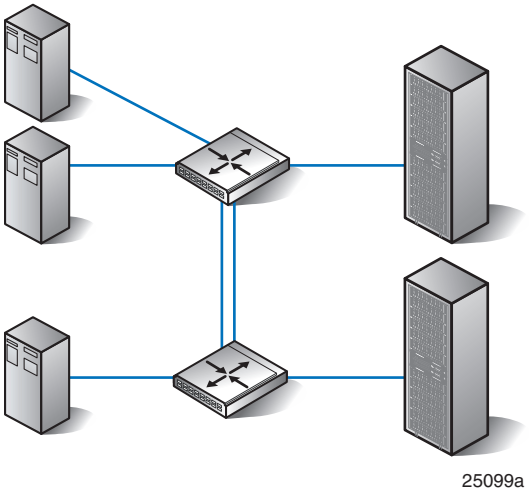



**Figure 92: HPE P6000 Continuous Access 3-site configuration with six mpx110 gateways, full peer-to-peer connectivity**

## 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. HPE does not support FCIP connectivity between other gateway models.
- The mpx110 gateway is supported using FCIP extension with P6000 Continuous Access, P9000 (XP) Continuous Access, and 3PAR Remote Copy. See **EVA storage system rules** on page 301 and **HPE P9000 (XP) Continuous Access** on page 323.
- 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 *HPE StorageWorks IP Distance Gateway User Guide*.

For current storage system support, see the SPOCK website at **http://www.hpe.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 HPE for P6000 Continuous Access, P9000 (XP) Continuous Access, and 3PAR Remote Copy. See IP Distance Gateway (mpx110) features and requirements table and MPX200 Multifunction Router using FCIP features and requirements table.

## EVA storage system rules

The EVA storage system rules follow:

- The mpx110 gateway configured for FCIP is supported for use with P6000 Continuous Access P6350/P6550 EVA using a minimum of XCS 11x, P6300/P6500 EVA using a minimum of XCS 10x.
- The mpx110 gateway is supported for use in all P6000 Continuous Access configurations, including the standard two-fabric, five-fabric, and six-fabric configurations. For more information, see the P6000 Continuous Access documentation.



https://t.me/learningnets

- FCIP is supported for P6000 Continuous Access DR group LUNs and non-DR LUNs.
- Supports the minimum IP bandwidth/maximum DR groups. See Network requirements for long-distance IP gateways and C-series switches with XCS 6.x table through Network requirements for long-distance IP gateways and C-series switches with VCS 4.x table.

### XP storage system rules

The XP storage system rules follow:

- The mpx110 gateway configured for FCIP is supported for use with P9000 (XP) Continuous Access Synchronous, Asynchronous, and Journal. For the latest support updates, see the Hewlett Packard Enterprise storage website: **http://www.hpe.com/info/StoreFabric**.
- Supported XP models are XP24000/20000 and XP12000/10000, using a minimum of 60x and 50x firmware levels respectively. Contact your Hewlett Packard Enterprise-authorized representative for supported firmware versions.
- The mpx110 gateway is supported for use in all Hewlett Packard Enterprise-supported P9000 (XP) Continuous Access FCIP configurations. For more information, see the P9000 (XP) Continuous Access documentation.
- Requires a minimum IP bandwidth of 16 Mb/s per path.
- For additional requirements, see HPE P9000 (XP) Continuous Access FCIP gateway support table.

### XP storage system software

The mpx110 gateway is supported with XP storage software applications, such as P9000 (XP) Continuous Access, Command View XP, P9000 (XP) Continuous Access Journal, Business Copy XP, and XP Array Manager.

### Fibre Channel switch and firmware support

See IP Distance Gateway (mpx110) features and requirements table.

## MPX200 Multifunction Router with FCIP

The 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:

- **HPE Storage Works MPX200 Half Chassis FCIP License** — Includes the license to enable FCIP functionality in one of two bays (slots) in an MPX200 Chassis.
- **HPE 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.



https://t.me/learningnets

**Table 129: MPX200 Multifunction Router using FCIP features and requirements**

| Feature | Requirements |
|---|---|
| Fibre Channel switches for FCIP | B-series switches — See B-series Fibre Channel switches table and B-series Fibre Channel switches and routers table. |
| IP network protocols | TCP/IP IPv4, Ethernet 10 Mb/s, 100 Mb/s, 1,000 Mb/s<br>Requires dedicated IP bandwidth. See Network requirements for long-distance IP gateways with XCS 6.x table and Network requirements for long-distance IP gateways with VCS 4.x table. |
| Storage systems | • P6000 Continuous Access<br>• P9000 (XP) Continuous Access<br>• 3PAR Remote Copy |
| FCIP-supported operating systems | All operating systems supported for P6000 Continuous Access and P9000 (XP) Continuous Access.<br>For P6000 Continuous Access, see the *HPE P6000 Enterprise Virtual Array Compatibility Reference Guide* at **http://www.hpe.com/info/P6000-ContinuousAccess**. |
| Documentation | For information about using the MPX200 Multifunction Router, see **http://www.hpe.com/info/mpx200**. |

## 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 one other license-enabled function—either data migration or FCIP for standalone or concurrent operation. 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 P6000/EVA storage** on page 359. For information about data migration, see **MPX200 Multifunction Router with data migration** on page 243.

**NOTE:** The MPX200 does not support simultaneous FCIP and data migration operation on the same blade (see MPX200 blade configurations table).

The MPX200 Multifunction Router supports the FCIP configurations shown in **Figure 93: MPX200 basic FCIP configuration with one or two long-distance links** on page 304, **Figure 96: MPX200 high-availability configuration with one or two long-distance links** on page 305, and **Figure 97: Local MPX200 basic FCIP configuration with remote IP Distance Gateway (mpx100)** on page 305, in addition to all configurations shown for the IP Distance Gateway, see **IP Distance Gateway configuration examples** on page 294. **Figure 97: Local MPX200 basic FCIP configuration with remote IP Distance Gateway (mpx100)** on page 305 shows an FCIP configuration using an MPX200 at the local site and an IP Distance Gateway at the remote site.

**Figure 93: MPX200 basic FCIP configuration with one or two long-distance links** on page 304 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.



https://t.me/learningnets

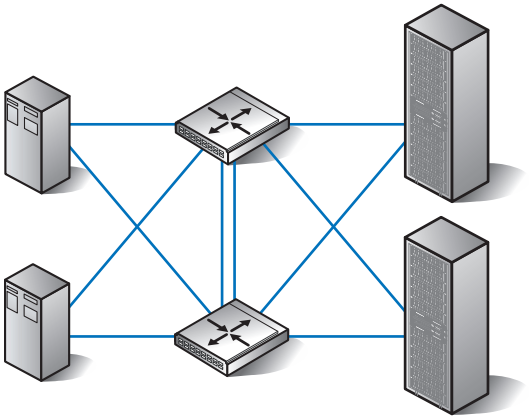



**Figure 93: MPX200 basic FCIP configuration with one or two long-distance links**

**Figure 94: MPX200 FCIP with B-series Integrated Routing** on page 304 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.

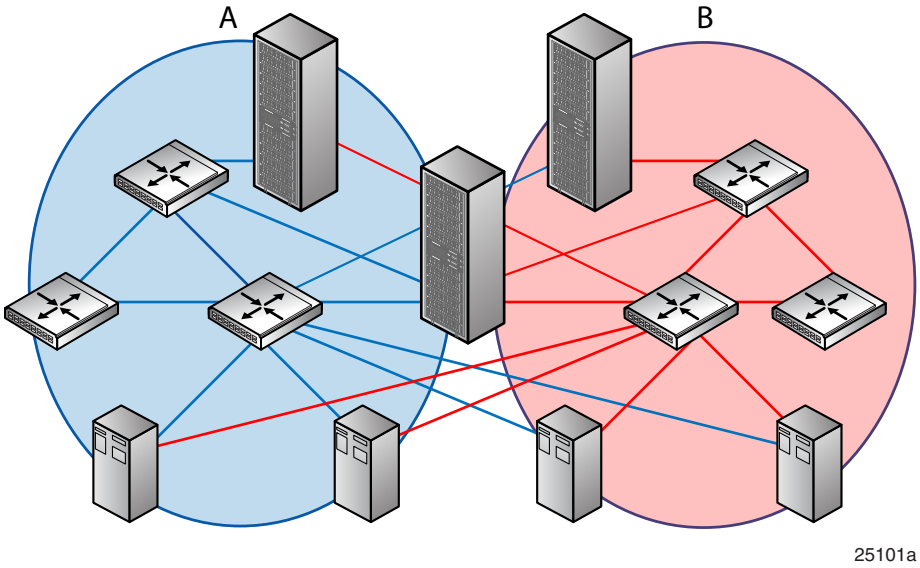



**Figure 94: MPX200 FCIP with B-series Integrated Routing**

**Figure 95: MPX200 FCIP with C-series IVR** on page 305 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.



https://t.me/learningnets

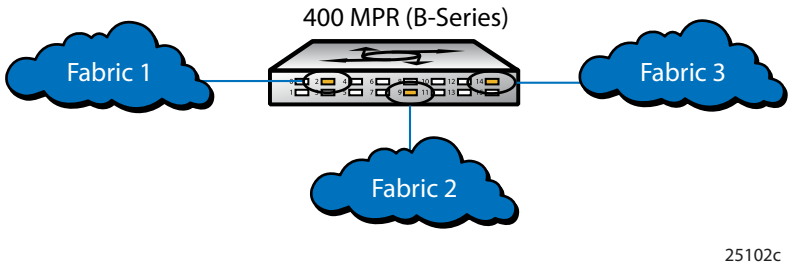



**Figure 95: MPX200 FCIP with C-series IVR**

**Figure 96: MPX200 high-availability configuration with one or two long-distance links** on page 305 shows a high-availability configuration using a dual-blade MPX200 chassis at the local site and the remote site for hardware and path redundancy.

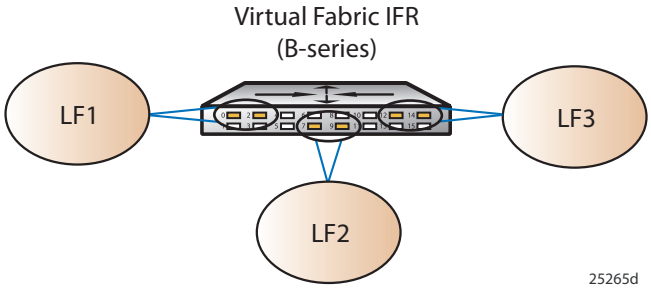



**Figure 96: MPX200 high-availability configuration with one or two long-distance links**

**Figure 97: Local MPX200 basic FCIP configuration with remote IP Distance Gateway (mpx100)** on page 305 shows a basic FCIP configuration with a local single-blade MPX200 chassis and a remote IP Distance Gateway (mpx100).

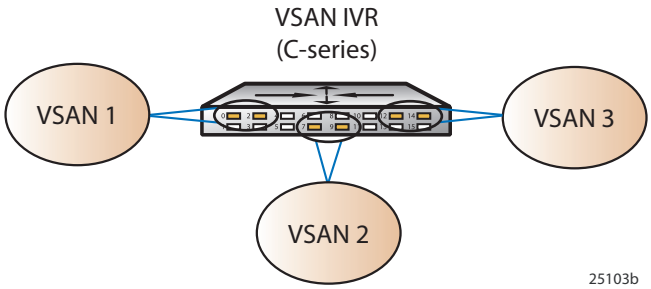



**Figure 97: Local MPX200 basic FCIP configuration with remote IP Distance Gateway (mpx100)**

## 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. HPE does not support FCIP connectivity between other gateway models.
- FCIP is supported on GbE ports only (GE1 and GE2).



https://t.me/learningnets

- The MPX200 is supported using FCIP extension with P6000 Continuous Access, P9000 (XP) Continuous Access, and 3PAR Remote Copy. See **EVA storage system rules** on page 301 and **XP storage system rules** on page 302.
- For performance-tuning information based on the link speed and delay, see the *HPE StorageWorks Multifunction Router User Guide*.

For current storage system support, see the SPOCK website at **http://www.hpe.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 HPE for P6000 Continuous Access, P9000 (XP) Continuous Access, and 3PAR Remote Copy. See IP Distance Gateway (mpx110) features and requirements table and MPX200 Multifunction Router using FCIP features and requirements table.

### Storage system rules

Observe the following storage system rules:

- The MPX200 Multifunction Router configured for FCIP is supported for use with P6000 Continuous Access and P6350/P6550 EVA using a minimum of XCS 11x, P6300/P6500 EVA using a minimum of XCS 10x

  For 3PAR, the MPX200 is supported for use with 3PAR Remote Copy for 3PAR StoreServ 10000 V-Class, 3PAR StoreServ 7000, 3PAR F-Class, and T-Class.

  For XP, the MPX200 is supported for use with P9000 (XP) Continuous Access for XP24000/20000 and XP12000/10000.

- The MPX200 Multifunction Router is supported for use in all supported P6000 Continuous Access SAN configurations, including the standard two-fabric, five-fabric, and six-fabric configurations. For more information, see the P6000 Continuous Access documentation.
- FCIP is supported for P6000 Continuous Access DR group LUNs and non-DR LUNs.
- The minimum IP bandwidth required for P9000 (XP) Continuous Access is 16 Mb/s per path.
- Supports the minimum IP bandwidth/maximum EVA DR groups. See Network requirements for long-distance IP gateways and C-series switches with XCS 6.x table through Network requirements for long-distance IP gateways and C-series switches with VCS 4.x table.

## HPE StoreFabric SN4000B SAN Extension Switch

The StoreFabric SN4000B SAN Extension Switch offers FCIP, Fibre Channel routing, and Fibre Channel switching. You can use all of these functions on the same StoreFabric SN4000B SAN Extension Switch simultaneously.

A StoreFabric SN4000B SAN Extension Switch can have:

- Up to 2 40-GbE IP ports for FCIP, see **SAN extension** on page 274.
- Up to 16 1-GbE/10-GbE IP ports for FCIP, see **SAN extension** on page 274.
- Up to 24 EX_Ports for Fibre Channel routing services, see **SAN fabric topologies** on page 24.
- Up to 24 F_Ports for Fibre Channel switching.



https://t.me/learningnets

For the StoreFabric SN4000B SAN Extension Switch, you can configure a Fibre Channel port as an F_Port, E_Port, D_Port (Diagnosis), M_Port (Mirror), U_Port (self discovery), and configure a GbE port as a VE_Port (FCIP).

**NOTE:** VEX_Ports and EX_Ports cannot connect to the same edge fabric.

VEX_Ports are not supported with the StoreFabric SN4000B SAN Extension Switch.

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.

## HPE StoreFabric SN4000B SAN Extension Switch features and requirements

### Table 130: HPE StoreFabric SN4000B SAN Extension Switch features and requirements

| Feature | Requirements |
|---|---|
| Fibre Channel switches for FCIP | The StoreFabric SN4000B SAN Extension Switch are supported with the B-series switches listed in **B-series switches and fabric rules** on page 105.<br>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, IPv6 Ethernet 10 Mb/s, 100 Mb/s, and 1,000 Mb/s<br>All StoreFabric SN4000B SAN Extension Switch FCIP implementations must have dedicated IP bandwidth configured as a committed bandwidth, or using the Adaptive Rate Limiting feature. HPE does not support the StoreFabric SN4000B SAN Extension Switch configurations with FCIP tunnels configured with nondedicated bandwidth. For more information about configuring and monitoring FCIP extension services, see *Brocade Fabric OS 7.3.x Administrator's Guide*.<br>For more information on dedicated IP bandwidth requirements, see Network requirements for long-distance IP gateways with XCS 6.x table and Network requirements for long-distance IP gateways with VCS 4.x table.<br>**NOTE:**<br>• P9000 (XP) Continuous Access storage systems require a minimum of 16 Mb/s IP bandwidth.<br>• FCIP is supported with IPsec data encryption (license required). This requires a minimum of firmware 7.3.0c.<br>• FCIP is supported with FCIP FastWrite acceleration.<br>• The connection to the StoreFabric SN4000B SAN Extension Switch FCIP port must be GbE compatible (10-GbE for the 10-GbE ports). |

*Table Continued*



https://t.me/learningnets

| Feature | Requirements |
|---|---|
| Storage systems | • P6000 Continuous Access<br>• P9000 (XP) Continuous Access (requires a minimum of firmware 7.3.0c)<br>• 3PAR Remote Copy<br><br>For operating system support, see **Heterogeneous server rules** on page 179.<br><br>FCIP FastWrite is supported with P9000 (XP) Continuous Access and is not supported with P6000 Continuous Access.<br><br>Supported routing modes:<br><br>• P9000 (XP) Continuous Access — Supported with port-based routing (aptpolicy = 1) or exchange-based routing (aptpolicy = 3), see HPE P9000 (XP) Continuous Access FCIP gateway support table.<br>• P6000 Continuous Access — Supported with port-based routing (all XCS versions) or exchange-based routing (XCS 11200000 or later). |
| FCIP-supported operating systems | For P6000 Continuous Access, see the *HPE P6000 Enterprise Virtual Array Compatibility Reference Guide* at **http://www.hpe.com/info/P6000-ContinuousAccess**. |
| Documentation | For information about the StoreFabric SN4000B SAN Extension Switch, see the SAN Infrastructure website **http://www.hpe.com/info/StoreFabric**. |

## HPE StoreFabric SN4000B SAN Extension Switch configuration examples

When connecting fabrics through IP, the StoreFabric SN4000B SAN Extension Switch serves as an FCIP gateway with Fibre Channel routing.

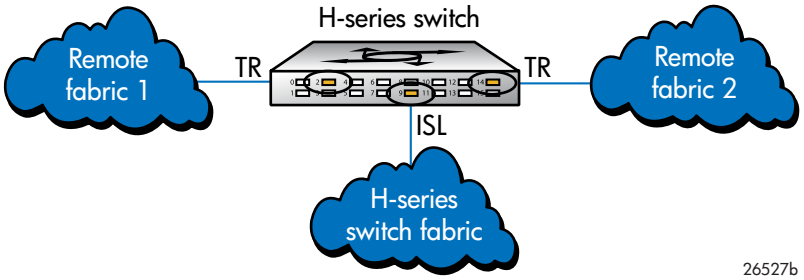



**Figure 98: NSPOF configuration with HPE StoreFabric SN4000B SAN Extension Switches providing Fibre Channel routing and FCIP**



https://t.me/learningnets

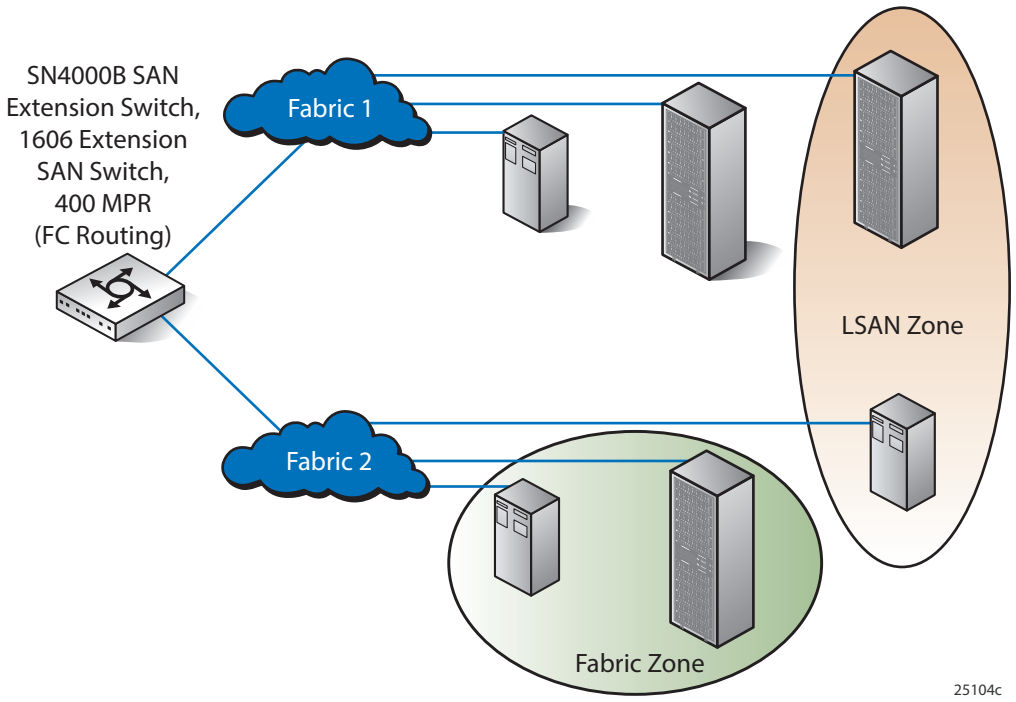



**Figure 99: Fibre Channel routing and FCIP using two HPE StoreFabric SN4000B SAN Extension Switches**

Hewlett Packard Enterprise supports FCIP configurations in which the StoreFabric SN4000B SAN Extension 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 StoreFabric SN4000B SAN Extension Switch.



https://t.me/learningnets

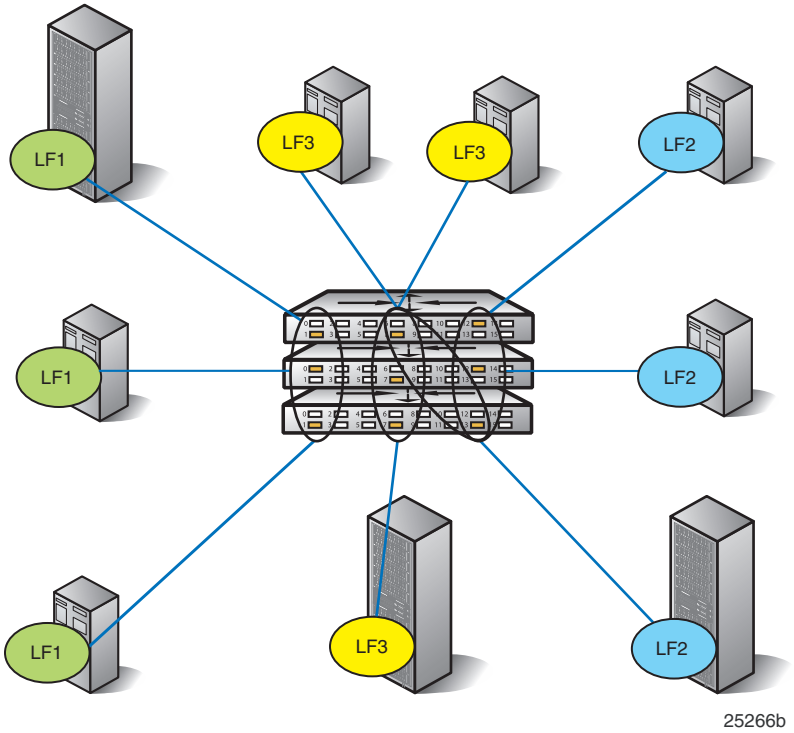



**Figure 100: NSPOF configuration with HPE StoreFabric SN4000B SAN Extension Switches providing FCIP with direct connect devices**

## 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** on page 274
- Up to 16 F_Ports for FC switching.
- Up to 16 EX_Ports for Fibre Channel routing services, see **SAN fabric topologies** on page 24

A DC Dir Switch MP Extension Blade can have:

- Up to 10 1-GbE IP ports for FCIP, see **SAN extension** on page 274
- Up to 2 10-GbE IP ports for FCIP, see **SAN extension** on page 274
- Up to 12 EX_Ports for Fibre Channel routing services, see **SAN fabric topologies** on page 24
- 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).


https://t.me/learningnets

**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 131: 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** on page 105.<br><br>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<br><br>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. HPE 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 *HPE StorageWorks Fabric OS 6.x Administrator's Guide*.<br><br>For more information on dedicated IP bandwidth requirements, see Network requirements for long-distance IP gateways with XCS 6.x table and Network requirements for long-distance IP gateways with VCS 4.x table.<br><br>**NOTE:**<br><br>• P9000 (XP) Continuous Access storage systems require a minimum of 16 Mb/s IP bandwidth.<br>• FCIP is supported with IPsec data encryption (license required). This requires a minimum of firmware 6.3.0.<br>• FCIP is supported with FCIP FastWrite acceleration.<br>• 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). |

*Table Continued*



https://t.me/learningnets

| Feature | Requirements |
|---|---|
| Storage systems | • P6000 Continuous Access<br>• P9000 (XP) Continuous Access (requires a minimum of firmware 6.3.0a)<br>• 3PAR Remote Copy<br><br>For operating system support, see **Heterogeneous server rules** on page 179.<br><br>FCIP FastWrite is supported with P9000 (XP) Continuous Access and is not supported with P6000 Continuous Access.<br><br>Supported routing modes:<br><br>• P9000 (XP) Continuous Access — Supported with port-based routing (aptpolicy = 1) or exchange-based routing (aptpolicy = 3), see HPE P9000 (XP) Continuous Access FCIP gateway support table.<br>• P6000 Continuous Access — Supported with port-based routing (all XCS versions) or exchange-based routing (XCS 09534000 or later). |
| FCIP-supported operating systems | For P6000 Continuous Access, see the *HPE P6000 Enterprise Virtual Array Compatibility Reference Guide* at **http://www.hpe.com/info/P6000-ContinuousAccess**. |
| Documentation | For information about the 1606 Extension SAN Switch and DC Dir Switch MP Extension Blade, see the SAN Infrastructure website **http://www.hpe.com/info/StoreFabric**. |

## 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 101: NSPOF configuration with 1606 Extension SAN Switches providing Fibre Channel routing and FCIP** on page 313) or as a single pair (**Figure 102: Fibre Channel routing and FCIP using two 1606 Extension SAN Switches** on page 313).



https://t.me/learningnets

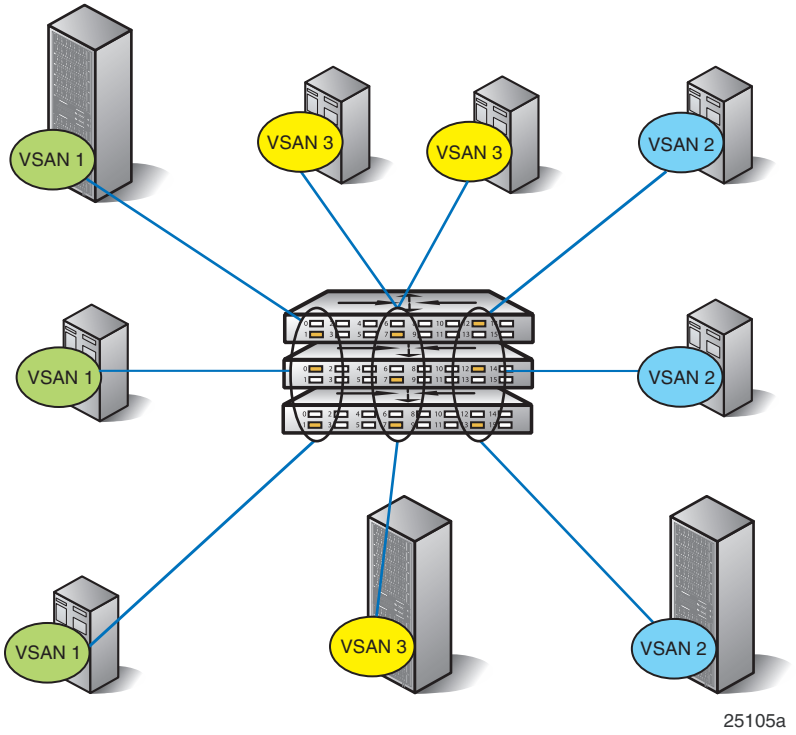



**Figure 101: NSPOF configuration with 1606 Extension SAN Switches providing Fibre Channel routing and FCIP**

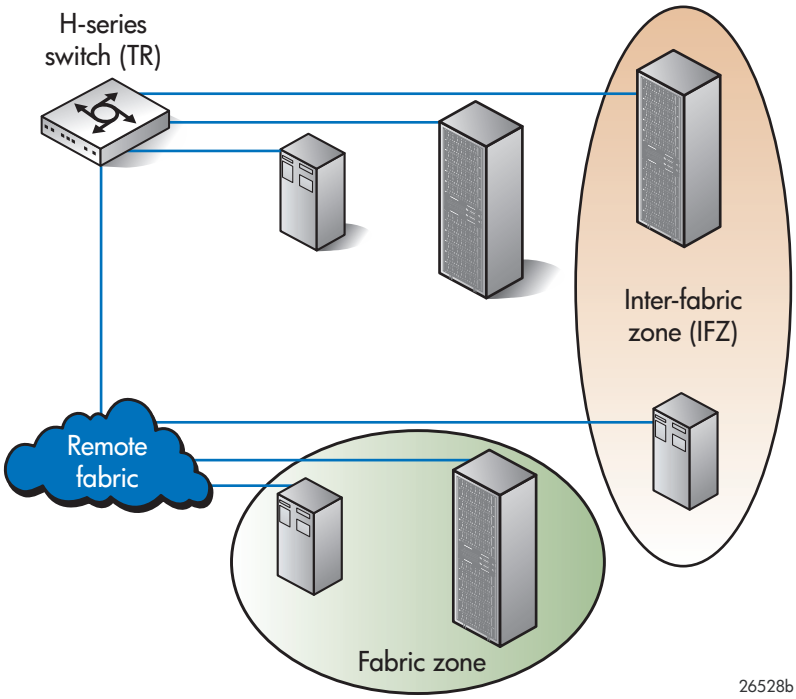



**Figure 102: Fibre Channel routing and FCIP using two 1606 Extension SAN Switches**

Hewlett Packard Enterprise 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 103: NSPOF configuration with direct connect devices** on page 314).



https://t.me/learningnets

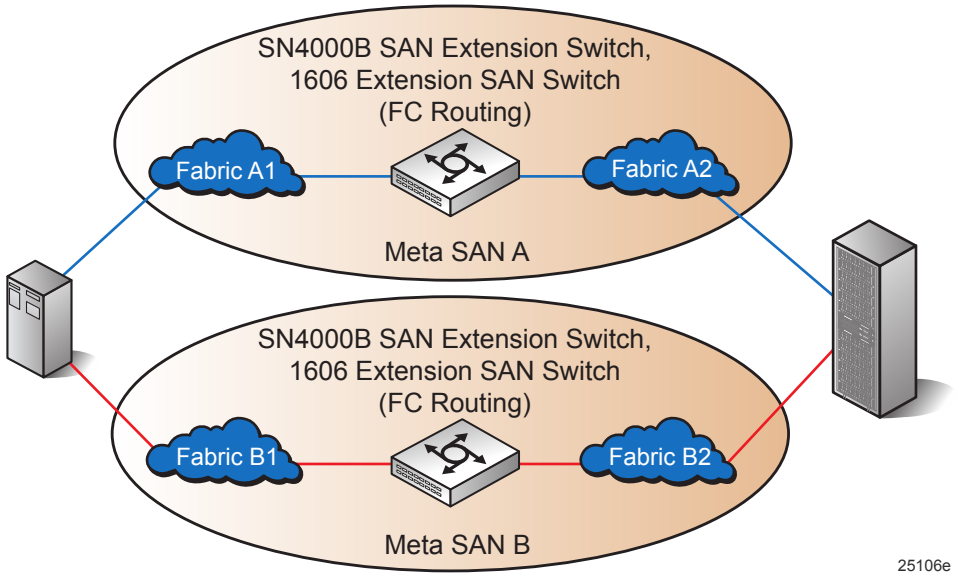



**Figure 103: NSPOF configuration with direct connect devices**

## C-series MDS 9222i, IPS-4, IPS-8, 18/4 Multiservice Modules

The C-series 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 SN8000C, 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 132: C-series MDS module features and requirements

| Feature | Requirements |
|---|---|
| Fibre Channel switch hardware support for iSCSI and FCIP with the IP services modules | C-series MDS 9222i Multiservice Fabric<br>SN8000C 6-Slot Supervisor 2A Director Switch<br>SN8000C 9-Slot Supervisor 2A Director Switch<br>SN8000C 13-Slot Supervisor 2A Fabric 2 Director Switch<br>C-series MDS 9513 Multilayer Director switch<br>C-series MDS 9506 Multilayer Director switch<br>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<br>Requires dedicated IP bandwidth, see Network requirements for long-distance IP gateways with XCS 6.x table and Network requirements for long-distance IP gateways with VCS 4.x table. |

*Table Continued*



https://t.me/learningnets

| Feature | Requirements |
|---|---|
| Storage systems | • P6000 Continuous Access<br>• P9000 (XP) Continuous Access<br>• 3PAR Remote Copy<br><br>Contact Hewlett Packard Enterprise storage representative for EVA and XP supported models.<br><br>For specific operating system support, see **Heterogeneous server rules** on page 179.<br><br>Supported load balance settings:<br><br>• P9000 (XP) Continuous Access — Supported with load balance setting `src-dst-id` or `src-dst-ox-id` on all C-series Fibre Channel switches.<br>• P6000 Continuous Access — Supported with load balance setting `src-dst-id` or `src-dst-ox-id` (XCS 09534000 or later) on all C-series Fibre Channel switches. |
| FCIP-supported operating systems | For P6000 Continuous Access, see the *HPE P6000 Enterprise Virtual Array Compatibility Reference Guide* at **http://www.hpe.com/info/P6000-ContinuousAccess**. |
| Documentation | For information about the C-series MDS modules, see **https://www.cisco.com/c/en/us/products/storage-networking/index.html**. |

# HPE storage replication products

HPE provides the following storage replication products:

- P6000 Continuous Access with the P63**xx**/P65**xx**EVA or EVA4100/4400/6100/6400/8100/8400
- P9000 (XP) Continuous Access with the XP family of storage systems
- OpenVMS host-based volume shadowing

The following products are qualified by HPE as Fibre Channel routers, network gateways, or iSCSI bridges with the EVA storage system:

- 1606 Extension SAN Switch, DC Dir Switch MP Extension Blade, StoreFabric SN4000B SAN Extension Switch, and MP Router Blade, are qualified as FCIP gateways and Fibre Channel routers. The 1606 Extension SAN Switch, DC Dir Switch MP Extension Blade, StoreFabric SN4000B SAN Extension Switch, and MP Router Blade are also qualified as Fibre Channel switches.
- C-series IP Storage Services Modules (IPS-4, IPS-8), and MDS 9222i switches are qualified as an FCIP gateway and an iSCSI bridge. MDS 9222i switches are also qualified as Fibre Channel switches.
- MPX200 Multifunction Router is qualified as an FCIP gateway.
- IP Distance Gateway is qualified as an FCIP gateway.

The following sections describe network requirements for replication products with qualified gateways:



https://t.me/learningnets

- **SAN extension best practices for HPE P6000 Continuous Access** on page 316
- **HPE P6000 Continuous Access with XCS 11x, XCS 10x, or XCS 09x** on page 316
- **HPE P6000 Continuous Access with XCS 6.x** on page 319
- **HPE P6000 Continuous Access with VCS 4.x** on page 321
- **HPE P9000 (XP) Continuous Access** on page 323
- **OpenVMS host-based volume shadowing** on page 339


## SAN extension best practices for HPE P6000 Continuous Access

Hewlett Packard Enterprise recommends you consider the following best practices when implementing SAN extension using FCIP with HPE 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 EVA8**x**00. 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 SCSI Fibre Channel protocol option, see the *HPE P6000 Continuous Access Implementation Guide*.

## HPE P6000 Continuous Access with XCS 11x, XCS 10x, or XCS 09x

This section describes the P6000 Continuous Access with P6350/P6550 EVA XCS 11**x**, P6300/P6500 EVA XCS 10**x** data replication specifications and the supported minimum and maximum transmission rates for qualified switch and gateway pairs.



https://t.me/learningnets

**Table 133: Network requirements for long-distance IP gateways with XCS 11x, XCS 10x, or XCS 09x**

| Specification | Description |
|---|---|
| IP bandwidth | • Must be dedicated to the P6000 Continuous Access storage replication function.<br>• The minimum IP bandwidth required for P6000 Continuous Access with P63xx/P65xx 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.<br>• There is no support for dynamic pacing of the gateway. |
| Maximum number of DR groups | For the maximum number of DR groups, see HPE P6000 Continuous Access heterogeneous SAN configuration rules table.<br>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 Network requirements for long-distance IP gateways and B-series switches with VCS 4.x table and Network requirements for long-distance IP gateways and C-series switches with VCS 4.x table. |
| MTU of the IP network | 1,500 bytes |
| Maximum latency[1] | 100 ms IP network delay one-way or 200 ms round-trip |
| Average packet-loss ratio | Low-loss network: 0.0012% average over 24 hours<br>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 | Must not exceed 10 ms over 24 hours |

[1] Pre-existing restriction

**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.



https://t.me/learningnets

### Table 134: Network requirements for long-distance IP gateways and B-series switches with XCS 11x, XCS 10x, or XCS 09x

| Gateway pair | Minimum supported firmware version | Minimum IP bandwidth[1] and maximum DR groups[2] | |
|---|---|---|---|
| | | Dual IP link maximum latency | Single or shared IP link maximum latency |
| | | 0 to 100 ms one-way | 0 to 100 ms one-way |
| IP Distance Gateway (mpx110) | See note[3] | At least 2 Mb/s for 1 DR group<br>Recommended: At least 5 Mb/s for 1 to 5 DR groups | At least 4 Mb/s for 1 DR group<br>Recommended: At least 10 Mb/s for 1 to 5 DR groups |
| MPX200 Multifunction Router | | | |
| B-series 1606 Extension SAN Switch | See B-series Fibre Channel switches and routers table. | | |
| DC Dir Switch MP Extension Blade | | | |

[1] HPE P6000 Continuous Access requires a minimum of 2 Mb/s of IP bandwidth per path, or 4 Mb/s for two paths.

[2] Assumes single-member DR groups (1 virtual disk). For the maximum number of DR groups supported based on maximum IP bandwidth, see HPE P6000 Continuous Access heterogeneous SAN configuration rules table.

[3] For current support, see the SPOCK website at **http://www.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access.

### Table 135: Network requirements for long-distance IP gateways and C-series switches with XCS11x, XCS 10x, or XCS 09x

| Gateway pair | Minimum supported firmware version | Minimum IP bandwidth[1] and maximum DR groups[2] | |
|---|---|---|---|
| | | Dual IP link maximum latency | Single or shared IP link maximum latency |
| | | 0 to 100 ms one-way | 0 to 100 ms one-way |
| IP Distance Gateway (mpx110) | See note[3] | At least 2 Mb/s for 1 DR group<br>Recommended: At least 5 Mb/s for 1 to 5 DR groups | At least 4 Mb/s for 1 DR group<br>Recommended: At least 10 Mb/s for 1 to 5 DR groups |
| MPX200 Multifunction Router | | | |
| C-series IPS-8, 18/4, MDS 9222i | | | |

[1] HPE P6000 Continuous Access requires a minimum of 2 Mb/s of IP bandwidth per path, or 4 Mb/s for two paths.



https://t.me/learningnets

[2] Assumes single-member DR groups (1 virtual disk). For the maximum number of DR groups supported based on maximum IP bandwidth, see HPE P6000 Continuous Access heterogeneous SAN configuration rules table.

[3] For current support, see the SPOCK website at **http://www.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access.

## HPE P6000 Continuous Access with XCS 6.x

This section describes the P6000 Continuous Access XCS 6.**x** data replication specifications and the supported minimum and maximum transmission rates for qualified switch and gateway pairs.

### Table 136: Network requirements for long-distance IP gateways with XCS 6.x

| Specification | Description |
|---|---|
| IP bandwidth | • Must be dedicated to the P6000 Continuous Access storage replication function.<br>• The minimum IP bandwidth required for P6000 Continuous Access with FCIP is 2 Mb/s per path, or 4 Mb/s for two paths when using one IP link.<br>• There is no support for dynamic pacing of the gateway. |
| Maximum number of DR groups | For the maximum number of DR groups, see HPE P6000 Continuous Access heterogeneous SAN configuration rules table.<br>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 Network requirements for long-distance IP gateways and B-series switches with XCS 6.x table and Network requirements for long-distance IP gateways and C-series switches with XCS 6.x table. |
| MTU of the IP network | 1,500 bytes |
| Maximum latency | 100 ms IP network delay one-way or 200 ms round-trip |
| Average packet-loss ratio | Low-loss network: 0.0012% average over 24 hours<br>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 | Must not exceed 10 ms over 24 hours |

**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.



https://t.me/learningnets

### Table 137: Network requirements for long-distance IP gateways and B-series switches with XCS 6.x

| Gateway pair | Minimum supported firmware version | Minimum IP bandwidth[1] and maximum DR groups[2] | |
|---|---|---|---|
| | | Dual IP link maximum latency | Single or shared IP link maximum latency |
| | | 0 to 100 ms one-way | 0 to 100 ms one-way |
| IP Distance Gateway (mpx110) | See note[3] | At least 2 Mb/s for 1 DR group<br>Recommended: At least 5 Mb/s for 1 to 5 DR groups | At least 4 Mb/s for 1 DR group<br>Recommended: At least 10 Mb/s for 1 to 5 DR groups |
| MPX200 Multifunction Router | | | |
| B–series 1606 Extension SAN Switch<br>DC Dir Switch MP Extension Blade<br>StoreFabric SN4000B SAN Extension Switch | See B-series Fibre Channel switches and routers table. | | |

1 HPE P6000 Continuous Access requires a minimum of 2 Mb/s of IP bandwidth per path, or 4 Mb/s for two paths.
2 Assumes single-member DR groups (1 virtual disk). For the maximum number of DR groups supported based on maximum IP bandwidth, see HPE P6000 Continuous Access heterogeneous SAN configuration rules table.
3 For current support, see the SPOCK website at **http://www.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access.

### Table 138: Network requirements for long-distance IP gateways and C-series switches with XCS 6.x

| Gateway pair | Minimum supported firmware version | Minimum IP bandwidth[1] and maximum DR groups[2] | |
|---|---|---|---|
| | | Dual IP link maximum latency | Single or shared IP link maximum latency |
| | | 0 to 100 ms one-way | 0 to 100 ms one-way |
| IP Distance Gateway (mpx110) | See note[3] | At least 2 Mb/s for 1 DR group<br>Recommended: At least 5 Mb/s for 1 to 5 DR groups | At least 4 Mb/s for 1 DR group<br>Recommended: At least 10 Mb/s for 1 to 5 DR groups |
| MPX200 Multifunction Router | | | |
| C-series IPS-8, 18/4,<br>MDS 9222i | | | |

*Table Continued*


https://t.me/learningnets

| Gateway pair | Minimum supported firmware version | Minimum IP bandwidth[1] and maximum DR groups[2] | |
|---|---|---|---|
| | | Dual IP link maximum latency | Single or shared IP link maximum latency |
| | | 0 to 100 ms one-way | 0 to 100 ms one-way |
| **Legacy IP gateways** | | | |
| C-series IPS-4 | See note[3]. | At least 2 Mb/s for 1 DR group<br>Recommended: At least 5 Mb/s for 1 to 5 DR groups | At least 4 Mb/s for 1 DR group<br>Recommended: At least 10 Mb/s for 1 to 5 DR groups |

1 HPE P6000 Continuous Access requires a minimum of 2 Mb/s of IP bandwidth per path, or 4 Mb/s for two paths.

2 Assumes single-member DR groups (1 virtual disk). For the maximum number of DR groups supported based on maximum IP bandwidth, see HPE P6000 Continuous Access heterogeneous SAN configuration rules table.

3 For current support, see the SPOCK website at **http://www.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access.

## HPE 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 139: Network requirements for long-distance IP gateways with VCS 4.x

| Specification | Description |
|---|---|
| IP bandwidth | • Must be dedicated to the P6000 Continuous Access storage replication function.<br>• The minimum IP bandwidth required for P6000 Continuous Access with FCIP is 2 Mb/s per path, or 4 Mb/s for two paths when using one IP link.<br>• There is no support for dynamic pacing of the gateway. |
| Maximum number of DR groups | For the maximum number of DR groups, see HPE P6000 Continuous Access heterogeneous SAN configuration rules table.<br>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 Network requirements for long-distance IP gateways and B-series switches with VCS 4.x table and Network requirements for long-distance IP gateways and C-series switches with VCS 4.x table. |
| MTU of the IP network | 1,500 bytes |
| Maximum latency[1] | 100 ms one-way or 200 ms round-trip |

*Table Continued*



https://t.me/learningnets

| Specification | Description |
|---|---|
| Average packet-loss ratio | Low-loss network: 0.0012% average over 24 hours<br>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 | Must not exceed 10 ms over 24 hours |

[1] Pre-existing restriction

**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 140: Network requirements for long-distance IP gateways and B-series switches with VCS 4.x

| Gateway pair | Minimum supported firmware version | Minimum IP bandwidth[1] and maximum DR groups[2] | |
|---|---|---|---|
| | | Dual IP link maximum latency | Single or shared IP link maximum latency |
| | | 0 to 100 ms one-way | 0 to 100 ms one-way |
| IP Distance Gateway (mpx110) | See note[3] | At least 2 Mb/s for 1 DR group<br>Recommended: At least 5 Mb/s for 1 to 5 DR groups | At least 4 Mb/s for 1 DR group<br>Recommended: At least 10 Mb/s for 1 to 5 DR groups |
| MPX200 Multifunction Router | | | |
| B-series 1606 Extension SAN Switch<br>DC Dir Switch MP Extension Blade<br>StoreFabric SN4000B SAN Extension Switch | See B-series Fibre Channel switches and routers table. | | |

[1] HPE P6000 Continuous Access requires a minimum of 2 Mb/s of IP bandwidth per path, or 4 Mb/s for two paths.
[2] Assumes single-member DR groups (1 virtual disk). For the maximum number of DR groups supported based on maximum IP bandwidth, see HPE P6000 Continuous Access heterogeneous SAN configuration rules table.
[3] For current support, see the SPOCK website at **http://www.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access.


https://t.me/learningnets

### Table 141: Network requirements for long-distance IP gateways and C-series switches with VCS 4.x

| Gateway pair | Minimum supported firmware version | Minimum IP bandwidth[1] and maximum DR groups[2] | |
|---|---|---|---|
| | | Dual IP link maximum latency | Single or shared IP link maximum latency |
| | | 0 to 100 ms one-way | 0 to 100 ms one-way |
| IP Distance Gateway (mpx110) | See note[3] | At least 2 Mb/s for 1 DR group | At least 2 Mb/s for 1 DR group |
| MPX200 Multifunction Router | | Recommended: At least 5 Mb/s for 1 to 5 DR groups | Recommended: At least 5 Mb/s for 1 to 5 DR groups |
| C-series IPS-8, 18/4, MDS 9222i | | | |
| **Legacy IP gateways** | | | |
| C-series IPS-4 | See note[3]. | At least 2 Mb/s for 1 DR group | At least 4 Mb/s for 1 DR group |
| | | Recommended: At least 5 Mb/s for 1 to 5 DR groups | Recommended: At least 10 Mb/s for 1 to 5 DR groups |

1 HPE P6000 Continuous Access requires a minimum of 2 Mb/s of IP bandwidth per path, or 4 Mb/s for two paths.
2 Assumes single-member DR groups (1 virtual disk). For the maximum number of DR groups supported based on maximum IP bandwidth, see HPE P6000 Continuous Access heterogeneous SAN configuration rules table.
3 For current support, see the SPOCK website at **http://www.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access.

## HPE P9000 (XP) Continuous Access

P9000 (XP) Continuous Access is supported on all XP storage systems. It supports three data replication modes: Synchronous, Asynchronous, and Journal.

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 142: HPE P9000 (XP) Continuous Access replication modes

| Storage system | Synchronous | Asynchronous[1] | Journal[2] | Recommended firmware versions | Minimum firmware versions |
|---|---|---|---|---|---|
| XP7 | . | - | . | 80.05.42.00/00 | 80 - 04 - 01 - 00/00 |
| P9500[3] | . | - | . | 70.06.46.00/00 | V01+1 |

*Table Continued*



https://t.me/learningnets

| Storage system | Synchronous | Asynchronous[1] | Journal[2] | Recommended firmware versions | Minimum firmware versions |
|---|---|---|---|---|---|
| XP24000<br>XP20000 | . | . | . | 60.08.62.00/00 | 60.01.68.00/00 |
| XP12000<br>XP10000 | . | . | . | 50.09.86.00/00 | 50.05.46.00/00 |

1 XP24000/20000/12000/10000: 32K pairs
2 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)
3 64K pairs

**Legend:** • = supported; — = not supported

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 143: HPE P9000 (XP) Continuous Access protocols and source-target replication pairing

| Storage system | Protocol | XP7 | P9500 | XP24000XP20000 | XP12000XP10000 |
|---|---|---|---|---|---|
| **P9500** | **Fibre Channel** | . | . | . | . |
| **XP7** | **Fibre Channel** | . | . | . | . |
| **XP24000** | **Fibre Channel** | . | . | . | . |
| **XP12000** | **Fibre Channel** | . | . | . | . |

**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 HPE P9000 (XP) Continuous Access replication modes.


https://t.me/learningnets

#### HPE P9000 (XP) Continuous Access Synchronous

### Table 144: HPE P9000 (XP) Continuous Access Synchronous replication rules

| Rule number | Description |
|---|---|
| 1 | Synchronous replication is supported on XP7, P9500. and XP24000/20000 storage systems. The maximum distance or latency supported for synchronous replication is 200 km when using dark fiber, or less than 5 ms using other extension methods. Distances greater than 200 km or latencies greater than 5 ms require HPE approval before implementation. Contact Hewlett Packard Enterprise product support if using greater distances or latencies.<br><br>**NOTE:** You must ensure that host-based applications utilizing replicated storage are able to tolerate the total of the following latencies:<br><br>• Local I/O servicing (local site storage system)<br>• Continuous Access Synchronous replication<br>• Remote I/O servicing (remote site storage system) |
| 2 | A minimum of 16 Mb/s of IP bandwidth per path is required. |

#### HPE P9000 (XP) Continuous Access Asynchronous

### Table 145: HPE P9000 (XP) Continuous Access Asynchronous replication rules

| Rule number | Description |
|---|---|
| 1 | Asynchronous replication is supported on XP24000/20000 and XP12000/10000 storage systems. The maximum latency supported for Asynchronous replication is 300 ms one-way or 600 ms round-trip.<br><br>**NOTE:** P9500 storage systems do not support Continuous Access Asynchronous replication. |
| 2 | A minimum of 16 Mb/s of IP bandwidth per path is required. |

#### HPE P9000 (XP) Continuous Access Journal

### Table 146: HPE P9000 (XP) Continuous Access Journal replication rules

| Rule number | Description |
|---|---|
| 1 | Journal replication is supported on and XP7, P9500, 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 **HPE P9000 (XP) Continuous Access Journal configuration**. |

*Table Continued*



https://t.me/learningnets

| Rule number | Description |
|---|---|
| 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 **HPE P9000 (XP) Continuous Access P-jnl and S-jnl groups**. |
| 7 | Hewlett Packard Enterprise 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:<br><br>Two topologies are supported; cascaded (1:1:1) and multi-target (1:2), see **HPE P9000 (XP) Continuous Access 3DC configuration**.<br><br>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.<br><br>Journal capacity must be capable of buffering data until off-peak times.<br><br>• For 256 Mb/s or higher, set the line speed to 256.<br>• For 100 to 256 Mb/s, set the line speed to 100.<br>• For 10 to 100 Mb/s, set the line speed to 10.<br><br>**NOTE:** A minimum of 16Mb/s of IP bandwidth per path is required. |

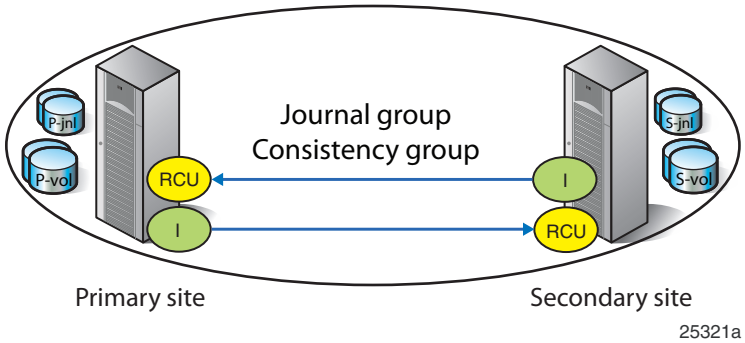



**Figure 104: HPE P9000 (XP) Continuous Access Journal configuration**



https://t.me/learningnets

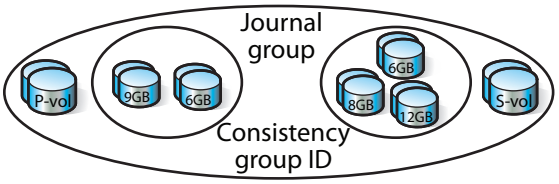



**Figure 105: HPE P9000 (XP) Continuous Access P-jnl and S-jnl groups**

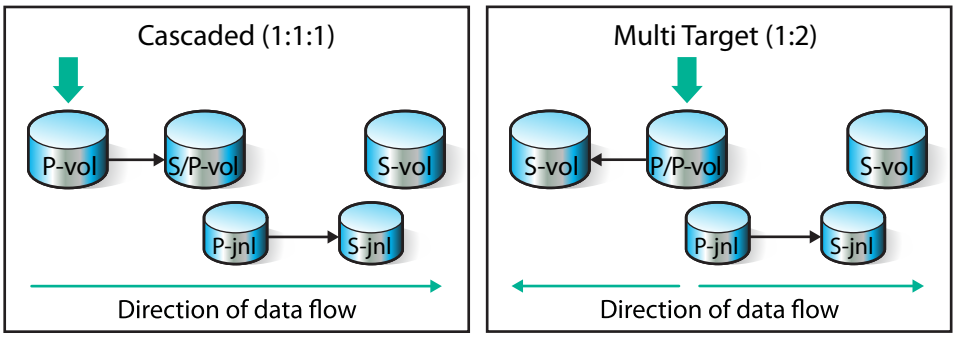



**Figure 106: HPE P9000 (XP) Continuous Access 3DC configuration**

### HPE P9000 (XP) Continuous Access configuration support

This section describes the products and maximum distances supported for the following 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



https://t.me/learningnets

**Direct storage-to-storage**

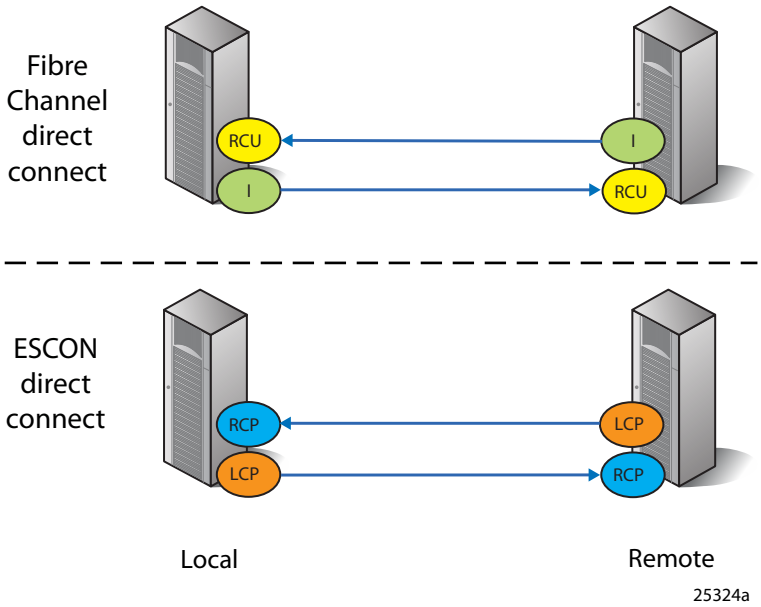



**Figure 107: HPE P9000 (XP) Continuous Access direct storage-to-storage configurations for Fibre Channel and ESCON**

## Table 147: HPE P9000 (XP) Continuous Access direct storage-to-storage distances

| Configuration | Maximum distance |
|---|---|
| Fibre Channel (Fibre CHIP) | 150 m at 4 Gb/s |
| Port parameters for direct connect: | 300 m at 2 Gb/s |
| Fabric = Off, Connection = FC-AL | 500 m at 1 Gb/s |
| | Long-wave SFPs: 10 km at 4 Gb/s |
| ESCON | 3 km |


https://t.me/learningnets

**Fibre Channel switches**

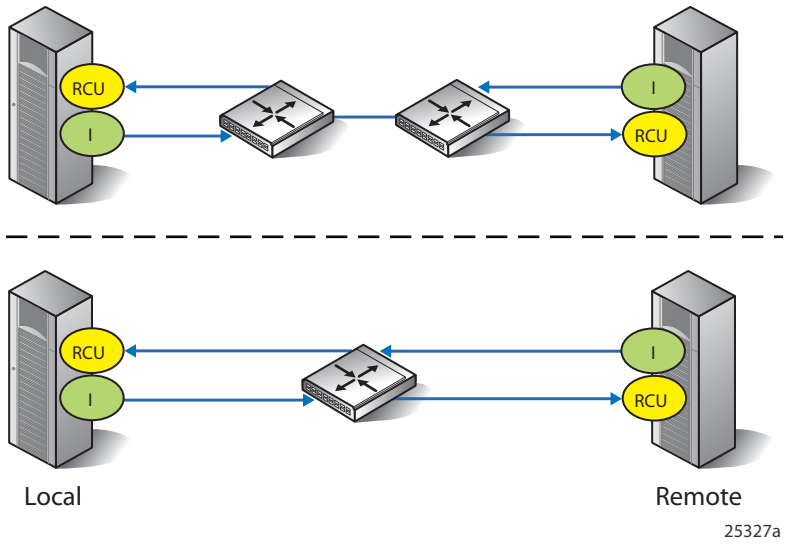



**Figure 108: HPE P9000 (XP) Continuous Access single-switch and multi-switch Fibre Channel configurations**

## Table 148: HPE P9000 (XP) Continuous Access Fibre Channel distances

| Configuration | Maximum distance |
|---|---|
| Single-switch or multi-switch, See **HPE P9000 (XP) Continuous Access single-switch and multi-switch Fibre Channel configurations**.<br><br>**NOTE:** For Fibre Channel switch model support, contact a Hewlett Packard Enterprise storage representative.<br><br>For B-series switches and routers:<br><br>Supports port-based routing (aptpolicy = 1) and exchange-based routing (aptpolicy = 3) see Table HPE P9000 (XP) Continuous Access FCIP gateway support.<br><br>Port parameters for switch/fabric connect:<br><br>• Fabric = On, Connection = Point to Point (recommended setting)<br>• Fabric = On, Connection = FC-AL<br><br>**NOTE:** For all switch configurations, P9000 (XP) Continuous Access ports and host ports must be in separate zones. | See Table 2 Gb/s Fibre Channel fiber optic cable loss budgets, 1 Gb/s Fibre Channel fiber optic cable loss budgets (nominal bandwidth), and Fibre Channel distance rules for 4 Gb/s switch models (B-series and C-series switches.<br><br>**NOTE:** For WDM distances, see Table HPE P9000 (XP) Continuous Access WDM distances and equipment. |



https://t.me/learningnets

**ESCON directors and repeaters**

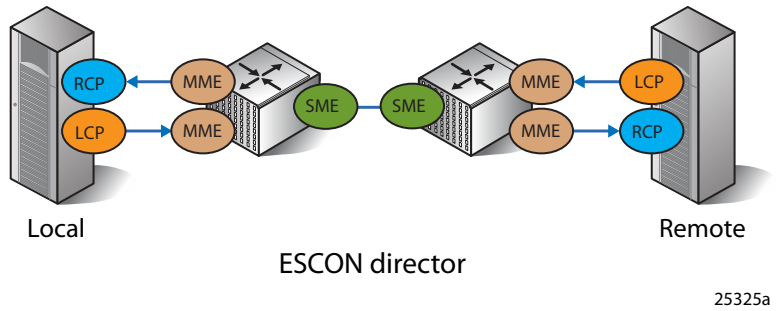



**Figure 109: HPE P9000 (XP) Continuous Access ESCON director configuration**

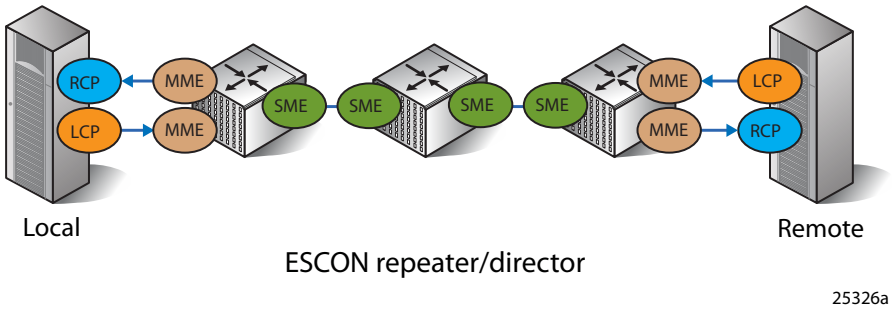



**Figure 110: HPE P9000 (XP) Continuous Access ESCON repeater configuration**

## Table 149: HPE 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 |
| **NOTE:** Supported only in a director-to-director configuration. | **NOTE:** The ESCON protocol is not converted in this process. |



https://t.me/learningnets

**WDM**

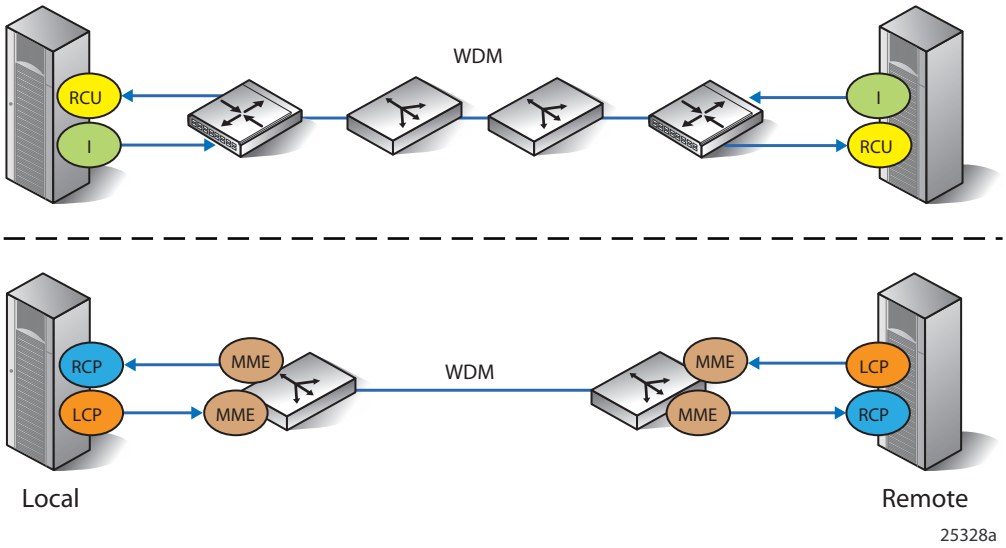



**Figure 111: HPE P9000 (XP) Continuous Access WDM configurations**

## Table 150: HPE P9000 (XP) Continuous Access WDM distances and equipment

| Configuration | Maximum distance and equipment |
|---|---|
| Fibre Channel | For supported distances, see Table 2 Gb/s Fibre Channel fiber optic cable loss budgets, 1 Gb/s Fibre Channel fiber optic cable loss budgets (nominal bandwidth), and Fibre Channel distance rules for 4 Gb/s switch models (B-series and C-series switches).<br>**NOTE:** For Synchronous replication, distance impacts performance. Hewlett Packard Enterprise recommends a maximum distance of 200 km when using dark fiber.<br>P9000 (XP) Continuous Access is supported with all WDM products supported by the switch vendors, see **Certified third-party WDM products**. |
| ESCON | 50 km<br>Nortel Optera Metro 5200/5100<br>Movaz RAYexpress |



https://t.me/learningnets

### ATM and SONET/SDH

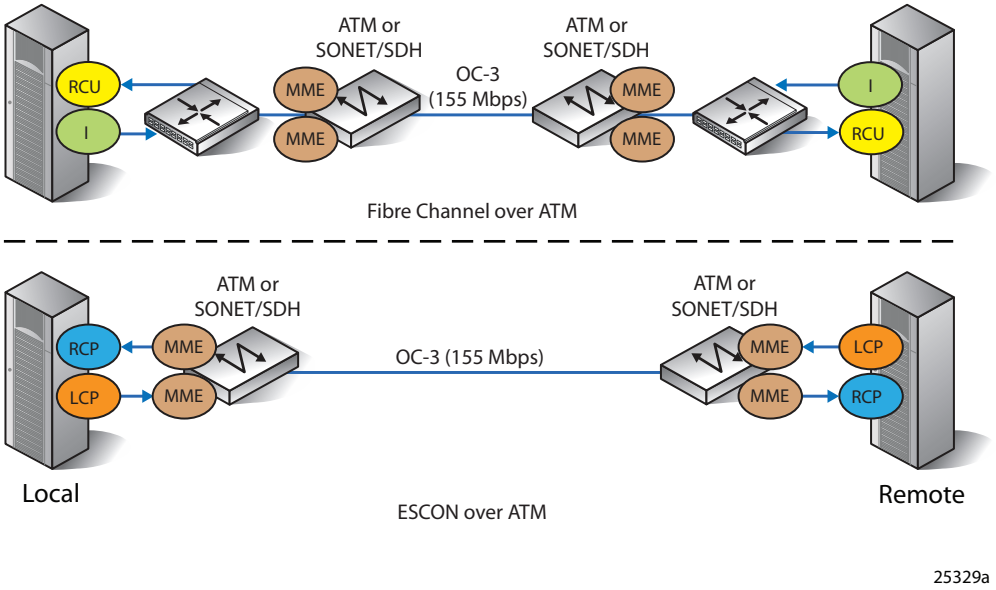



**Figure 112: HPE P9000 (XP) Continuous Access ATM and SONET/SDH configurations**

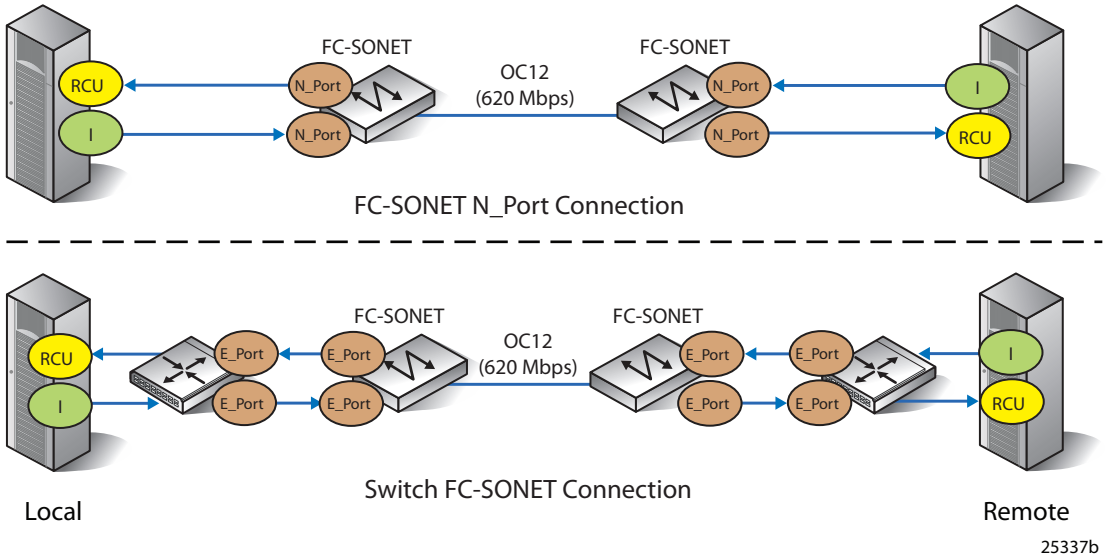



**Figure 113: HPE P9000 (XP) Continuous Access SONET/SDH FC direct and switch configurations**



https://t.me/learningnets

### Table 151: HPE P9000 (XP) Continuous Access ATM and SONET/SDH products (Cisco FCMR-4 only)

| Configuration | XP storage system | Supported bandwidths and products |
|---|---|---|
| ESCON (ATM) | XP24000<br>XP20000 | OC-3-Brocade-CNT Ultranet Storage Director (USD) firmware 2.7 or 3.2.1 |
| ESCON (ATM) | XP12000<br>XP10000 | OC-3-Brocade-CNT-Inrange 9801 SNS firmware 2.3 (Build 27) or 2.4 (Build 13)<br>Brocade-CNT-Inrange 9811H FW ACP-3; 3.1.1702<br>**Requirements:**<br>• Buffer size = 64 credits/port<br>• 1 x ATM OC-3 (155 Mb/s) port<br>• Interoperability mode = 1<br>• E_D_TOV = 5,000<br>**NOTE:** This converter has been discontinued. There is limited support for existing installations. |
| Fibre Channel | XP12000<br>XP10000 | OC-3-Brocade-CNT (Inrange) 9801H or L, 2.3 (Build 28) or 2.4 (Build 13)<br>**Requirements:**<br>• Buffer size = 64 credits/port<br>• 1 x ATM OC-3 (155 Mb/s) port<br>• Interoperability mode = 1<br>• E_D_TOV = 5,000<br>**NOTE:** This converter has been discontinued. There is limited support for existing installations. |
| Fibre Channel | For a list of supported storage systems, contact your Hewlett Packard Enterprise representative. | Cisco FCMR-4 (800-22030-03) for FC over SONET/SDH<br>**Requirements:**<br>• Firmware 57-5391-03, hardware revision 03<br>• FC/FICON (1 Gb/s or 2 Gb/s) over SONET/SDH (OC-3,OC-12, OC-48, or OC-192)<br>• 255 buffer credits (ingress) and 1,200 buffer credits (egress)<br>• ML-Series, CE-Series, G-Series, or E-Series |

*Table Continued*



https://t.me/learningnets

| Configuration | XP storage system | Supported bandwidths and products |
|---|---|---|
| Fibre Channel | XP24000<br>XP20000<br>XP12000 | Ciena CN 2000 for FC over SONET/SDH or DWDM, see **HPE P9000 (XP) Continuous Access SONET/SDH FC direct and switch configurations**<br>**Requirements:**<br>• Firmware 5.1.0, 5.0.1<br>• SONET/SDH OC-3<br>• Port parameters:<br>◦ For direct connect, Fabric = Off, Connection = FC-AL<br>◦ For switch/fabric connect (B-series switches), Fabric = On, Connection = Point to Point |

| Distances |
|---|
| • Distance is unlimited using converters and Asynchronous replication.<br>• Converter and Synchronous replication have practical distance limitations.<br>• Supported distance for an ATM connection depends on network latency delays, packet loss, and application performance requirements. |

### FCIP and routing extension

This section describes the P9000 (XP) Continuous Access data replication IP specifications and the supported minimum and maximum transmission rates for qualified switch and IP gateway pairs.

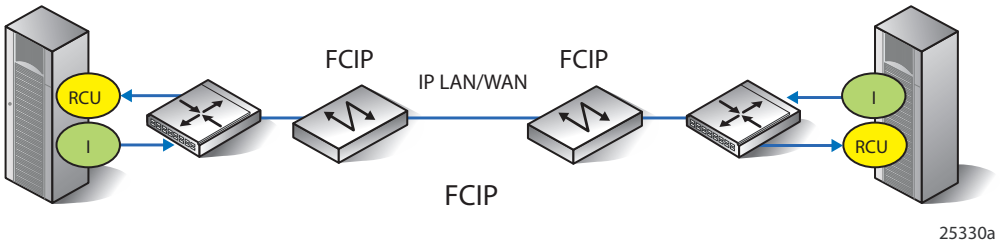



**Figure 114: HPE P9000 (XP) Continuous Access FCIP configuration**



https://t.me/learningnets

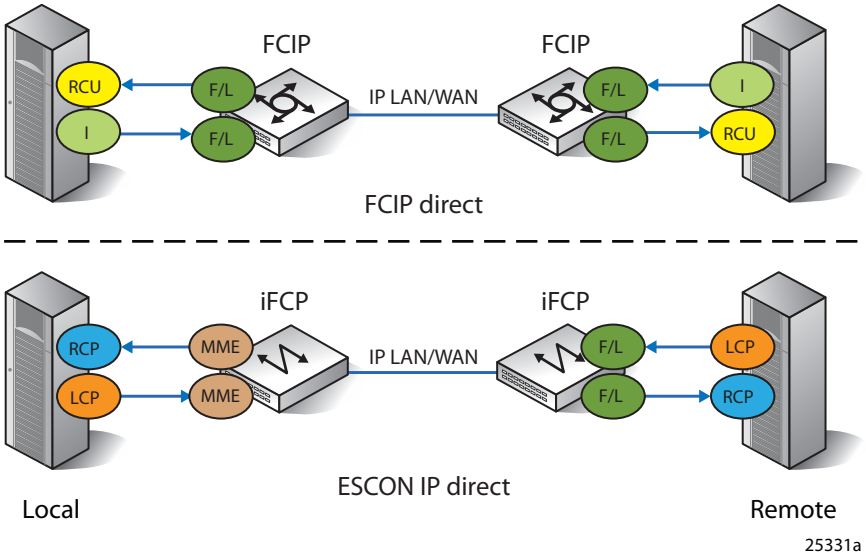



**Figure 115: HPE P9000 (XP) Continuous Access FCIP and iFCP configurations**

## Table 152: Network requirements for long-distance IP gateways with HPE P9000 (XP) Continuous Access

| Specification | Requirement |
|---|---|
| Bandwidth[1] | • Must be dedicated to the P9000 (XP) Continuous Access storage replication function.<br>• The minimum IP bandwidth required for P9000 (XP) Continuous Access is 16 Mb/s per path.<br>• 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.<br>• There is no support for dynamic pacing of the gateway. |
| MTU of the IP network | 1,500 bytes: StorageWorks IP Distance Gateway (mpx100)<br>3,000 bytes: C-series MDS 9222i, IPS-8, IPS-4<br>**NOTE:** 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. |
| Maximum latency[1] | See Tables HPE P9000 (XP) Continuous Access Synchronous replication rules, HPE P9000 (XP) Continuous Access Asynchronous replication rules, and HPE P9000 (XP) Continuous Access Journal replication rules. |

*Table Continued*



https://t.me/learningnets

| Specification | Requirement |
|---|---|
| Average packet-loss ratio[2] | Low-loss network: 0.0012% average over 24 hours<br>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[3] | Must not exceed 10 ms over 24 hours |

1 Pre-existing restriction

2 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 HPE P9000 (XP) Continuous Access FCIP gateway support are supported in both low-loss and high-loss networks.

3 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 *HPE StorageWorks XP Continuous Access User Guide*.

### Table 153: HPE P9000 (XP) Continuous Access FCIP gateway support

| Product and minimum supported firmware version | XP storage system | Notes |
|---|---|---|
| IP Distance Gateway (mpx110) [1] | P9500<br>XP24000<br>XP20000<br>XP12000<br>XP10000<br>XP7 | Requires a B-series or C-series switch between the XP storage systems and the gateways with the following settings and packet-loss criteria:<br>• TCP window size = 5 or 6<br>• Average packet-loss ratio maximum = 0.1%<br>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 *HPE StorageWorks IP Distance Gateway User Guide*. |
| MPX200 Multifunction Router with FCIP[1] | P9500<br>XP24000<br>XP20000<br>XP12000<br>XP10000<br>XP7 | Requires a B-series or C-series switch between the XP storage systems and the gateways with the following settings and packet-loss criteria:<br>• TCP window size = 5 or 6<br>• Average packet-loss ratio maximum = 0.1%<br>For compression usage recommendations, see the *HPE MPX200 Multifunction Router User Guide* at **http://www.hpe.com/info/mpx200**. |

*Table Continued*



https://t.me/learningnets

| Product and minimum supported firmware version | XP storage system | Notes |
|---|---|---|
| B-series1606 Extension SAN Switch<br>DC Dir Switch MP Extension Blade<br>For firmware versions, see Table B-series Fibre Channel switches and routers. | P9500<br>XP24000<br>XP20000<br>XP12000<br>XP10000<br>XP7 | • Can be used as a switch and FCIP gateway with XP storage systems directly connected<br>• Can be used for Fibre Channel to Fibre Channel routing<br>• Requires `iodSET` (in-order delivery)<br>• Supports the following routing settings:<br>Port-based routing:<br>◦ `aptpolicy = 1`<br>◦ `aptpolicy -ap = 0` (port/asic load sharing)<br>◦ `dlsReset` (no dynamic load sharing)<br>Exchange-based routing:<br>◦ `aptpolicy = 3`<br>◦ `aptpolicy -ap = 0` (port/asic load sharing)<br>◦ `dlsSet` (dynamic load sharing)<br><br>For additional requirements, see Table 1606 Extension SAN Switch and DC Dir Switch MP Extension Blade features and requirements and HPE P9000 (XP) Continuous Access Fibre Channel distances. |

*Table Continued*



https://t.me/learningnets

| Product and minimum supported firmware version | XP storage system | Notes |
|---|---|---|
| B-series MP Router Blade<br>For firmware versions, see Table B-series legacy Fibre Channel switches and routers. | P9500<br>XP24000<br>XP20000<br>XP12000<br>XP10000<br>XP7 | • Can be used as a switch and FCIP gateway with XP storage systems directly connected<br>• Can be used for Fibre Channel to Fibre Channel routing<br>• Maximum transmission rate from 16 Mb/s to 1 Gb/s<br>• Requires `iodSET` (in-order delivery)<br>• Supports the following routing settings:<br>Port-based routing:<br>◦ aptpolicy = 1<br>◦ aptpolicy -ap = 0 (port/asic load sharing)<br>◦ `dlsReset` (no dynamic load sharing)<br>Exchange-based routing:<br>◦ aptpolicy = 3<br>◦ aptpolicy -ap = 0 (port/asic load sharing)<br>◦ `dlsSet` (dynamic load sharing)<br>For additional requirements, see Tables MP Router Blade features and requirements and HPE P9000 (XP) Continuous Access Fibre Channel distances. |
| C-series SN8000C, 9500 and 9200 series switches with IPS-8, IPS-4, 14/2, 18/4, and MDS 9222i<br>For firmware versions, see Table C-series Fibre Channel legacy switches. | P9500<br>XP24000<br>XP20000<br>XP12000<br>XP10000<br>XP7 | • Can be used as a switch and FCIP gateway with XP storage systems directly connected<br>• Distributed services TOV (D_S_TOV) = 5,000<br>• Error detect TOV (E_D_TOV) = 2,000<br>• Resource allocation TOV (R_A_TOV) = 10,000<br>• DataFieldSize = 2,112<br>• BB_credit from 16 (default) to 255, 14/2 and BB_credit maximum = 3,500 |

[1] For current support, see the SPOCK website at **http://www.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access.



https://t.me/learningnets

**Table 154: HPE P9000 (XP) Continuous Access iFCP gateway support**

| Product and minimum supported firmware version | XP storage system | Notes |
|---|---|---|
| Brocade-CNT Edge 3000 | XP24000<br>XP20000<br>XP12000<br>XP10000 | • Can be used in FCIP direct configuration<br>• Data compression possible<br>• No remote switch license required<br>• Maximum transmission rate from 16 Mb/s to 1 Gb/s |
| Nishan IPS3300 | XP12000<br>XP10000 | Contact a Hewlett Packard Enterprise storage representative. |
| Nishan IPS4300 | XP12000<br>XP10000 | Contact a Hewlett Packard Enterprise storage representative. |

**Table 155: HPE P9000 (XP) Continuous Access ESCON IP gateway support**

| Configuration | Product and minimum supported firmware version | Minimum IP bandwidth per path |
|---|---|---|
| P9000 (XP) Continuous Access ESCON over IP | CNT Ultranet Storage Director (USD) ESCON<br>Firmware: 2.6.2-0 or 3.1 | 100 Mb/s |

## OpenVMS host-based volume shadowing

Hewlett Packard Enterprise supports the following HPE and third-party devices and features for OpenVMS host-based volume shadowing:

- B-series1606 Extension SAN Switch and DC Dir Switch MP Extension Blade
- StoreFabric SN4000B SAN Extension Switch
- B-series FC and FCIP Fastwrite
- Cisco PA-FC-1G
- C-series MDS IP Storage Services Module (IPS-4, IPS-8)
- Cisco MDS 18/4 Multiservice Module (including the MDS 9222i Multiservice Fabric switch and gateway)
- C-series Write Acceleration



https://t.me/learningnets

# Certified third-party WDM, iFCP, and SONET products

This section describes the following topics:

**Certified third-party WDM products** on page 340

## Certified third-party WDM products

Hewlett Packard Enterprise supports P6000 Continuous Access, P9000 (XP) Continuous Access and 3PAR Remote Copy on all WDM products, including DWDM and CWDM, certified by the Fibre Channel switch vendors for the equivalent HPE switch models.

- **B-series switch products** — All Brocade WDM-certified products, listed by Brocade switch or router model number, are supported. Contact a Hewlett Packard Enterprise storage representative for the equivalent HPE 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 HPE model numbers are equivalent.

**NOTE:** CWDM support is based on the switch vendors' support for CWDM SFPs only. For information about Hewlett Packard Enterprise-supported standard (non-CWDM) SFPs, see the HPE switch model QuickSpecs or the Fibre Channel switch Streams on the SPOCK website at **http://www.hpe.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** on page 279 and **WDM system architectures** on page 279.



https://t.me/learningnets

# iSCSI storage

This chapter describes iSCSI storage in a SAN environment:


- **iSCSI overview** on page 341
- **iSCSI concepts** on page 342
- **iSCSI storage network requirements** on page 345
- **HPE Native iSCSI products** on page 346
- **HPE iSCSI bridge products** on page 358


## 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** on page 341
- **iSCSI bridge to Fibre Channel** on page 359


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. Hewlett Packard Enterprise 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 HPE 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



https://t.me/learningnets

# iSCSI concepts

This section describes key iSCSI concepts:


- **Initiator and target devices** on page 342
- **iSCSI naming** on page 342
- **Discovery mechanisms** on page 343
- **Sessions and logins** on page 344
- **Security** on page 344
- **Software and hardware iSCSI initiators** on page 345
- **Bridging and routing** on page 358
- **iSCSI boot** on page 345


## 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 116: iSCSI node definition** on page 342 shows an iSCSI node definition in the context of an IP network.

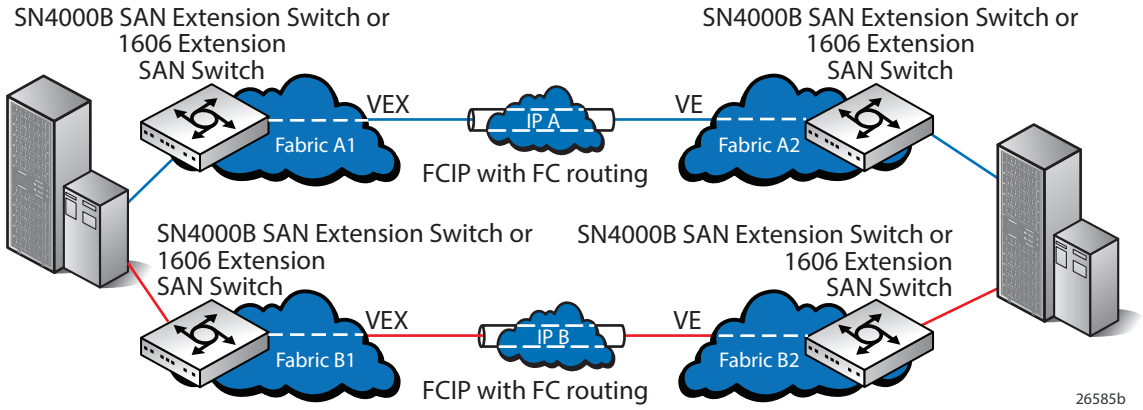



**Figure 116: iSCSI node definition**

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.



https://t.me/learningnets

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. Hewlett Packard Enterprise 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. Hewlett Packard Enterprise 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:



https://t.me/learningnets

- **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:

**Procedure**

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:
   a. Success means the target sends a login accept to the initiator; the session continues.
   b. 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:

- An initiator creates a unique ID by combining its iSCSI name with an ISID.
- During login, the initiator sends the ISID to the target.
- 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:


https://t.me/learningnets

- 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

Hewlett Packard Enterprise 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:


- **HPE StoreVirtual Storage** on page 355
- **MPX200 Multifunction Router with iSCSI for P6000/EVA storage** on page 359
- **C-series iSCSI** on page 377




https://t.me/learningnets

# HPE Native iSCSI products

This section describes HPE Native iSCSI storage system products. These products provide an iSCSI interface within the storage system hardware.

## HPE 3PAR StoreServ 20000 and StoreServ 8000

**HPE 3PAR StoreServ 20000 and 8000 iSCSI overview** on page 346

### HPE 3PAR StoreServ 20000 and 8000 iSCSI overview

The 3PAR StoreServ 20000 and StoreServ 8000 storage systems are available with different SAN or host interface options as described in HPE 3PAR StoreServ 20000 and 8000 host ports table.

This section describes 3PAR StoreServ 20000 and StoreServ 8000 10 GbE iSCSI support. For information about 3PAR StoreServ 20000 and StoreServ 8000 Fibre Channel support, see **HPE 3PAR StoreServ storage rules** on page 260.

### HPE 3PAR StoreServ 20000 and HPE StoreServ 8000 10 GbE iSCSI support

The 3PAR StoreServ 20000 and StoreServ 8000 10 GbE iSCSI interface option provides 10 GbE DCB iSCSI support (DCB iSCSI module).

### Operating system and multipath software support

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

#### Table 156: 10 GbE iSCSI operating system and multipath software support

| Operating System | Multipathing Software | Clusters |
|---|---|---|
| Windows 2008 R2 SP1 & 2012 R2 | Native MS MPIO | MS Failover cluster |
| RHEL 5.x, 6.x and 7.x | Linux inbox Multipath | RHEL KVM Cluster |
| ESX 5.x and 6.x | Native ESX Multipath | VMware cluster |
| Citrix Xenserver 5.x and 6.x | Citrix XenServer inbox Multipath | XenServer HA cluster |
| SLES 11.x and 12.x | Linux inbox Multipath | SLES HA cluster |

All hosts must have the appropriate Host Operation System type parameter set (Host Persona) and the required host settings described in the corresponding 3PAR Implementation Guide. For more information, see the specific host operating system 3PAR StoreServ 20000 and StoreServ 8000 Implementation Guide. These guides are located at **https://support.hpe.com/hpesc/public/home/documentHome?document_type=5000001%2C5000002&language_code=1871440&sort_by=relevance&sp4ts.oid=1008622482**.

**NOTE:** A multipathing driver is required to perform online firmware upgrades on 3PAR StoreServ 20000 and StoreServ 8000 storage systems.



https://t.me/learningnets

# HPE 3PAR StoreServ10000 and 7000


- **HPE 3PAR StoreServ 10000 iSCSI overview** on page 347
- **HPE 3PAR StoreServ 10000 and StoreServ 7000 10 GbE iSCSI support** on page 347


## HPE 3PAR StoreServ 10000 iSCSI overview

The 3PAR StoreServ 10000 V-Class storage system is available with different SAN or host interface options as described in Table HPE 3PAR StoreServ 10000 and 7000 host ports.

This section describes 3PAR StoreServ 10000 and StoreServ 7000 10 GbE iSCSI support. For information about 3PAR StoreServ 10000 and StoreServ 7000 Fibre Channel support, see **HPE 3PAR StoreServ storage rules** on page 260.

## HPE 3PAR StoreServ 10000 and StoreServ 7000 10 GbE iSCSI support

The 3PAR StoreServ 10000 and StoreServ 7000 10 GbE iSCSI interface option provides 10 GbE iSCSI support (iSCSI module).

### Operating system and multipath software support

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

### Table 157: 10 GbE iSCSI operating system and multipath software support

| Operating system | Multipath software | Clusters |
|---|---|---|
| Windows 2012 R2 & 2008 R2 SP1, 2003 | 3PAR MPIO (Windows 2003)<br>Microsoft MPIO DSM (Windows 2008, 2012) | MSCS (Windows 2003)<br>Failover Cluster (Windows 2008, 2012) |
| Red Hat Linux, SUSE Linux, Citrix XenServer | Device Mapper | Red Hat native cluster suite<br>Citrix XenServer native cluster suite |
| Solaris | Solaris MPxIO | Native Solaris cluster |
| VMware | VMware MPxIO | Native ESX/ESXi cluster solution |

All hosts must have the appropriate Host Operation System type parameter set (Host Persona) and the required host settings described in the corresponding 3PAR Implementation Guide.

For more information, refer to the specific host operating system 3PAR StoreServ 10000 V-Class and StoreServ 7000 Implementation Guide. These guides are located at **https://support.hpe.com/hpesc/public/home/documentHome?document_type=5000001%2C5000002&language_code=1871440&sort_by=relevance&sp4ts.oid=5335766**.

**NOTE:** A multipathing driver is required to perform online firmware upgrades on 3PAR StoreServ 10000 V-Class and StoreServ 7000 storage systems.



https://t.me/learningnets

# HPE 3PAR F-Class, T-Class


- **HPE 3PAR F-Class, T-Class iSCSI overview** on page 348
- **HPE 3PAR F-Class, T-Class iSCSI support** on page 348


## HPE 3PAR F-Class, T-Class iSCSI overview

The HPE 3PAR F-Class and T-Class storage systems are available with different SAN or host interface options as described in Table 3PAR F-Class, T-Class host ports.

This section describes 3PAR iSCSI support. For information about 3PAR Fibre Channel support, see **HPE 3PAR StoreServ storage rules** on page 260.

## HPE 3PAR F-Class, T-Class iSCSI support

The 3PAR F-Class and T-Class iSCSI interface option provides 1 GbE iSCSI support (iSCSI module).

### Operating system and multipath software support

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

### Table 158: iSCSI operating system and multipath software support

| Operating system | Multipath software | Clusters |
|---|---|---|
| Microsoft Windows Server 2008, 2003 | 3PAR MPIO (Windows 2003)<br>Microsoft MPIO DSM (Windows 2008) | MSCS (Windows 2003)<br>Failover Cluster (Windows 2008) |
| Red Hat Linux, SUSE Linux, Citrix XenServer | Device Mapper | Red Hat native cluster suite<br>Citrix XenServer native cluster suite |
| Solaris | Solaris MPxIO | Native Solaris cluster |
| VMware | VMware MPxIO | Native ESX/ESXi cluster solution |

All hosts must have the appropriate Host Operation System type parameter set (Host Persona) and the required host settings described in the corresponding 3PAR Implementation Guide.

For more information, refer to the specific host operating system 3PAR Implementation Guide. These guides are located at **https://support.hpe.com/hpesc/public/home/documentHome?document_type=5000001%2C5000002&language_code=1871440&sort_by=relevance&sp4ts.oid=5335766**.

**NOTE:** A multipathing driver is required to perform online firmware upgrades on 3PAR storage systems.

# P6300/P6350/P6500/P6550 EVA


- **P6300/P6350/P6500/P6550 EVA overview** on page 349
- **P6300/P6350/P6500/P6550 EVA iSCSI support** on page 349




https://t.me/learningnets

### P6300/P6350/P6500/P6550 EVA overview

The P6300/P6350/P6500/P6550 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/P6350/P6500/P6550 EVA Fibre Channel support, see **P6000 storage system rules** on page 238. For information about P6300/P6350/P6500/P6550 EVA FCoE support, see **Fibre Channel over Ethernet** on page 65.

### P6300/P6350/P6500/P6550 EVA iSCSI support

The P6300/P6350/P6500/P6550 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 P6000 (see **Figure 39: P63xx/P65xx FCoE/iSCSI end-to-end configuration** on page 89). 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 SPOCK website at **http://www.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access.

#### Table 159: iSCSI/FCoE operating system and multipath software support

| Operating system | Multipath software | Clusters |
|---|---|---|
| Apple Mac OS X (1GbE iSCSI only) | None | None |
| Microsoft Windows Server 2008, 2003 | MPIO with HPE DSMMPIO 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 *HPE P6300/P6500 Enterprise Virtual Array User Guide*.

## HPE StorageWorks MSA family of iSCSI SAN arrays

This section describes the following topics:



https://t.me/learningnets

- **HPE MSA 2050/2052 SAN Overview** on page 350
- **HPE MSA 1050 iSCSI SAN Overview** on page 350
- **HPE MSA 2040 SAN overview** on page 350
- **HPE MSA 1040 iSCSI overview** on page 350
- **P2000 G3 FC/iSCSI, P2000 G3 10Gb iSCSI, P2000 G3 iSCSI overview** on page 351
- **MSA iSCSI storage family maximum configurations** on page 352


## HPE MSA 2050/2052 SAN Overview

The flash-ready HPE MSA 2050/2052 Storage system with the Gen10 ProLiant Branding is designed for affordable application acceleration that is ideal for small and remote office deployments. But do not let the low cost fool you. The HPE MSA 2050/2052 Storage system gives you the combination of simplicity, flexibility to grow now and into the future, and advanced features you may not expect to find in an entry-priced array. Start small and scale as needed with any combination of solid state disks (SSD), high-performance enterprise, or lower-cost midline SAS-based drives.

The MSA 2050 SAN controller provides customers a wide range of protocol options, supporting 8Gb or 16Gb Fibre Channel along with 1GbE or 10GbE iSCSI. The controller connectivity options extend to supporting all Fibre Channel, all iSCSI or a combination of protocols for the most flexibility. See the QuickSpecs for valid configurations. For host OS and HBA connectivity requirements, see the SPOCK website at **http://www.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access.

## HPE MSA 1050 iSCSI SAN Overview

The HPE MSA 1050 Storage brings affordable flash storage down to ever lower customer segments. The MSA 1050 is designed to meet entry-level storage requirements and budgetary constraints. With the lowest storage array price points in the HPE Storage portfolio and field-proven ProLiant compatibility, the MSA 1050 is the platform of choice for smaller SAS, iSCSI, and FC deployments.

The HPE MSA 1050 features 8Gb Fibre Channel, 12 Gb SAS, 1GbE, and 10GbE iSCSI at previously unattainable entry price points. The array allows you to take advantage of the latest storage technologies in simple and efficient ways by providing a good balance between performance and budget resulting in a highly favorable $/GB return on their investment.

## HPE MSA 2040 SAN overview

The HPE MSA 2040 SAN storage system is a high-performance storage array designed for entry-level HPE customers desiring 8 and/or 16Gb Fibre Channel, 1 and/or 10GbE iSCSI connectivity with 4 host ports per controller. The MSA 2040 storage array provides an excellent value for customers needing performance balanced with price to support initiatives such as consolidation and virtualization.

The MSA 2040 SAN controller allows customers to create their own combo controller by mixing FC and iSCSI SFPs. See the QuickSpecs for valid configurations. Options include 1Gb RJ-45 SFP+, 10Gb optical transceiver SFP+, and Direct Attach Copper (DAC) cables. For host OS and HBA connectivity requirements, see the SPOCK website at **http://www.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access.

## HPE MSA 1040 iSCSI overview

The MSA 1040 storage is designed for entry level market needs, and features 1GbE and 10GbE iSCSI at previously unattainable entry price points. The new array allows users to take advantage of the latest storage technologies in simple and efficient ways by providing a good balance between performance and budget resulting in a highly favorable $/GB return on their investment.



https://t.me/learningnets

The MSA 1040 1GbE controllers include 1GbE RJ-45 SFP+ for host connectivity. The MSA 1040 10GbE controllers include 10GbE optical transceivers for host connectivity. DAC cables are supported only with the MSA 1040 10GbE controllers. For host OS and HBA connectivity requirements, see the SPOCK website at **http://www.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access.

## P2000 G3 FC/iSCSI, P2000 G3 10Gb iSCSI, P2000 G3 iSCSI overview

The HPE MSA 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 160: P2000 G3 controller types

| Controller | FW version | Redundancy | Interface | Expansion |
|---|---|---|---|---|
| P2000 G3 FC/iSCSI | TS251P0**xx** | Single or dual controller | Two 1-GbE iSCSI[1] | P2000 LFF drive enclosure |
| P2000 G3 10-GbE iSCSI | | | Two 10-GbE SFP +[2] | D2700 SFF disk enclosure |
| P2000 G3 iSCSI | | | Four 1-GbE iSCSI | MSA2000 LFF drive enclosure |
| | | | | MSA70 SFF drive enclosure |

1 Controller also includes two 8Gb Fibre Channel ports described in **MSA storage system rules** on page 227.

2 SFP+ optical transceivers or SFP+ direct attached cables must be purchased separately.



https://t.me/learningnets

## MSA iSCSI storage family maximum configurations

### Table 161: MSA iSCSI storage family maximum configurations

| Model | Array chassis | Expansion[1] | Expansion maximum |
|---|---|---|---|
| MSA 2050/2052 SAN | Preconfigured dual controller SFF or LFF | MSA 2050 SFF disk enclosure<br>MSA 2050 LFF disk enclosure<br>MSA 2040 LFF disk enclosure[2]<br>D2700 SFF disk enclosure[1] | 8 total enclosures<br>96 LFF drives<br>192 SFF drives |
| MSA 1050 1G iSCSI<br>MSA 1050 10G iSCSI | Preconfigured dual controller SFF or LFF | MSA 2050 SFF Disk Enclosures<br>MSA 2050 LFF Disk Enclosure<br>MSA 2040 LFF Disk Enclosure[1]<br>D2700 SFF Disk Enclosure 4 total enclosures[1] | 48 LFF drives<br>96 SFF drives |
| MSA 2040/2042 SAN | MSA 2040 SFF<br>MSA 2040 LFF | MSA 2040 LFF disk enclosure<br>D2700 SFF disk enclosure<br>P2000 LFF drive enclosure | 8 total enclosures<br>96 LFF drives<br>199 SFF drives |
| HPE MSA 1040 1GbE iSCSI<br>HPE MSA 1040 10GbE iSCSI | Preconfigured dual controller SFF or LFF | MSA 2040 LFF disk enclosure<br>D2700 SFF disk enclosure<br>P2000 LFF drive enclosure | 4 total enclosures<br>48 LFF drives<br>99 SFF drives |
| P2000 FC/iSCSI<br>P2000 G3 10-GbE iSCSI<br>P2000 G3 iSCSI | P2000 3.5-in drive bay (LFF)<br>P2000 2.5-in drive bay (SFF)<br>2012 3.5-in drive bay (LFF, upgrade only)<br>2024 2.5-in drive bay (SFF, upgrade only) | P2000 LFF drive enclosure<br>D2700 SFF disk enclosure<br>MSA2000 LFF drive enclosure<br>MSA70 SFF drive enclosure | 8 total enclosures or 149 drives<br>96 LFF drives<br>149 SFF drives |

[1] Confirm cabling requirements and limitation in the Cable Configuration Guide at **http://www.hpe.com/info/msa**



https://t.me/learningnets

[2] Supported in upgrade from MSA 2040 only

## Server support

The MSA2000i support the following ProLiant servers:

- HPE ProLiant DL, ML
- HPE ProLiant c-Class BladeSystem

**NOTE:** BL20p G1 servers are not supported.

## Operating system support

### Table 162: MSA iSCSI storage family operating system support

| Storage system | Firmware | Operating systems |
|---|---|---|
| MSA 2050/2052 | VL100P001 | Microsoft Windows<br>Red Hat Enterprise Linux<br>SUSE Linux Enterprise Server<br>VMware ESX |
| MSA 1050 | VE100R002 | Microsoft Windows<br>Red Hat Enterprise Linux<br>SUSE Linux Enterprise Server<br>VMware ESX |
| MSA 2040/2042 SAN | GL220P009 (or later) | Microsoft Windows<br>OpenVMS<br>Red Hat Enterprise Linux<br>SUSE Linux Enterprise Server<br>VMware ESX<br>XenServer CITRIX |

*Table Continued*



https://t.me/learningnets

| Storage system | Firmware | Operating systems |
|---|---|---|
| MSA 1040 | GL220P010 (or later) | Microsoft Windows<br>Red Hat Enterprise Linux<br>SUSE Linux Enterprise Server<br>VMware ESX |
| P2000 G3<br>• FC<br>• FC/iSCSI | TS252P005 (or later) | Apple OS X<br>Microsoft Windows<br>Red Hat Enterprise Linux<br>SUSE Linux Enterprise Server<br>VMware ESX<br>XenServer CITRIX |

**Legend:** X: Supported; P: Planned - check SPOCK for details **http://www.hpe.com/storage/spock**

## Path failover software

The MSA iSCSI storage family supports the following:

- Microsoft MPIO basic failover software for Windows operating systems
- Device Mapper for path failover with Linux operating systems
- VMware embedded multipath

## Management software support

The MSA iSCSI storage family supports target-based management interfaces, including Telnet (CLI), SSH (CLI), FTP, SFTP, SMI-S, SNMP, HTTP and HTTPS web-based interface. The web-based interface is supported with Microsoft Internet Explorer, Google Chrome and Mozilla Firefox.


https://t.me/learningnets

### Maximum configurations

**Table 163: MSA iSCSI storage family maximum configurations**

| Storage systems | Operating systems | Drives | Hosts | LUNs | LUN size | Snapshots and clones[1] |
|---|---|---|---|---|---|---|
| MSA 2050 / 2052 SAN | Microsoft Windows<br>Red Hat Enterprise Linux<br>SUSE Linux Enterprise Server<br>VMware ESXi | 96 3.5-in LFF<br>192 2.5-in SFF | 1024 | 1024 | 128TiB | 64 Standard (maximum 512 snapshots) |
| MSA 1050 SAN | Microsoft Windows<br>Red Hat Enterprise Linux<br>SUSE Linux Enterprise Server<br>VMware ESXi | 48 3.5-in LFF<br>96 2.5-in SFF | 1024 | 1024 | 128TiB | 64 Standard (maximum 512 snapshots) |
| MSA 2040 SAN | Microsoft Windows<br>Red Hat Enterprise Linux<br>SUSE Linux Enterprise Server<br>VMware ESXi | 96 3.5-in LFF<br>199 2.5-in SFF | 1024[2] | 512 | Up to 64 TB, depending on vdisk configuration | 64 standard (maximum 512 snapshots) |
| MSA 1040 FC | Microsoft Windows<br>Red Hat Enterprise Linux<br>SUSE Linux Enterprise Server<br>VMware ESXi | 48 3.5-in LFF<br>99 2.5-in SFF | 1024[2] | 512 | Up to 64 TB, depending on vdisk configuration | 64 standard (maximum 512 snapshots) |
| P2000 G3 FC<br>P2000 G3 FC/ iSCSI | Apple OS X<br>Citrix XenServer<br>Microsoft Windows<br>Red Hat Enterprise Linux<br>SUSE Linux Enterprise Server<br>VMware ESX , ESXi | 96 3.5-in LFF<br>149 2.5-in SFF | 64 | 512 | Up to 64 TB, depending on vdisk configuration | 64 standard<br>(maximum 512 snapshots, clones or remote snaps) |

1 Snapshots and volume copies require additional licenses.
2 Using Firmware GL200 or later and Virtualized Storage

## HPE StoreVirtual Storage

This section describes the following topics:



https://t.me/learningnets

- **HPE StoreVirtual Storage overview** on page 356
- **HPE StoreVirtual 3200 Storage support** on page 357
- **HPE StoreVirtual 4000 Storage support** on page 357


In this section, StoreVirtual Storage refers to StoreVirtual Storage, StoreVirtual 3200, HPE LeftHand Storage, P4000 G2, and HPE LeftHand P4000 products, as well as StoreVirtual VSA, P4000 VSA, and HPE LeftHand VSA products.

## HPE StoreVirtual Storage overview

**IMPORTANT:** StoreVirtual Storage is the new name for LeftHand Storage, and P4000 SAN solutions.

StoreVirtual Storage is composed of multiple storage nodes consolidated into single or multiple pools of storage. 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 presents the volumes as targets to iSCSI initiators.

Key features include:

- A cluster of pooled storage. For product specific configurations, contact a Hewlett Packard Enterprise storage representative.
- Multiple layers of high availability:
  - Multiple RAID levels within the storage node (RAID5, 6, and 10)
  - Multiple Network RAID levels across the storage nodes on a per volume basis (NetworkRAID 0, 10, 10+1, 10+2, 5, 6)
- 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.
- StoreVirtual Storage comes with an all-inclusive enterprise feature set which includes: storage clustering, Network RAID, application integrated snapshots, Remote Copy, and thin provisioning.
- Multi-site cluster support — StoreVirtual Storage can be deployed in clusters that are stretched across multiple racks, data rooms or data centers. These installations require:
  - A stretched 1 GbE or 10 GbE Ethernet network

**NOTE:** 50 MB/s (400 Mb/s) of bandwidth per storage node pair needs to be allocated on the 1 GbE network between the locations. 200 MB/s (1,600 Mb/s) of bandwidth per storage node pair needs to be allocated on the 10 GbE network between the locations. Network latency among storage nodes cannot exceed 1 ms.

For more information, including links to manuals, see **http://www.hpe.com/info/StoreVirtual3200Manuals** and **http://www.hpe.com/support/storevirtualmanuals**.



https://t.me/learningnets

## HPE StoreVirtual 3200 Storage support

For the latest information on version support, see the **SPOCK website**. For information on manuals, see **http://www.hpe.com/info/StoreVirtual3200Manuals**.

## HPE StoreVirtual 4000 Storage support

For the latest information on version support, see the *HPE StoreVirtual 4000 Storage Compatibility Matrix* at **http://www.hpe.com/info/P4000Compatibility** or the SPOCK website at **http://www.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access to SPOCK.

## Multi-pathing software

StoreVirtual 4000 Storage supports the following:

- HPE StoreVirtual DSM for MPIO
- Microsoft DSM and Windows built-in MPIO support for Microsoft Windows Server 2008 and higher
- VMware native MPIO (Round-robin is preferred policy)
- Device Mapper MPIO for Red Hat Enterprise Linux
- Citrix/Linux/Unix bonding of network interfaces performed at the networking layer
- For more information, see the following documents:
  - *HPE StoreVirtual 4000 Storage Application Aware Snapshot Manager Deployment Guide* at **http://www.hpe.com/info/P4000Support** (in the Manuals section)
  - *HPE StoreVirtual 4000 Storage with VMware vSphere: Design Considerations and Best Practices* at **https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/techpaper/vmw-vsphere-p4000-lefthand-san-solutions-white-paper.pdf**

## Management software support

StoreVirtual 4000 Storage supports target-based management interfaces, including the StoreVirtual Centralized Management Console for Windows and Linux(CMC), Command Line Interface for Windows (CLIQ), and on-node command line interface (SSH to storage node using port 16022).

For more information, see the following documents, available at **http://www.hpe.com/info/P4000Support** (see Manuals section):

- *HPE StoreVirtual Command-Line Interface User Manual*
- *HPE StoreVirtual 4000 Storage User Guide*

## Maximum configurations

StoreVirtual Centralized Management Console shows a Configuration Summary which 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.



https://t.me/learningnets

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 management groups and clusters, applications accessing the volumes, and purpose of using snapshots. The following sections contain guidelines that can help you manage StoreVirtual to obtain the best and safest performance and scalability for your circumstances. These guidelines are in line with tested limits for common 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 164: 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.

# HPE 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:

**Procedure**

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.



https://t.me/learningnets

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 Hewlett Packard Enterprise-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 Hewlett Packard Enterprise-supported iSCSI to Fibre Channel bridge products are as follows:

- B-series iSCSI Director Blade
- C-series IP Storage Services Modules
- C-series 18/4 Multiprotocol Module
- MPX200 Multifunction Router iSCSI
- EVA iSCSI Connectivity Option
- EVA4400 iSCSI Connectivity Option

## MPX200 Multifunction Router with iSCSI for P6000/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, 3PAR (see **MPX200 Multifunction Router with iSCSI for 3PAR StoreServ Storage** on page 366), or XP 24000/20000 storage systems (see **MPX200 Multifunction Router with iSCSI for XP storage** on page 370), and 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 Hewlett Packard Enterprise 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 one other license-enabled function — 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** on page 65. For information about using the MPX200 data migration feature, see **MPX200 Multifunction**



https://t.me/learningnets

**Router with data migration** on page 243. For information about using the MPX200 FCIP feature, see **MPX200 Multifunction Router with FCIP** on page 302. For information about configuring the MPX200 for multiple functions concurrently, see the *HPE 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 117: MPX200-EVA single-blade fabric-attached configuration** on page 360 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.

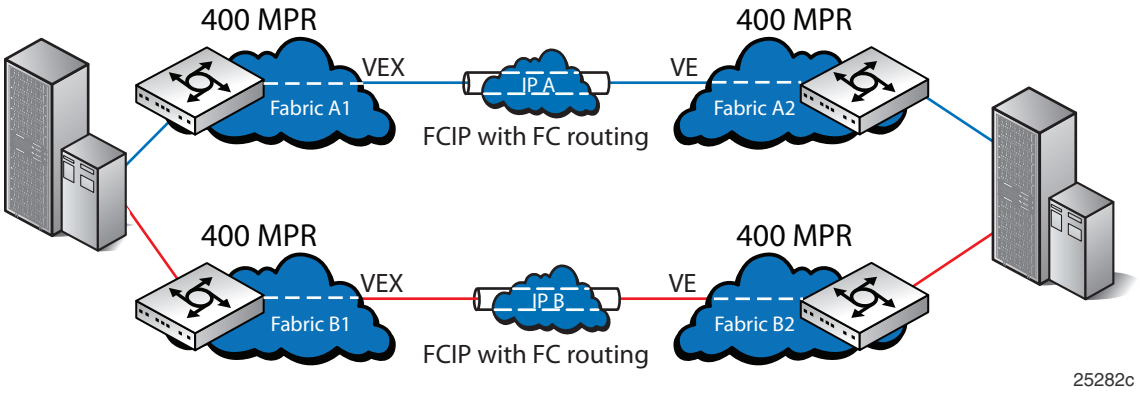



**Figure 117: MPX200-EVA single-blade fabric-attached configuration**

**Figure 118: MPX200-EVA dual-blade fabric-attached configuration** on page 360 shows an MPX200-EVA dual-blade fabric-attached configuration. This configuration provides high availability with failover between blades.

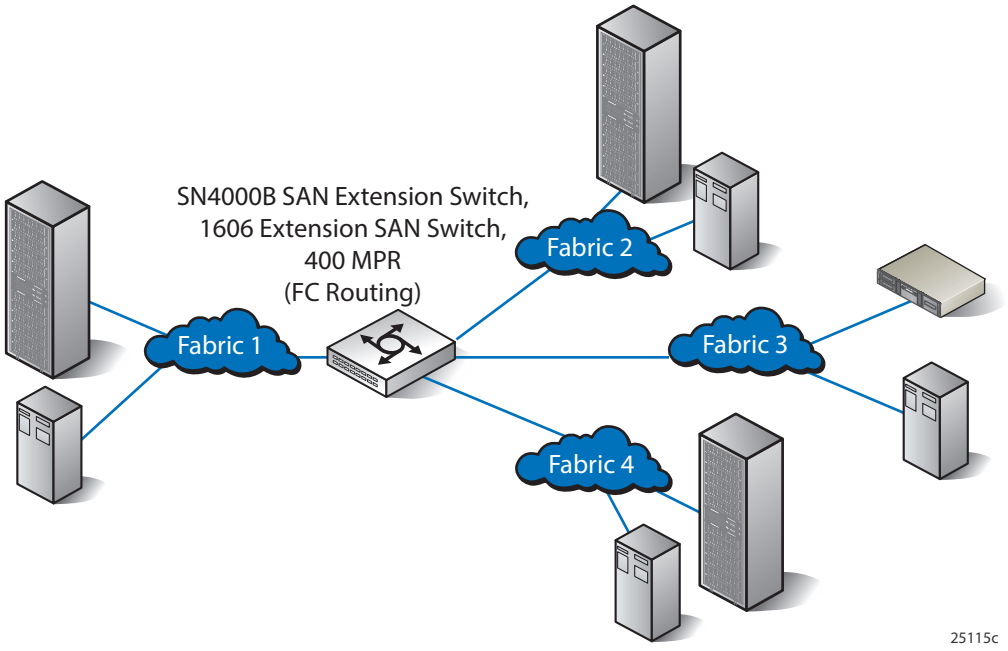



**Figure 118: MPX200-EVA dual-blade fabric-attached configuration**

**Figure 119: MPX200 single-blade multi-EVA configuration** on page 361 shows a multi-EVA configuration with connectivity for up to four EVA storage systems from a single MPX200 blade.



https://t.me/learningnets

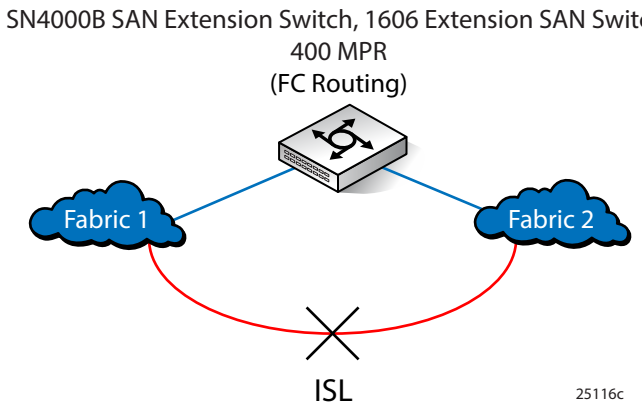



**Figure 119: MPX200 single-blade multi-EVA configuration**

**Figure 120: MPX200 dual-blade multi-EVA configuration** on page 361 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.

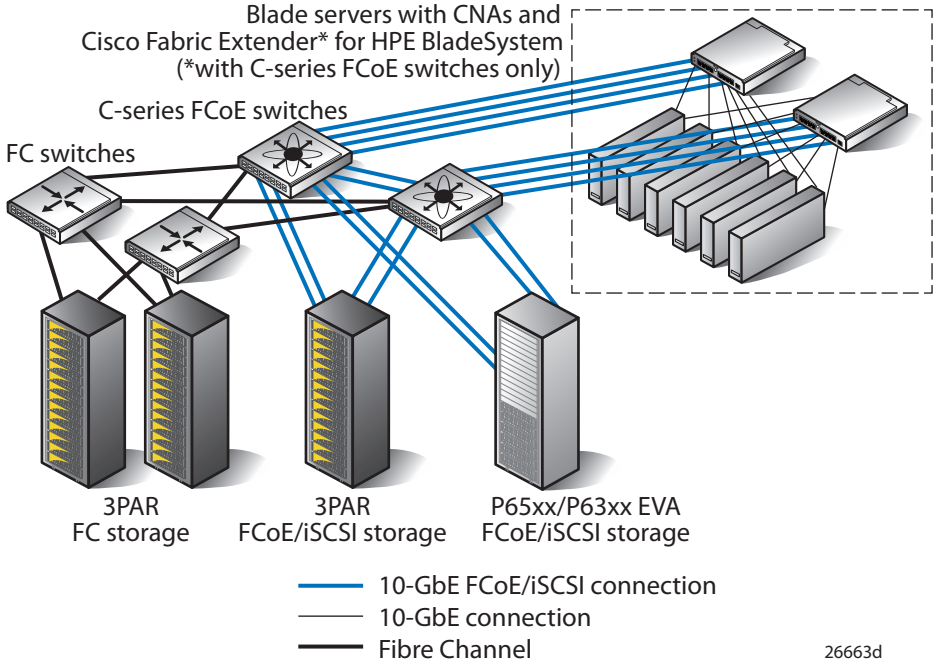



**Figure 120: MPX200 dual-blade multi-EVA configuration**

**Figure 121: MPX200 dual-blade direct connect to one EVA configuration** on page 362, **Figure 122: MPX200 single-blade direct connect to one EVA configuration** on page 362, and **Figure 123: MPX200 dual-blade direct connect to two EVA configuration** on page 363 illustrate EVA direct connect configurations.



https://t.me/learningnets

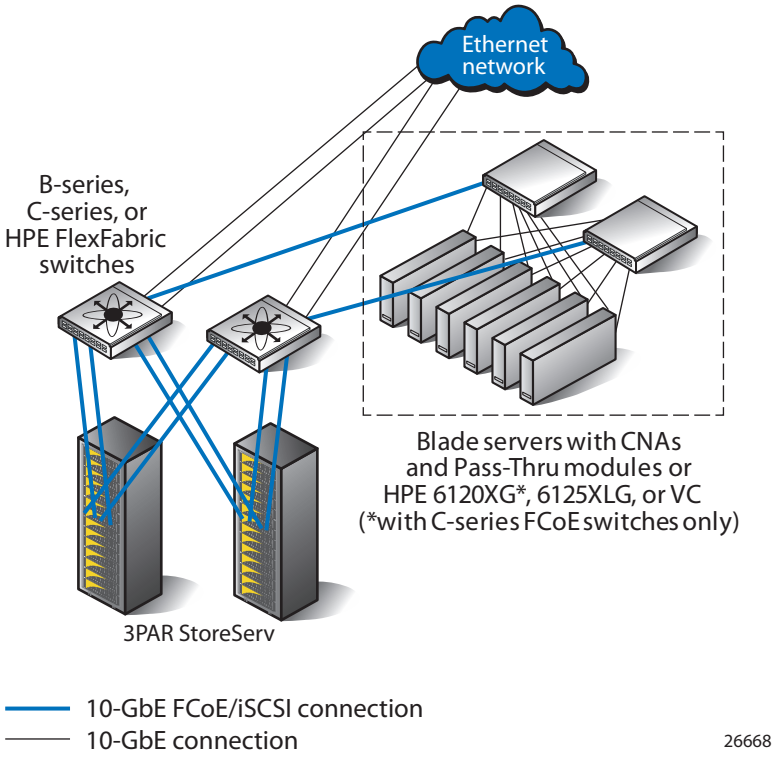



**Figure 121: MPX200 dual-blade direct connect to one EVA configuration**

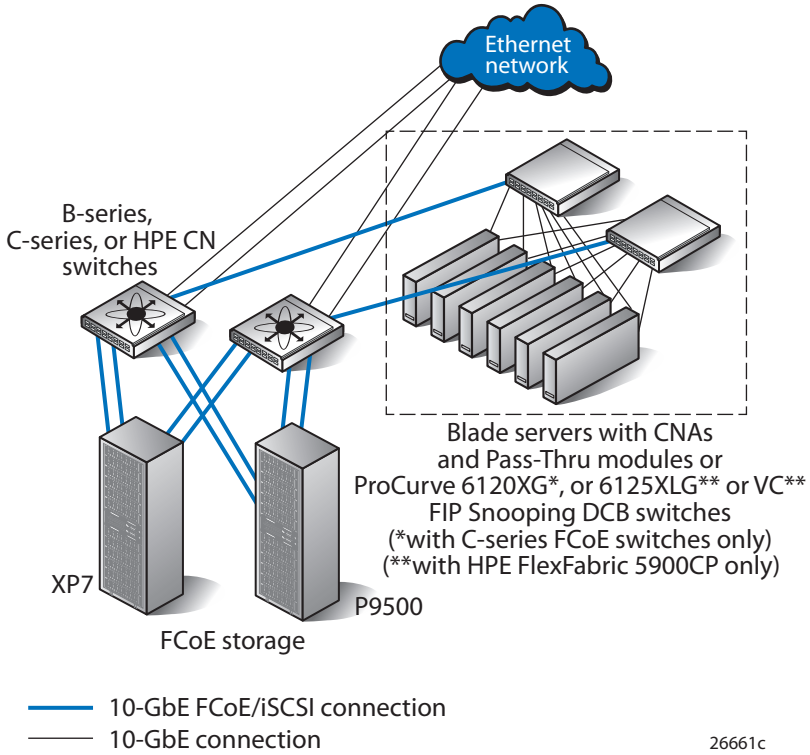



**Figure 122: MPX200 single-blade direct connect to one EVA configuration**



https://t.me/learningnets

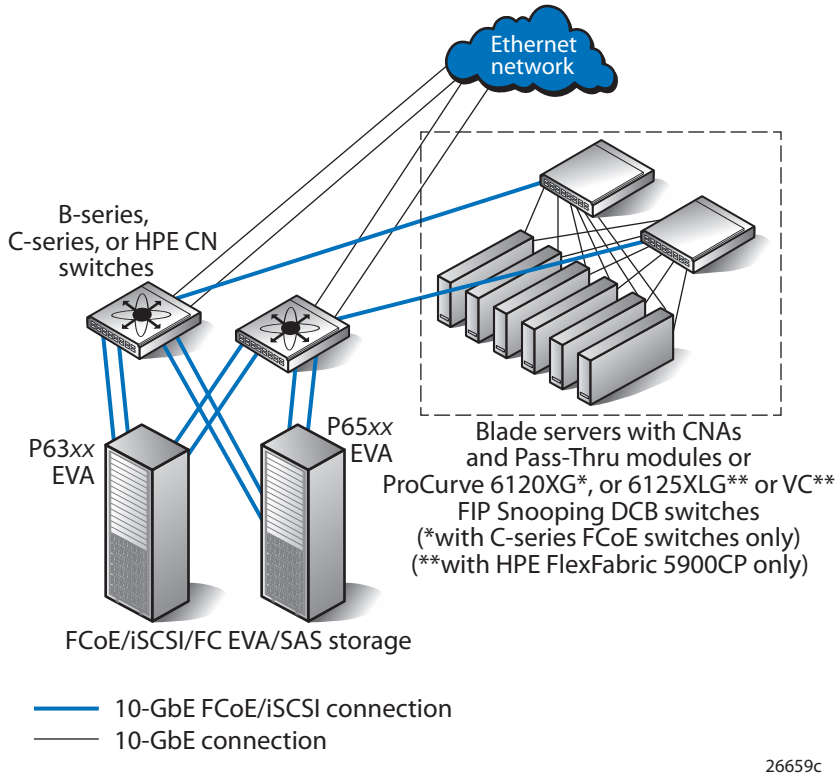



**Figure 123: MPX200 dual-blade direct connect to two EVA configuration**

## 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 165: Supported MPX200 maximums

| Description | Maximum per MPX200 solution[1] |
|---|---|
| **Hardware** | |
| P6000/EVA and/or XP24000/20000 storage systems (see **MPX200 Multifunction Router with iSCSI for XP storage** on page 370) | 4 total (any combination) |
| MPX200 | 1 chassis with up to 2 blades |
| MPX200 iSCSI port connections | See the *HPE MPX200 Multifunction Router User Guide* at **http://www.hpe.com/info/mpx200**. |
| **Configuration parameter** | |

*Table Continued*



https://t.me/learningnets

| Description | Maximum per MPX200 solution[1] |
|---|---|
| Total number of iSCSI initiators | 300 per chassis for 1-GbE (1 or 2 blades)<br>600 per chassis for 10-GbE (1 or 2 blades) |
| 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 |

[1] For mixed blade type chassis configurations that include one 1-GbE blade and one 10-GbE blade, the maximum values for a 1-GbE blade prevail.

## MPX200 blade configurations

The MPX200 supports the following functions: iSCSI-FCoE, FCIP, Data Migration. For simultaneous operation, you can configure the MPX200 chassis with a single blade or dual blades to run up to two functions per blade in combinations.

### Table 166: MPX200 blade configurations

| Single blade chassis (blade1/empty) | Dual-blade chassis (blade1/blade2) |
|---|---|
| iSCSI-FCoE/empty | iSCSI-FCoE/iSCSI-FCoE |
| iSCSI-FCoE-FCIP/empty | iSCSI-FCoE-FCIP/iSCSI-FCoE-FCIP |
| iSCSI-FCoE-DMS/empty | iSCSI-FCoE-DMS/iSCSI-FCoE-DMS |
| FCIP/empty | FCIP/FCIP |
| DMS/empty | DMS/DMS |
| | iSCSI-FCoE-DMS/iSCSI-FCoE-FCIP |
| | iSCSI-FCoE-FCIP/iSCSI-FCoE-DMS |

- Simultaneous iSCSI and FCoE are considered one function. FCoE is only supported with 10-GbE models.
- When configuring for blade redundancy, you must configure both blades. To add a redundant blade, you must un-present/re-present existing LUN presentations to gain access through the second blade.
- Dual-blade iSCSI-FCoE configurations are always configured for high availability.
- Dual-blade FCIP configurations can be configured for separate operation or high availability.
- A license is required for FCIP, half-chassis or full chassis.
- A license is required for data migration, 1TB, 5TB, or 1 Array. FCIP is not required for remote data migration.

**NOTE:** For more information on data migration, see *HPE Data Migration Services User's Guide*.



https://t.me/learningnets

## P6000/EVA storage system rules and guidelines

The MPX200 is supported with the following P6000/EVA storage systems:

P6300/P6350/P6500/P6550

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 167: 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<br>1,320 with 4 EVAs | 852 with 1 EVA<br>1,608 with 4 EVAs |
| Maximum number of initiators | 1,024 | 1,308 with 1 EVA<br>4,332 with 4 EVAs | 1,608 with 1 EVA<br>4,632 with 4 EVAs |

## HPE P6000 Command View and MPX200 management rules and guidelines

The HPE 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 P6000 Command View. For more information, see the *HPE MPX200 Multifunction Router User Guide* at **http://www.hpe.com/info/mpx200**. From the website, under storage, click **Storage Networking**, and then under Routers Gateways/Multiplexers, click **HP MPX200 Multifunction Router**.

Observe the following MPX200 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.
- P6000 Command View manages the MPX200 out of band (IP) through the MPX200 management IP port. The 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 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, Oracle Solaris, VMware, and Windows 2008 host mode setting is configured with P6000 Command View.



https://t.me/learningnets

**NOTE:** Communication between P6000 Command View and the MPX200 is not secured by the communication protocol. If this unsecured communication is a concern, Hewlett Packard Enterprise recommends a confined or secured IP network within a data center.

### P6000/EVA storage system software

For FCIP, the MPX200 is supported with P6000 Continuous Access, Business Copy, SSSU, or Replication Solutions Manager.

### Fibre Channel switch and fabric support

The MPX200 is supported with 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.hpe.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.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access.

#### Table 168: MPX200-EVA operating system, multipath software, and cluster support

| Operating system | Multipath software | Clusters | P6000/EVA storage system |
|---|---|---|---|
| Apple Mac OS X | None | None | P6300/P6350/P6500/P6550 |
| Microsoft Windows Server 2012, 2008, 2003, Hyper-V | MPIO with HPE DSM<br>MPIO with Microsoft DSM | MSCS[1] | |
| Oracle VM Server | Native DM | None | |
| Linux Red Hat, SUSE | Device Mapper | None | |
| Oracle Solaris | Solaris MPxIO | None | |
| VMware | VMware MPxIO | None | |

[1] MSCS is supported with Microsoft DSM only.

## MPX200 Multifunction Router with iSCSI for 3PAR StoreServ Storage

The 3PAR StoreServ family of FC storage systems supports integrated iSCSI connectivity using the MPX200 Multifunction Router. The MPX200 hardware is integrated with up to four 3PAR StoreServ, P6000 EVA (see **MPX200 Multifunction Router with iSCSI for P6000/EVA storage** on page 359, or XP24000/20000 storage systems (see **MPX200 Multifunction Router with iSCSI for XP storage** on page 370) 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 HPE as an option for a new 3PAR purchase or as a field upgrade to existing 3PAR StoreServ 10000 V-Class; 3PAR StoreServ 7000,3PAR F-Class, or T-Class storage systems. With this product, iSCSI connectivity to the 3PAR StoreServ Storage systems is provided for servers through a standard 1-GbE or 10-GbE NIC.



https://t.me/learningnets

## MPX200 configuration options

A 3PAR storage system can be configured for simultaneous connectivity to iSCSI and FC attached hosts. Support for iSCSI to a 3PAR storage system is provided through the MPX200 and an existing FC switch fabric port (fabric-attached).

**Figure 124: MPX200-3PAR single-blade fabric-attached configuration** on page 367 shows an MPX200-3PAR single-blade fabric-attached configuration. This is the lowest-cost configuration and is used when high availability for iSCSI hosts is not required.

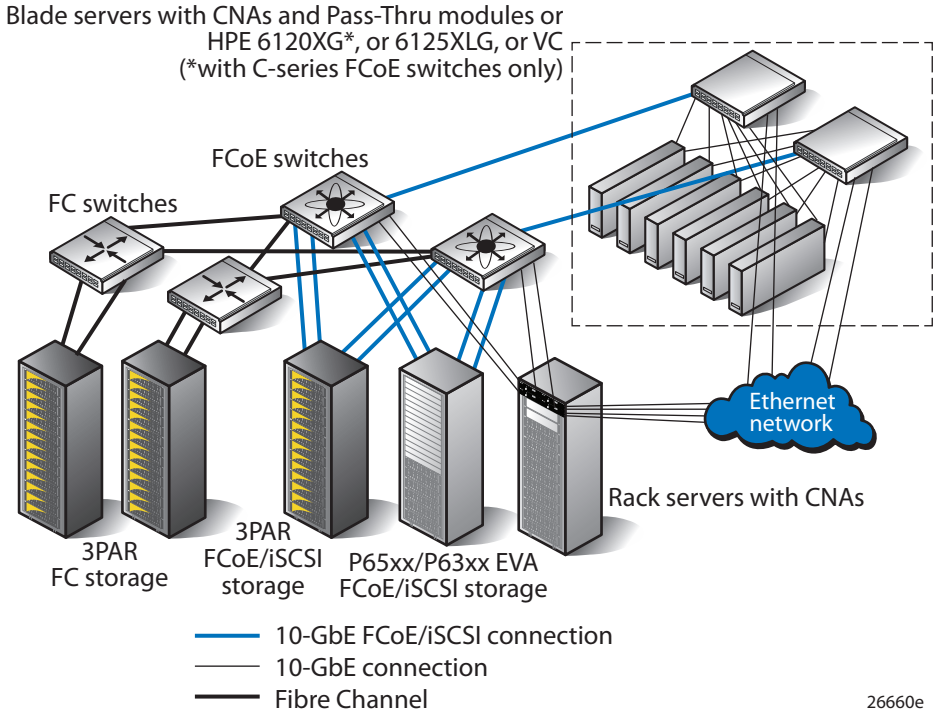



**Figure 124: MPX200-3PAR single-blade fabric-attached configuration**

**Figure 125: MPX200-3PAR dual-blade fabric-attached configuration** on page 368 shows an MPX200-3PAR dual-blade fabric-attached configuration. This configuration provides high availability with failover between blades.



https://t.me/learningnets

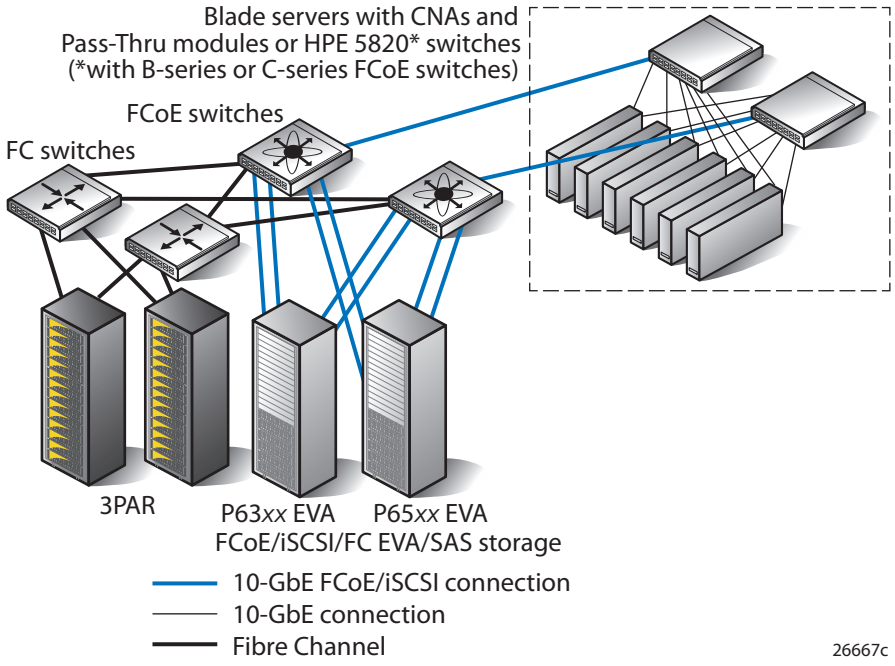



**Figure 125: MPX200-3PAR dual-blade fabric-attached configuration**

**Figure 126: MPX200-3PAR multi-3PAR fabric-attached configuration** on page 369 shows a multi-3PAR configuration with connectivity for up to four 3PAR storage systems from a single MPX200 blade.



https://t.me/learningnets

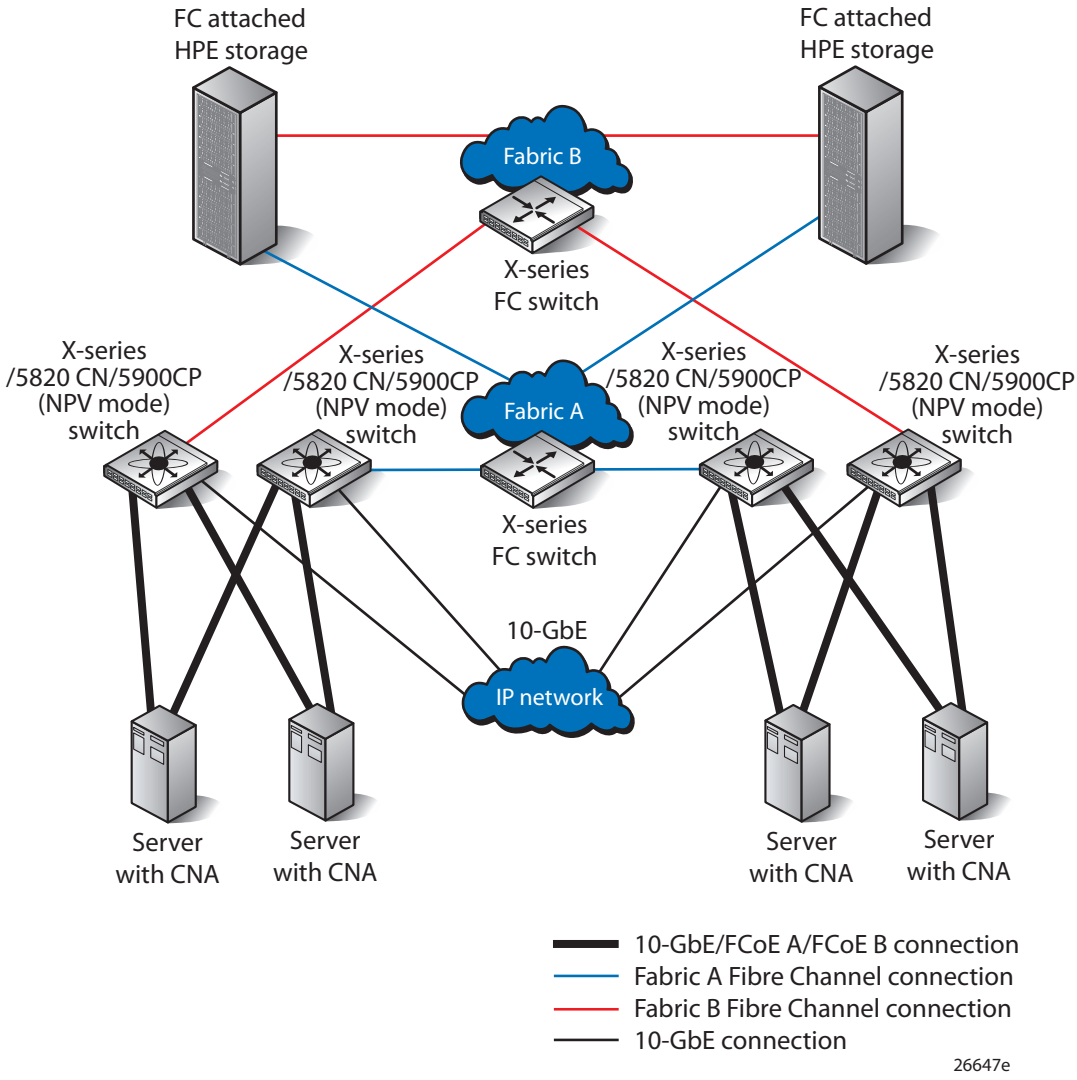



**Figure 126: MPX200-3PAR multi-3PAR fabric-attached configuration**

## 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.

For information about 3PAR FC host connectivity, see the 3PAR documentation.

## 3PAR storage system rules and guidelines

The MPX200 is supported with the following 3PAR storage systems:



https://t.me/learningnets

- 3PAR StoreServ 10000 V-Class
- 3PAR StoreServ 7000
- 3PAR F-Class, T-Class

All MPX200 configurations must follow these connectivity rules:

- When using the MPX200 for iSCSI, MPX200 FC connections can be fabric-attached through an FC switch or direct-connect to a 3PAR FC port.
- Multiple MPX200 chassis can be connected to a single 3PAR array. However, Hewlett Packard Enterprise recommends that array FC ports are not shared between different chassis.
- Hewlett Packard Enterprise recommends a maximum of eight 3PAR array ports be connected to a single MPX200 chassis.
- A maximum of one MPX200 chassis (two blades) can be zoned with up to four 3PAR storage systems.
- 3PAR, XP and P6000 EVA storage systems can connect to the same MPX200. The total allowable number of storage systems is four per MPX200 chassis.
- A 3PAR storage system can present LUNs to iSCSI initiators and FC hosts concurrently. 3PAR does not support presenting the same LUN to both iSCSI and FC initiators at the same time.

**Table 169: MPX200-3PAR StoreServ operating system and multipath support**

| Operating system | Multipath software | 3PAR storage system |
|---|---|---|
| Citrix Xen | Native MPxIO | 3PAR StoreServ 10000 V-Class; 3PAR StoreServ 7000; 3PAR F-Class, T-Class |
| Microsoft Windows Server 2012, 2008, 2003 | 3PAR MPIO (Windows 2003)<br>Microsoft MPIO DSM (Windows 2012, 2008) | |
| Oracle VM Server | Native Device Mapper | |
| Linux Red Hat, SUSE | Device Mapper | |
| Oracle Solaris | Solaris MPxIO | |

## 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, P6000/EVA (see **MPX200 Multifunction Router with iSCSI for P6000/EVA storage** on page 359), or 3PAR StoreServ Storage systems (see **MPX200 Multifunction Router with iSCSI for 3PAR StoreServ Storage** on page 366) 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 HPE 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.



https://t.me/learningnets

## 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 127: MPX200-XP single-blade fabric-attached configuration** on page 371 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.

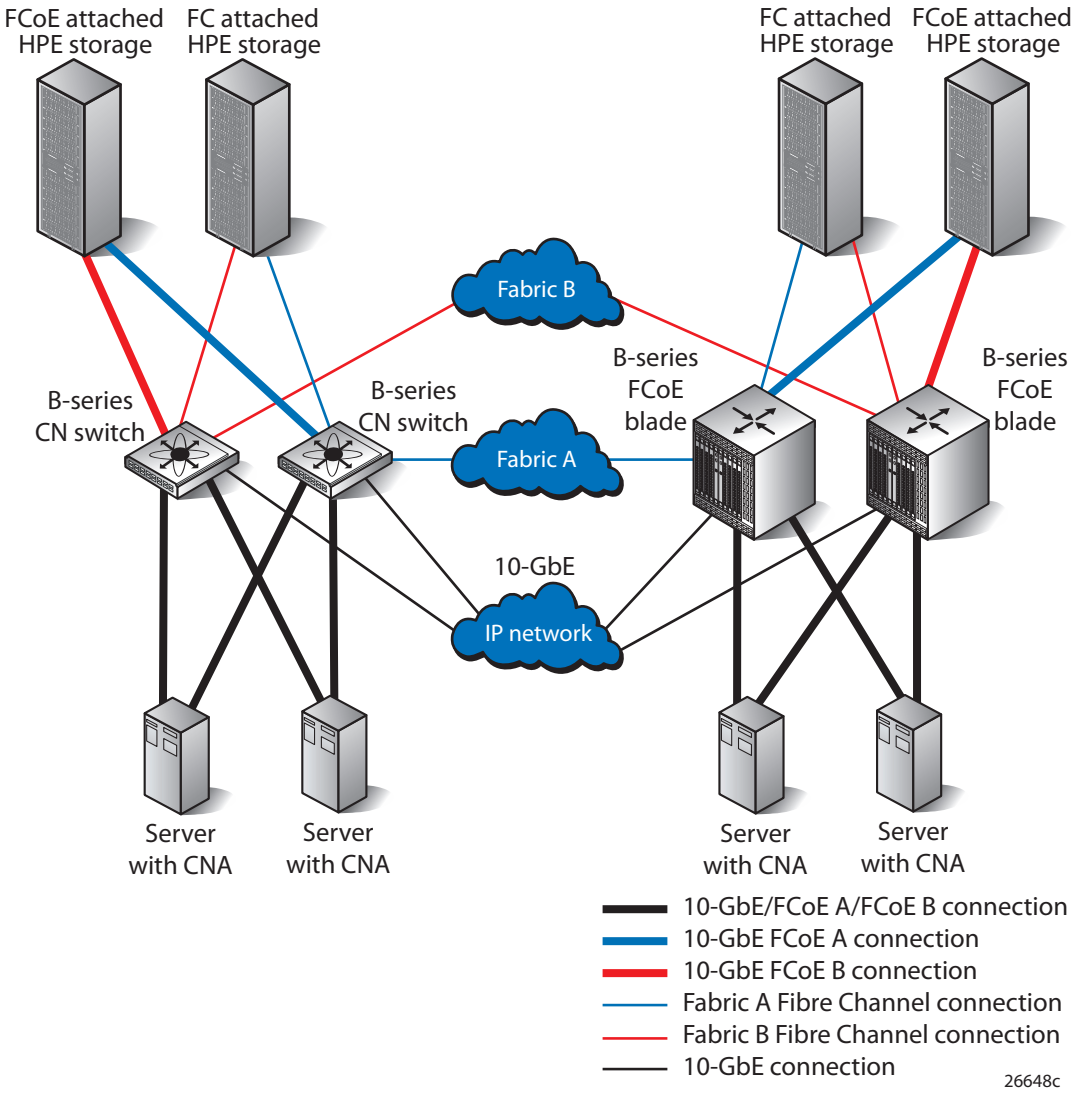



**Figure 127: MPX200-XP single-blade fabric-attached configuration**

**Figure 128: MPX200-XP dual-blade fabric-attached configuration** on page 372 shows an MPX200-XP dual-blade fabric-attached configuration. This configuration provides high availability with failover between blades.



https://t.me/learningnets

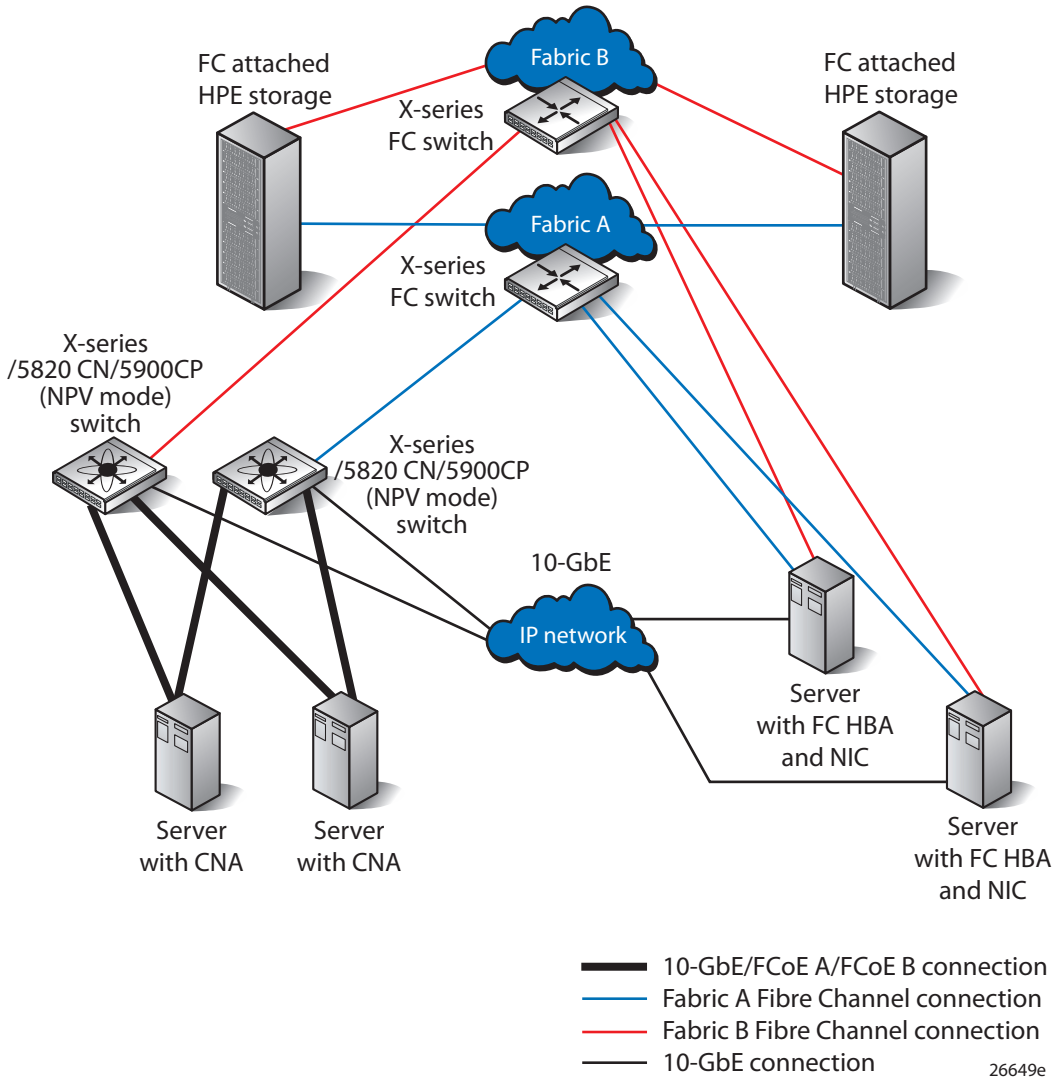



**Figure 128: MPX200-XP dual-blade fabric-attached configuration**

**Figure 129: MPX200-XP multi-XP fabric-attached configuration** on page 373 shows a multi-XP configuration with connectivity for up to four XP storage systems from a single MPX200 blade.



https://t.me/learningnets

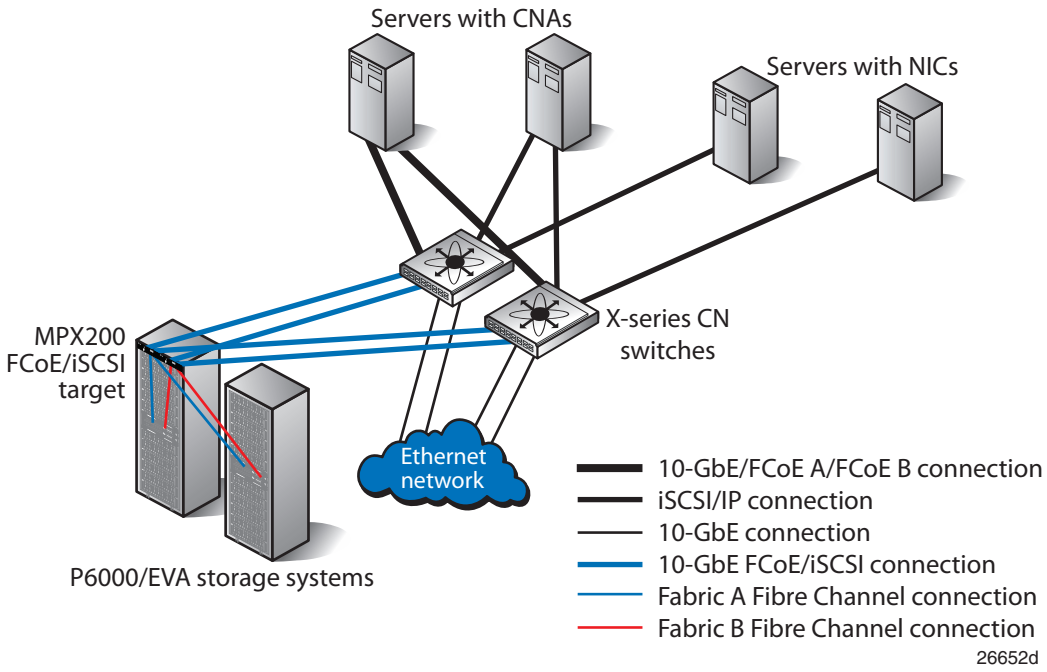



**Figure 129: MPX200-XP multi-XP fabric-attached configuration**

## 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.

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



https://t.me/learningnets

### 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.hpe.com/storage/spock**. You must sign up for an HP Passport to enable access.

#### Table 170: MPX200-XP operating system and multipath support

| Operating system | Multipath software | XP storage system |
|---|---|---|
| Microsoft Windows Server 2012, 2008, 2003 | MPIO with Microsoft DSM | XP24000/20000 |

## 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** on page 374
- **Hardware support** on page 375
- **Software support** on page 375
- **Scalability rules** on page 376


For the latest B-series documentation, see **http://www.hpe.com/info/StoreFabric**.

### 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 130: Fibre Channel and IP configuration with the B-series iSCSI Director Blade** on page 375 shows a sample configuration in which a B-series iSCSI Director Blade bridges an IP network and Fibre Channel network.



https://t.me/learningnets

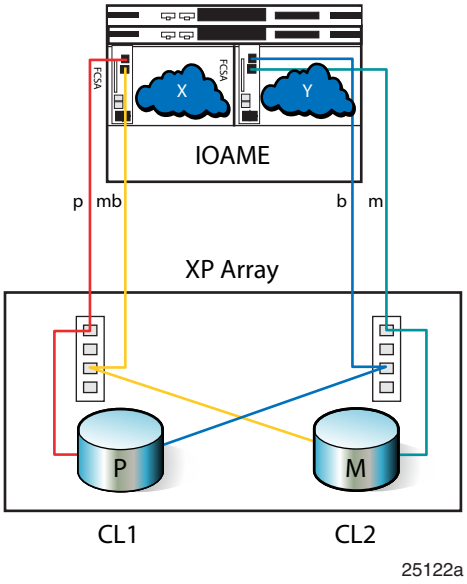



**Figure 130: Fibre Channel and IP configuration with the B-series iSCSI Director Blade**

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 system is supported with the B-series iSCSI Director Blade. Contact a Hewlett Packard Enterprise storage representative for specific support information.

XP10000/12000

**NOTE:**

Hewlett Packard Enterprise 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** on page 105.

## 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 Hewlett Packard Enterprise-supported NICs are also supported.



https://t.me/learningnets

- 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 a Hewlett Packard Enterprise storage representative for specific support information.

### Scalability rules

#### Table 171: 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 |

*Table Continued*



https://t.me/learningnets

| Rule | Maximum |
|---|---|
| iSCSI sessions per fabric | 4,096 |
| TCP connections per fabric | 4,096 |
| iSCSI targets per fabric | 4,096 |
| 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 |

# C-series iSCSI

The C-series IP Storage Services Modules, 18/4 Multiservice Module, 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, 18/4 Multiservice Module, 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** on page 377
- **Hardware support** on page 378
- **Software support** on page 379
- **Configuration rules** on page 380


For the latest C-series documentation, see **http://www.hpe.com/info/StoreFabric**.

## Modules overview

The C-series IP Storage Services Modules (IPS-4, IPS-8), 14/2 Multiprotocol Services Module, 18/4 Multiservice Module, 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 131: C-series Fibre Channel and IP configuration with the IP module** on page 378 shows a sample configuration where an IP module bridges an IP network and a Fibre Channel network.



https://t.me/learningnets

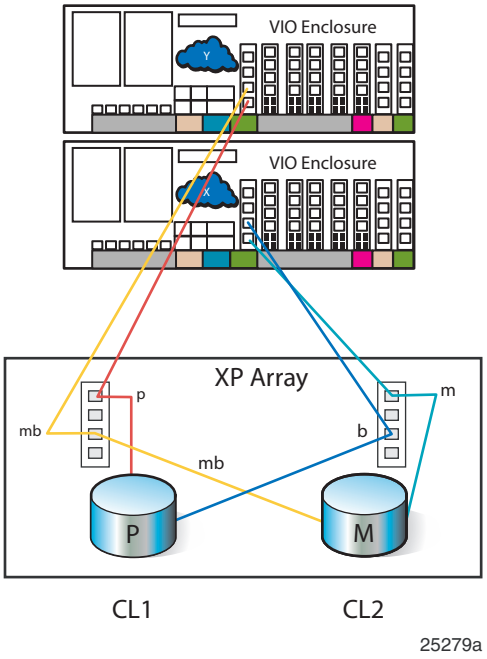



**Figure 131: C-series Fibre Channel and IP configuration with the IP module**

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, 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 a Hewlett Packard Enterprise storage representative for specific support information.

The following storage systems are supported with iSCSI:

- XP10000/12000
- XP20000/24000

**NOTE:** Hewlett Packard Enterprise supports direct connection of storage arrays to the Fibre Channel ports on the 18/4 Multiservice Module, and MDS 9222i Multiservice Fabric. HPE 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 18/4 Multiservice Module, 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 C-series switches listed in **C-series switches and fabric rules** on page 137.



https://t.me/learningnets

**Table 172: C-series switches that support iSCSI**

| Switch | Maximum number of Fibre Channel ports | Supported IP modules |
|---|---|---|
| MDS 9222i Multiservice Fabric | 66 | IPS-8, 18/4 |
| SN8000C 6-Slot Supervisor 2A Director Switch | 192 | IPS-4, IPS-8, 14/2, and 18/4 |
| SN8000C 9-Slot Supervisor 2A Director Switch | 336 | IPS-4, IPS-8, 14/2, and 18/4 |
| SN8000C 13-Slot Supervisor 2A Fabric 2 Director Switch | 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), 18/4 Multiservice Module, 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 Hewlett Packard Enterprise-supported NICs are also supported with iSCSI.

- Red Hat Enterprise Linux AS, ES, WS
- WS2008
- WS2012
- SUSE Linux Enterprise Server
- IBM AIX
- Oracle Solaris
- VMware ESX

## HPE Network Teaming

Windows 2008 and Windows 2012 support HPE Network Teaming.

## C-series management applications

C-series management applications are as follows:

- Cisco Data Center Network Manager (DCNM)
- Cisco Fabric Manager
- Cisco Device Manager
- CLI

## iSCSI initiators

C-series modules support these iSCSI software initiators:



https://t.me/learningnets

- 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
- Oracle 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, and 18/4 Multiservice Modules in the following:

- SN8000C 6-Slot Supervisor 2A Director Switch, SN8000C 9-Slot Supervisor 2A Director Switch, SN8000C 13-Slot Supervisor 2A Fabric 2 Director Switch
- MDS 9222i 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. HPE Secure Path is not supported with iSCSI initiators. Contact a Hewlett Packard Enterprise 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, and MDS 9222i Multiservice Fabric switch are supported on fabrics constructed with C-series switches. See **C-series switches and fabric rules** on page 137 for the latest C-series fabric rules.

### Table 173: 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


https://t.me/learningnets

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.

# HPE ProLiant Storage Server iSCSI Feature Pack

This section describes the following topics:


- **Overview** on page 381
- **HPE ProLiant Storage Server iSCSI Feature Pack support** on page 381
- **HPE ProLiant Storage Server iSCSI license upgrade options** on page 382
- **Designing a Microsoft Exchange solution with iSCSI Feature Pack** on page 383
- **Sample iSCSI NAS Microsoft Exchange Server 2003 configuration** on page 385


## Overview

ProLiant Storage Servers can be configured as iSCSI targets using the ProLiant Storage Server iSCSI Feature Pack.

The 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 iSCSI Feature Pack (T3669A only) includes the ProLiant Application Storage Manager, which provides storage management for HPE NAS servers hosting Microsoft Exchange 2003 storage groups. It reduces the time and training required to set up and monitor email stores.

The 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://www.hpe.com/support/storageworks-discontinued**.

## HPE ProLiant Storage Server iSCSI Feature Pack support

This section describes support for the ProLiant Storage Server iSCSI Feature Pack.

## Hardware support

The following storage servers support the iSCSI Feature Pack:

- HPE ProLiant DL100 Storage Server
- HPE ProLiant ML110 Storage Server
- HPE ProLiant ML350 G4 Storage Server
- HPE ProLiant ML370 G4 Storage Server



https://t.me/learningnets

- HPE ProLiant DL380 G4 Storage Server (Base, External SCSI, and External SATA models, SAN Storage model Gateway Edition only)
- HPE StorageWorks NAS 500s
- HPE StorageWorks NAS 1200s
- HPE StorageWorks NAS 1500s
- HPE StorageWorks NAS 2000s
- HPE StorageWorks NAS 4000s (Gateway Edition only)

## Application support

The following host applications support the 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 Storage Server GUI.

## iSCSI Initiator support rules

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.

## HPE ProLiant Storage Server iSCSI license upgrade options

The ProLiant Storage Server iSCSI Feature Pack has three licensed options:

- **Snapshot** on page 383
- **Clustering** on page 383
- **Direct Backup** on page 383



https://t.me/learningnets

### Snapshot

The Snapshot option is an upgrade license for 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 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 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 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.

Hewlett Packard Enterprise 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.



https://t.me/learningnets

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

Hewlett Packard Enterprise 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 Storage Server with the 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

Hewlett Packard Enterprise recommends that you:



https://t.me/learningnets

- 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 a Hewlett Packard Enterprise-tested iSCSI NAS configuration. .

## HPE ProLiant DL380 G4 Storage Server configuration

With up to six internal drives and fault-tolerant options for fans and power supplies, an HPE ProLiant DL380 server acts as a high-capacity Exchange server. Other 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. Hewlett Packard Enterprise 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 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.



https://t.me/learningnets

The Exchange server accesses its database on the iSCSI NAS RAID 1+0 volume.



https://t.me/learningnets

# HPE SAN Product Programs and Applications

The following chapters describe the HPE Complete Program and the HPE Smart SAN application.


- **HPE Complete Program** on page 388
- **HPE Smart SAN for 3PAR Application** on page 389




https://t.me/learningnets

# HPE Complete Program

The HPE Complete program enables the purchase of best in class Third Party branded products with the added reliability of HPE interop assurance for a complete validated solution, all via one HPE purchase order. Through HPE Complete, HPE provides a broader portfolio and choice of products that complement and complete HPE Storage and Server Solutions.

**The HPE Complete Advantage**

- Buying Third Party branded products from HPE provides added confidence and reliability of HPE interop assurance when deploying into the HPE ecosystem.
- Direct Third party expertise provided for installation, support, and warranty of their products.
- Troubleshooting to isolate issue with respect to HPE product or Third Party branded product and partners with Third Party as needed on customer issue resolution.
- Completing solutions with Third Party branded products purchased through HPE Complete minimizes customer risk, complexity, and cost when deploying multi-vendor products and technologies into the HPE ecosystem.

**Quoting and Ordering Process**

All HPE Complete Third Party branded products have HPE Part Numbers, are on the HPE Price list and can be ordered utilizing standard quoting and ordering processes and tools such as Watson and SBW.

**Support and Warranty**

Direct Third Party expertise is provided for the installation, support, and warranty of their products. HPE ensures complete customer satisfaction by troubleshooting to isolate issue to HPE or Third Party branded product and partners with Third Party as needed on issue resolution.

**Resources**

For HPE Complete information, see the **HPE Complete Portal**.



https://t.me/learningnets

# HPE Smart SAN for 3PAR Application

HPE Smart SAN for 3PAR is a target driven orchestration application that provides automated management control of the SAN for all-flash, hybrid, and traditional storage deployments. HPE 3PAR Smart SAN reduces the number of tedious configuration process steps that a SAN administrator is required to perform from hours to minutes and is error free in making SAN infrastructure for all-Flash more SAN aware.

With Smart SAN for 3PAR 1.0, HPE introduced Fibre Channel Industry Association (FCIA) T-11 Target Driven Peer Zoning (TDPZ) for 3PAR StoreServ Storage and utilized standards based FDMI FC registration for gathering SAN configuration data from HPE StoreFabric (32Gb, 16Gb, and 8Gb) B-series FC switches and StoreFabric (32Gb, 16Gb, and 8Gb) FC HBAs. TDPZ reduces the complexity of manually performing Fibre Channel (FC) zoning tasks by automating these tasks through the HPE 3PAR Storage System. The HPE 3PAR StoreServ Storage System communicates with the SAN fabric switches and creates the required Target Driven Peer Zones. TDPZ utilizes FCIA standards-based peer zoning to reduce the number of steps required to provision new storage hosts by up to 90% and error-free.

Apart from HPE 3PAR StoreServ Management Console (SSMC) support for provisioning with Smart SAN for 3PAR, a few key points related to Smart SAN for 3PAR are:

- Smart SAN for 3PAR and HPE OneView 3.1 can co-exist when using HPE StoreFabric B-series FC Switches (Brocade fabrics) with BNA 14.0.2 or above.
- Both regular and TDPZ zones can co-exist in the same fabric. If devices are listed in multiple zones, whether regular zones and/or TDPZ zones, the devices will abide by the access rules as defined for the individual zones.

**Smart SAN for 3PAR resources**

- **HPE Smart SAN for 3PAR Solutions Brief**
- **Smart SAN for 3PAR Web Page**
- **Blog-Welcome to SmartTopia:**
- **Smart SAN for 3PAR QuickSpecs**
- **HPE Smart SAN for 3PAR WP**
- **Demartek 16Gbps FC with 8450**
- **Comware supports HPE Smart SAN for 3PAR**
- **HPE Smart SAN for 3PAR Users Guide**



https://t.me/learningnets

# Storage security, best practices, and support information

The following chapters describe storage security and SAN best practices:


- **Storage security** on page 391
- **Best practices** on page 407




https://t.me/learningnets

# Storage security

This chapter describes storage security best practices. It describes the following topics:


- **Storage security threats** on page 391
- **Storage security compliance** on page 393
- **Security technologies** on page 393
- **HPE security strategy** on page 395
- **Storage security best practices** on page 397
- **Assessing security risks** on page 398
- **HPE storage security solutions** on page 398


## 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.



https://t.me/learningnets

**Table 174: Security attack classes**

| Attack class | Description |
|---|---|
| Passive | Attacks that can disclose information to an attacker.<br>Passive attacks include:<br>• Analyzing traffic<br>• Monitoring unprotected communications<br>• Decrypting weakly encrypted traffic<br>• Capturing authentication information (passwords)<br>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.<br>Active attacks include:<br>• Attempting to circumvent or break protection features<br>• Introducing malicious code<br>• Stealing or modifying information<br>• Attacking a network backbone<br>• Exploiting in-transit information<br>• Penetrating an enclave<br>• Attacking when a remote user attempts to connect to an enclave |
| 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. |

*Table Continued*



https://t.me/learningnets

| Attack class | Description |
|---|---|
| Insider | Unauthorized attacks by an authorized user can be malicious or nonmalicious.<br>Malicious attackers can:<br>• Eavesdrop<br>• Steal or damage data<br>• Use data for fraudulent purposes<br>• Deny authorized users access<br>Nonmalicious attacks can result from:<br>• Carelessness<br>• Lack of knowledge<br>• Circumventing security for nonmalicious purposes to perform tasks |
| 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 175: U.S. and international security regulations

| U.S. federal and state regulations | International regulations |
|---|---|
| • Sarbanes-Oxley (SOX) Act of 2002<br>• Gramm-Leach-Bliley Act (GLBA) of 1999<br>• Securities and Exchange Commission Act (SEC) rules 17a-3 and 17a-4<br>• Department of Energy (DOE) 10 CFR 600.153 Retention and access requirements for records<br>• California Data Security Act (SB 1386/AB 1950)<br>• New York Regulation 173 Standards for safeguarding customer information | • European Union Data Protection Directive of 1995<br>• Canada: Personal Information Protection and Electronic Documents Act (PIPEDA)<br>• Australia: Privacy Act 1988<br>• Japan: Personal Information Protection Act<br>• UK: Data Protection Act 1998<br>• New Zealand: Privacy Act 1993 |

# Security technologies

This section describes security technologies for IP SAN, Fibre Channel SAN, and encryption.



https://t.me/learningnets

## 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:

- When the server is first connected, it sends a challenge message to the peer.
- The peer responds by sending a value generated by a one-way hash function.
- The server compares this value to its own generated value.
- If the values match, the connection is allowed to continue; if they do not match, the connection is terminated.
- 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



https://t.me/learningnets

## 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:

### Procedure

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.

# HPE security strategy

This section describes the HPE Secure Advantage, the HPE security strategy.



https://t.me/learningnets

# HPE Secure Advantage

Secure Advantage allows you to combine HPE 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. Hewlett Packard Enterprise 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 HPE 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



https://t.me/learningnets

#### 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:



https://t.me/learningnets

- **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 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.

# HPE storage security solutions

This section describes HPE storage security solutions for the following products:



https://t.me/learningnets

- **C-series Storage Media Encryption** on page 399
- **C-series NX-OS security** on page 399
- **C-series IP SAN security** on page 400
- **B-series Encryption Switch and Encryption FC Blade security** on page 401
- **B-series Fabric OS security** on page 402
- **Key management** on page 405


## 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.

## C-series NX-OS security

This section describes the C-series SAN-OS security features.



https://t.me/learningnets

#### 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:



https://t.me/learningnets

- **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** on page 105.

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*.

#### 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 HPE 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:



https://t.me/learningnets

- 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 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:** To ensure a secure network, avoid using Telnet or any other unprotected applications to communicate with switches.



https://t.me/learningnets

### 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 supports the following browsers for SSL connections:

- Internet Explorer (Microsoft Windows)
- Mozilla (Oracle 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.



https://t.me/learningnets

- 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 176: 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:

- 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.
- In the default configuration, FCAP authentication is tried first, then DHCHAP authentication. Each switch can be configured to negotiate one or both types.


https://t.me/learningnets

- 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** on page 134.

#### B-series IPsec

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 StoreFabric SN4000B SAN Extension Switch, 1606 Extension SAN Switch, DC Dir Switch MP Extension Blade, or MP Router Blade. IPsec does not require that you configure security for each application that uses TCP/IP. When configuring IPsec, you must ensure that a compatible StoreFabric SN4000B SAN Extension Switch, 1606 Extension SAN Switch, DC Dir Switch MP Extension Blade, or MP Router Blade with the same version of FOS is at each end of the FCIP tunnel.

  For compatible B-Series FCIP gateways, see Table HPE FCIP gateways.

- 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

StorageWorks Secure Key Manager for HPE LT04 tape libraries is part of the Secure Advantage solution.

Secure Key Manager features include:

- Centralized encryption key management for LT04 tape libraries



https://t.me/learningnets

- 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 HPE 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.



https://t.me/learningnets

# Best practices

This chapter describes HPE best practices for SAN design and implementation. It describes the following topics:


- **SAN planning** on page 407
- **Design specification** on page 408
- **SAN topology** on page 408
- **SAN configuration** on page 409
- **Storage-based LUN masking** on page 411
- **Zoning** on page 411
- **FCoE switch configuration quick-setup instructions** on page 415
- **SAN scaling** on page 426
- **SAN fabric merging** on page 427


For SAN design assistance, see the Storage Services website **http://www.hpe.com/support/HPEservices**.

## 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.



https://t.me/learningnets

# 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** on page 408
- **Failover protection** on page 409
- **Data access patterns** on page 409
- **ISL ratio** on page 409
- **Incremental SAN expansion** on page 409


## Multi-fabric SANs

Hewlett Packard Enterprise 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, Hewlett Packard Enterprise recommends that you utilize intelligent fabric monitoring tools such as HPE Intelligent Infrastructure Analyzer Software (IIAS),**SAN infrastructure monitoring** on page 430, 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, Hewlett Packard Enterprise recommends that you implement proper cable management practices and end-to-end SAN monitoring.



https://t.me/learningnets

## 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** on page 40.

## 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 Table Data access performance by SAN fabric topology.

## 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 **Core-edge fabric types** on page 32.

## 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**



https://t.me/learningnets

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. HPE 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 Table Rules for fiber optic cable connections. 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** on page 410
- **Server setup** on page 410
- **Storage system configuration** on page 410


## Fibre Channel switch configuration

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 HPE-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:



https://t.me/learningnets

- 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 EVA storage systems, use SSP.
- For HPE XP storage systems, use LUN Configuration and Security Manager XP.

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** on page 411
- **Zoning guidelines** on page 412
- **EBS zoning** on page 414
- **Zone naming** on page 414


## 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



https://t.me/learningnets

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:

- Zoning enforcement for B-series Fibre Channel switches and MP Router LSANs (B-series)
- Zoning enforcement for C-series Fibre Channel switches (C-series)
- Zoning enforcement for H-series switches (H-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
- HBA port
- NPIV port
- 3PAR persistent ports
- 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:

- **Zoning limits and enforcement** on page 133

  (B-series)
- **Zoning limits and enforcement** on page 149

  (C-series)
- **Zoning limits and enforcement** on page 160



https://t.me/learningnets

(H-series)

- **Common server access, different storage system types** on page 224

## 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, Hewlett Packard Enterprise 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** on page 225.

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 HBA port

Zoning by HBA port applies when you are utilizing dual-ported HBAs. From a zoning perspective, you can view each port as if it were a separate HBA. As such, you would use the same criteria as described above in the Zoning by HBA section. In this case however, each HBA port should be thought of as a separate HBA.

## Zoning by NPIV port

With NPIV, one physical link is shared by multiple virtual ports and each is assigned a different WWPN. Similar to zoning by HBA port, from a zoning perspective, each virtual port should be viewed as if it were a separate HBA.

## Zoning with 3PAR persistent ports

3PAR Persistent Ports technology allows for a completely non-disruptive software upgrade environment (from the host pathing point of view) where host-based multipathing software will not be involved in the software upgrade process. Additionally, 3PAR Persistent Ports technology renders an array node failure transparent to hosts using the array, avoiding the need for the multipathing software of the host to maintain host connectivity for node failure recovery. For information about zoning when using this feature, see the HPE 3PAR Persistent Ports technical whitepaper (4AA4-4545ENW), available on the Hewlett Packard Enterprise storage website: **http://www.hpe.com/info/StoreFabric**. The whitepaper has a separate section on "Best practice and zoning and multipathing considerations".

## 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.



https://t.me/learningnets

# EBS zoning

For EBS zoning recommendations, see the *HP Enterprise Backup Solution Design Guide*, available at **http://www.hpe.com/info/ebs**.

# 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.



https://t.me/learningnets

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.

## HPE 5820 FCoE Converged Network Switch quick setup

This procedure is intended for users who are familiar with HPE A-series switches. Use this procedure to enable servers with CNAs attached to the HPE 5820 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 A-series switches, refer to the 5820 product documentation for detailed instructions on setting up the switch.



https://t.me/learningnets

**Procedure**

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).

   a. 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
```



https://t.me/learningnets

```
[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
```

**b.** 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
```

**c.** Internal 5820X 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 5820 FCoE module.

**a.** Connect to FCoE module from 5820 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 5820 ToSAN VLAN.

```
H3C (admin-config) #> vlan 4001 create
H3C (admin-config) #> vlan 4001 add port 4-7
```



https://t.me/learningnets

**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 5820 prompt.

```
H3C #> quit
```

Press **CTRL+K** to return to 5820 prompt.

```
<switchname>
```

## 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 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.

**Procedure**

**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
```



https://t.me/learningnets

```
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 converged[1]
    switch(config-if-te-0/0)# switchport converged allowed vlan add
5[1]      FCoE Required (Step 3 VLAN)
    switch(config-if-te-0/0)# vlan classifier activate group 1 vlan 5[1]
    switch(config-if-te-0/0)# cee default[1]
    switch(config-if-te-0/0)# exit
```

5. Repeat **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 **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
=========                   ==========       ======
```

[1] 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.



https://t.me/learningnets

```
========
TenGigabitEthernet 0/0     unassigned       up                    up

switch#show vlan brief

VLAN     Name             State   Ports
                                   (u)-Untagged, (t)-Tagged
======= ================ ======= ================================
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)
```



https://t.me/learningnets

**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>
```

## HPE C-series Nexus 5548UP/5596UP 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 HPE C-series Nexus 5548UP/5596UP 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.

### Table 177: Recommended VFC port assignments

| Model | Slot | Number of FCoE ports (maximum) | VFC port assignments (recommended) |
|---|---|---|---|
| Cisco Nexus 5548UP Converged Network Switch | 1 (Standard) | 32 | VFC ports 1 through 32 |
| | 2 (Optional expansion module slot[1]) | 16 | VFC ports 33 through 48 |
| Cisco Nexus 5596UP Converged Network Switch | 1 (Standard) | 48 | VFC ports 1 through 48 |
| | 2 (Optional expansion module slots[1]) | 16 | VFC ports 49 through 64 |

*Table Continued*



https://t.me/learningnets

| Model | Slot | Number of FCoE ports (maximum) | VFC port assignments (recommended) |
|---|---|---|---|
| | 3 (Optional expansion module slots[1]) | 16 | VFC ports 65 through 80 |
| | 4 (Optional expansion module slots[1]) | 16 | VFC ports 81 through 96 |

[1] 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 5548UP/5596UP 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.

To establish CNA connectivity and enable login to the C-series Nexus 5000 Converged Network Switch, configure the IEEE DCB ports as follows:

**Procedure**

**1.** Enable FCoE (disabled by default).

**NOTE:** The C-series Nexus 5500 Converged Network Switch will require a reload (reboot).

```
Nexus5548# configure terminal
Nexus5548(config)# feature fcoe
Nexus5548(config)# feature npiv
Nexus5548(config)# 2009 Apr  1 12:05:06 Nexus5548 %$ VDC-1 %$ %PLATFORM-2-
FC_LICENSE_DESIRED: FCoE/FC feature will be enabled after the
configuration
is saved followed by a reboot
Nexus5548(config)# exit
Nexus5548# 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 **3**.

```
Nexus5548#
                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.



https://t.me/learningnets

**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.

```
Nexus5548# configure terminal
Nexus5548(config)# vlan 200
Nexus5548(config-vlan)# exit
Nexus5548(config)# exit
Nexus5548# 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

Nexus5548# configure terminal
Nexus5548(config)# interface ethernet 1/1-20
Nexus5548(config-if-range)# switchport mode trunk
Nexus5548(config-if-range)# switchport trunk allowed vlan 1, 200[2]
Nexus5548(config-if-range)# interface vfc 1-20
Nexus5548(config-if-range)# exit
```

4. Create a new VSAN that includes the FC and VFC ports.

   By default, all ports are in VSAN 1. Hewlett Packard Enterprise 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.

```
Nexus5548# configure terminal
Nexus5548(config)# vsan database
Nexus5548(config-vsan-db)# vsan 2
Nexus5548(config-vsan-db)# vsan 2 interface fc2/1-8
Nexus5548(config-vsan-db)# vsan 2 interface vfc 1-20
Nexus5548(config-vsan-db)# exit
Nexus5548(config)# exit
Nexus5548# 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
```

[2] 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.



https://t.me/learningnets

```
vsan 4094(isolated_vsan) interfaces:
```

5. Associate the VLAN with the VSAN.

   In this example, VLAN 200 is associated with VSAN 2.

```
Nexus5548# configure terminal
Nexus5548(config)# vlan 200
Nexus5548(config-vlan)# fcoe vsan 2
Nexus5548(config-vlan)# exit
Nexus5548(config)# exit
Nexus5548# 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

```
Nexus5548# configure terminal
Nexus5548(config)# interface vfc 1
Nexus5548(config-if)# bind interface ethernet 1/1
Nexus5548(config-if)# exit
Nexus5548(config)# interface vfc 2
Nexus5548(config-if)# bind interface ethernet 1/2
Nexus5548(config-if)# exit
Nexus5548(config)# interface vfc 3
Nexus5548(config-if)# bind interface ethernet 1/3
Nexus5548(config-if)# exit
Nexus5548(config)# interface vfc 4
Nexus5548(config-if)# bind interface ethernet 1/4
Nexus5548(config-if)# exit
Nexus5548(config)# interface vfc 5
Nexus5548(config-if)# bind interface ethernet 1/5
Nexus5548(config-if)# exit
Nexus5548(config)# interface vfc 6
Nexus5548(config-if)# bind interface ethernet 1/6
Nexus5548(config-if)# exit
Nexus5548(config)# interface vfc 7
Nexus5548(config-if)# bind interface ethernet 1/7
Nexus5548(config-if)# exit
Nexus5548(config)# interface vfc 8
Nexus5548(config-if)# bind interface ethernet 1/8
```



https://t.me/learningnets

```
Nexus5548(config-if)# exit
Nexus5548(config)# interface vfc 1-8
Nexus5548(config-if)# no shutdown
Nexus5548(config-if)# exit
```

**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

```
Nexus5548# configure terminal
Nexus5548(config)# interface vfc 1
Nexus5548(config-if)# bind interface ethernet 1/1
Nexus5548(config-if)# exit
Nexus5548(config)# interface vfc 2
Nexus5548(config-if)# bind interface ethernet 1/2
Nexus5548(config-if)# exit
Nexus5548(config)# interface vfc 3
Nexus5548(config-if)# bind interface ethernet 1/3
Nexus5548(config-if)# exit
Nexus5548(config)# interface vfc 4
Nexus5548(config-if)# bind interface ethernet 1/4
Nexus5548(config-if)# exit
Nexus5548(config)# interface vfc 5
Nexus5548(config-if)# bind interface ethernet 1/5
Nexus5548(config-if)# exit
Nexus5548(config)# interface vfc 6
Nexus5548(config-if)# bind interface ethernet 1/6
Nexus5548(config-if)# exit
Nexus5548(config)# interface vfc 10
Nexus5548(config-if)# bind interface ethernet 1/10
Nexus5548(config-if)# exit
Nexus5548(config)# interface vfc 20
Nexus5548(config-if)# bind interface ethernet 1/20
Nexus5548(config-if)# exit
Nexus5548(config)# interface vfc 21
Nexus5548(config-if)# bind interface ethernet 1/21
Nexus5548(config-if)# exit
Nexus5548(config)# interface vfc 22
Nexus5548(config-if)# bind interface ethernet 1/22
Nexus5548(config-if)# exit
Nexus5548(config)# interface vfc 23
Nexus5548(config-if)# bind interface ethernet 1/23
Nexus5548(config-if)# exit
Nexus5548(config)# interface vfc 24
Nexus5548(config-if)# bind interface ethernet 1/24
Nexus5548(config-if)# exit
Nexus5548(config)# interface vfc 25
Nexus5548(config-if)# bind interface ethernet 1/25
Nexus5548(config-if)# exit
Nexus5548(config)# interface vfc 1-6
Nexus5548(config-if)# no shutdown
```



https://t.me/learningnets

```
Nexus5548(config-if)# exit
Nexus5548(config)# interface vfc 10
Nexus5548(config-if)# no shutdown
Nexus5548(config-if)# exit
Nexus5548(config)# interface vfc 20-25
Nexus5548(config-if)# no shutdown
Nexus5548(config-if)# exit
```

7. Enable the FC ports.

   In this example, FC ports 2/1 through 2/8 are enabled.

```
Nexus5548(config)# interface fc 2/1-8
Nexus5548(config-if)# no shutdown
Nexus5548(config-if)# exit
Nexus5548(config)# exit
```

8. Copy the running configuration to the startup configuration location.

```
Nexus5548# copy running-config startup-config
[########################################] 100%
```

9. Copy the running configuration to a backup location.

```
Nexus5548# copy running-config ftp://10.10.20.1/backup.txt
```

10. Verify the configuration.

```
Nexus5548# show interface brief

Nexus5548# 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** on page 24.

**CAUTION:** Hewlett Packard Enterprise 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:



https://t.me/learningnets

- **Fabric rules** on page 156, for H-series switches
- **Fibre Channel switch fabric rules** on page 118, for B-series switches
- **Fibre Channel switch fabric rules** on page 145, for C-series switches

- Add another fabric as a high-availability NSPOF solution.
- Implement Fibre Channel routing, see **Fibre Channel routing** on page 46.
- Deploy multiple independent SANs.
- Migrate to a different SAN topology, see **Topology migration** on page 44.

When adding new switches, avoid fabric segmentation. For more information, see **Fabric segmentation errors** on page 428.

## 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.



https://t.me/learningnets

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** on page 428.

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, Hewlett Packard Enterprise strongly recommends that the fabric be quiesced. Zone membership should not be changed for devices that are actively performing I/O in a fabric. After 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** on page 428
- **Switch configuration parameters** on page 429
- **Independent fabric merge** on page 429
- **High-availability redundant fabric merge** on page 429


## 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.



https://t.me/learningnets

## 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.

- Identify and resolve any issues that can cause fabric segmentation.
- Verify that each fabric provides a redundant path to all attached devices.
- Verify that paths are open to each device that must remain online during the merge.
- Select fabrics for merging, for example, fabric A with fabric C.
- 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.
- Verify that the fabric selected for merging has no I/O activity.
- Connect the selected fabrics (fabric A and fabric C).
- Verify that the newly merged fabric contains all switches and that the zoning has merged correctly.
- Restore I/O operations on the new fabric from the multipathing software console.
- Verify that paths are open and restored for each device.
- Ensure that all paths and I/O operations have been restored.



https://t.me/learningnets

## SAN infrastructure monitoring

Hewlett Packard Enterprise recommends use of the Intelligent Infrastructure Analyzer Software (IIAS) to monitor and diagnose the physical layer of SANs in real-time, with an emphasis on the SFP transceivers utilized in Fibre Channel switches. The solution is highly beneficial for enterprise SANs, where SAN management can be cumbersome and tedious. IIAS is intended for customers who utilize HPE storage hardware in their SAN.

HPE IIAS uses industry-standard protocols (such as SNMP, SMI-S, Telnet) to monitor the physical layer (SFP transceivers) of a SAN. It diagnoses changes/events in SFP states and characteristics, and presents SAN topology, inventory, and diagnostic information to the user in real-time. IIAS enables you to:

- Discover and collect data for B-Series and H-Series switches in a SAN
- Periodically monitor a SAN at configured intervals using an active profile
- Monitor or diagnose degrading or failing SFPs in an active profile
- Generate current and historical reports
- Notify the user of any change in the topology or SAN component state
- Provide a SAN hardware summary in terms of component inventory

For additional information about IIAS, see the IIAS website at **http://www.hpe.com/info/iias**.



https://t.me/learningnets

# Support and other resources

## Accessing Hewlett Packard Enterprise Support

- For live assistance, go to the Contact Hewlett Packard Enterprise Worldwide website:

  **http://www.hpe.com/assistance**

- To access documentation and support services, go to the Hewlett Packard Enterprise Support Center website:

  **http://www.hpe.com/support/hpesc**

**Information to collect**

- Technical support registration number (if applicable)
- Product name, model or version, and serial number
- Operating system name and version
- Firmware version
- Error messages
- Product-specific reports and logs
- Add-on products or components
- Third-party products or components

## Accessing updates

- Some software products provide a mechanism for accessing software updates through the product interface. Review your product documentation to identify the recommended software update method.
- To download product updates:

  **Hewlett Packard Enterprise Support Center**

  **www.hpe.com/support/hpesc**

  **Hewlett Packard Enterprise Support Center: Software downloads**

  **www.hpe.com/support/downloads**

  **Software Depot**

  **www.hpe.com/support/softwaredepot**

- To subscribe to eNewsletters and alerts:

  **www.hpe.com/support/e-updates**

- To view and update your entitlements, and to link your contracts and warranties with your profile, go to the Hewlett Packard Enterprise Support Center **More Information on Access to Support Materials** page:



https://t.me/learningnets

**www.hpe.com/support/AccessToSupportMaterials**

**IMPORTANT:** Access to some updates might require product entitlement when accessed through the Hewlett Packard Enterprise Support Center. You must have an HPE Passport set up with relevant entitlements.

# Customer self repair

Hewlett Packard Enterprise customer self repair (CSR) programs allow you to repair your product. If a CSR part needs to be replaced, it will be shipped directly to you so that you can install it at your convenience. Some parts do not qualify for CSR. Your Hewlett Packard Enterprise authorized service provider will determine whether a repair can be accomplished by CSR.

For more information about CSR, contact your local service provider or go to the CSR website:

**http://www.hpe.com/support/selfrepair**

# Remote support

Remote support is available with supported devices as part of your warranty or contractual support agreement. It provides intelligent event diagnosis, and automatic, secure submission of hardware event notifications to Hewlett Packard Enterprise, which will initiate a fast and accurate resolution based on your product's service level. Hewlett Packard Enterprise strongly recommends that you register your device for remote support.

If your product includes additional remote support details, use search to locate that information.

**Remote support and Proactive Care information**

**HPE Get Connected**

**www.hpe.com/services/getconnected**

**HPE Proactive Care services**

**www.hpe.com/services/proactivecare**

**HPE Proactive Care service: Supported products list**

**www.hpe.com/services/proactivecaresupportedproducts**

**HPE Proactive Care advanced service: Supported products list**

**www.hpe.com/services/proactivecareadvancedsupportedproducts**

**Proactive Care customer information**

**Proactive Care central**

**www.hpe.com/services/proactivecarecentral**

**Proactive Care service activation**

**www.hpe.com/services/proactivecarecentralgetstarted**

# Warranty information

To view the warranty information for your product, see the links provided below:

**HPE ProLiant and IA-32 Servers and Options**

**www.hpe.com/support/ProLiantServers-Warranties**

**HPE Enterprise and Cloudline Servers**

**www.hpe.com/support/EnterpriseServers-Warranties**



https://t.me/learningnets

**HPE Storage Products**

**www.hpe.com/support/Storage-Warranties**

**HPE Networking Products**

**www.hpe.com/support/Networking-Warranties**

# Regulatory information

To view the regulatory information for your product, view the *Safety and Compliance Information for Server, Storage, Power, Networking, and Rack Products*, available at the Hewlett Packard Enterprise Support Center:

**www.hpe.com/support/Safety-Compliance-EnterpriseProducts**

### Additional regulatory information

Hewlett Packard Enterprise is committed to providing our customers with information about the chemical substances in our products as needed to comply with legal requirements such as REACH (Regulation EC No 1907/2006 of the European Parliament and the Council). A chemical information report for this product can be found at:

**www.hpe.com/info/reach**

For Hewlett Packard Enterprise product environmental and safety information and compliance data, including RoHS and REACH, see:

**www.hpe.com/info/ecodata**

For Hewlett Packard Enterprise environmental information, including company programs, product recycling, and energy efficiency, see:

**www.hpe.com/info/environment**

# Documentation feedback

Hewlett Packard Enterprise is committed to providing documentation that meets your needs. To help us improve the documentation, send any errors, suggestions, or comments to Documentation Feedback (**docsfeedback@hpe.com**). When submitting your feedback, include the document title, part number, edition, and publication date located on the front cover of the document. For online help content, include the product name, product version, help edition, and publication date located on the legal notices page.



https://t.me/learningnets

# Related information

**Table 178: Related documentation**

| Topic | Information source |
|---|---|
| Switches | For the latest information on B-series, C-series, and H-series switches and firmware versions, see the SAN Infrastructure website:<br>**http://www.hpe.com/info/StoreFabric** |
| Fabric interoperability | *HPE StorageWorks Fabric Interoperability: Merging Fabrics Based on C-series and B-series Fibre Channel Switches Application Notes*<br>See this document on the SAN Infrastructure website:<br>**http://www.hpe.com/info/StoreFabric** |
| HPE Data Availability, Protection and Retention | The HPE solution is the *HPE Data Availability, Protection and Retention Compatibility Matrix* (**www.hpe.com/storage/daprcompatibility**). |
| Storage arrays | For the latest information on storage arrays, see the Disk Storage Systems website:<br>**http://www.hpe.com/storage/disk-storage** |



https://t.me/learningnets