

HP StorageWorks

XPath OS 7.4.x administrator guide

Legal and notice information

© Copyright 2005 Hewlett-Packard Development Company, L.P.

© Copyright 2005 Brocade Communications Systems, Incorporated.

Hewlett-Packard Company makes no warranty of any kind with regard to this material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

This document contains proprietary information, which is protected by copyright. No part of this document may be photocopied, reproduced, or translated into another language without the prior written consent of Hewlett-Packard. The information is provided "as is" without warranty of any kind and is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft® is a U.S. registered trademark of Microsoft Corporation.

UNIX is a registered trademark of The Open Group.

Linux is a U.S. registered trademark of Linus Torvalds.

Java™ is a U.S. trademark of Sun Microsystems, Inc.

XPath OS 7.4.x administrator guide

Contents

About this guide	7
Intended audience	7
Related documentation	7
Document conventions and symbols	8
Rack stability	8
HP technical support	9
HP-authorized reseller.	9
Helpful web sites	9
1 Introducing XPath OS features	11
MP Routing Services	11
FC-FC routing service	11
FCIP Tunneling Service	12
iSCSI Gateway Service	12
Combining services	12
Fibre Channel features	13
HP Fabric OS support	14
PID mode requirements.	15
Feature compatibility	15
Daemon overseer service.	15
High availability	17
2 Performing basic configuration.	19
Viewing router information.	19
Verifying the PID mode	20
Creating interswitch links.	21
Configuring a long-distance connection.	22
Verifying connectivity	23
Synchronizing time with an NTP server	23
Licensed features	24
Changing account passwords	24
Enabling and disabling switches.	25
Enabling and disabling ports	25
Activating Ports on Demand: upgrading an 8-port base model to a 16-port full model.	26
Setting the domain ID	27
Controlling routing within a fabric	28
Specifying frame delivery order.	28
Using dynamic load sharing	28
Viewing routing information	29
Displaying command help	31
3 Performing basic maintenance	33
Maintaining the router configuration	33
Maintaining firmware	35
Performing hardware checks	41
Viewing hardware status	41
Modifying the power supply status threshold	43
4 Using the FC-FC Routing Service.	45
About Fibre Channel routing	45
Proxy devices	47
LSANs and zoning.	48
Fibre Channel NAT and phantom domains	50
Connecting multiple EX_Ports to an edge fabric	51

Matching fabric parameters	51
SAN scalability	51
Configuring an interfabric link	52
XPath OS and Secure Fabric OS	54
Configuring a secure XPath OS DH-CHAP secret	54
Setting a proxy PID	57
Monitoring resources	58
Routing ECHO	58
Connecting to McDATA SANs	59
Connectivity features	59
Interconnectivity benefits	59
Connectivity	59
Scalability	59
Supported modes	59
Configuring the fabrics for interconnectivity	61
Configuring the MP Router	61
Preparing the HP StorageWorks switch for connectivity	63
Configuring the McDATA fabric for interconnection	64
LSAN zoning	67
Completing the configuration	68
5 Using the FCIP Tunneling Service	71
Synchronizing time	71
Configuring an FCIP interswitch link	72
Disabling and enabling an FCIP interswitch link	76
6 Using the iSCSI Gateway Service	77
Summary of configuration steps	78
Configuring an iSCSI portal	78
Configuring iSCSI gateway zones	79
Configuring CHAP	80
Administering iSCSI configurations	80
Enabling and disabling iFCS	81
Displaying iFCS information	81
iFCS behavior during configuration download	81
Working with the WWN mapping table	82
Enabling and disabling failover	82
7 Creating and maintaining zones	85
Zoning terminology	86
Zoning enforcement	86
Soft zoning	86
Hard zoning	86
Zone server compatibility	86
Standards and zoning compatibility	87
Configuring zones	87
Mapping iSCSI names	88
Zoning commands	89
8 Using ISL trunking	91
How exchange-based trunking works	91
Enabling trunking	92
Managing trunking	92
Trunking commands	93
9 Monitoring system logs	95
System error log	95
Message severity levels	95
Viewing the system error log	96
Sample system error log message	96

Clearing the system error log	97
Port log	97
Port log management	97
Sample port log	98
Using the syslog daemon	98
XPath OS syslogd CLI commands	98
Enabling syslogd	99
Disabling syslogd	99

A Hard zoning background	101
---	------------

B Recovery kernel for XPath OS 7.4.x	103
---	------------

Software installation support environment	103
Using the recovery kernel	103

Glossary	105
---------------------------	------------

Index	119
------------------------	------------

Figures

1 Simple FCIP tunneling configuration	12
2 Combining FCIP tunneling and FC-FC routing services	13
3 Simple meta-SAN	46
4 Meta-SAN with interfabric links	46
5 Edge SANs connected through a backbone fabric	47
6 Proxy topology	48
7 SAN Pilot	63
8 SAN Pilot and EFCM zones	65
9 Modify Zone tab	65
10 World Wide Name box	66
11 Modify zone window	66
12 Activate zone set	67
13 Network using FCIP	71
14 Example for configuration procedures	72
15 Simple application of the iSCSI Gateway Service	77
16 iSCSI gateway configuration	78
17 iSCSI high availability configuration	83
18 Sample zone configuration	85
19 Sample exchange-based trunking configuration	92
20 Host and target directly connected to the MP Router	101
21 Host and target in a homogeneous fabric	101
22 Host and target in a heterogeneous fabric, combination 1	102
23 Host and target in a heterogeneous fabric, combination 2	102

Tables

1 Document conventions	8
2 XPath OS Fibre Channel switch features	13
3 Monitored XPath OS daemons	16
4 ISL modes	22
5 Component status rules	41
6 portCfgExPort -m interop parameters	60
7 iFCS information	81
8 Effect of downloading a configuration file with iFCS enabled or disabled	81
9 Zoning commands	89
10 Trunking commands	93
11 Managing the system error log	95
12 Message severity levels	96
13 System error log message field descriptions	97
14 Port log management commands	98
15 syslogd configuration commands	99

About this guide

This document provides procedures for Storage Area Network (SAN) administrators to set up and manage HP StorageWorks SANs. It is specific to XPath Operating System 7.4.x and the Multi-protocol (MP) Router running XPath OS 7.4.x

Intended audience

This guide is intended for system administrators and technicians who are experienced with the following:

- HP StorageWorks Fibre Channel SAN switches
- XPath OS 7.4.x or earlier

Related documentation

Documentation, including white papers and best practices documents, is available on the HP web site:

<http://www.hp.com/country/us/eng/prodserv/storage.html>

 **IMPORTANT:** For late-breaking, supplemental information, access the latest version of the *HP StorageWorks XPath OS 7.4.x release notes* using the following steps.

To access current XPath OS documents:

1. Locate the **IT storage products** section of the web page.
2. Under **Networked storage**, click **SAN infrastructure**.
3. From the **SAN Infrastructure** web page, locate the **SAN Infrastructure products** section.
4. Click **Multi-protocol Routers and Gateways**.
5. To access XPath OS 7.4.x documents (such as this document), click **B-Series Multi-Protocol Router**.
The **HP StorageWorks B-Series Multi-Protocol Router** overview page is displayed.
6. Go to the **Product Information** section, located on the right side of the web page.
7. Click **Technical documentation**.
8. Follow the onscreen instructions to download XPath OS 7.4.x documents.

Document conventions and symbols

Table 1 Document conventions

Convention	Element
Medium blue text: Figure 1	Cross-reference links and e-mail addresses
Medium blue, underlined text (http://www.hp.com)	Web site addresses
Bold font	<ul style="list-style-type: none">• Key names• Text typed into a GUI element, such as into a box• GUI elements that are clicked or selected, such as menu and list items, buttons, and check boxes
<i>Italics font</i>	Text emphasis
Monospace font	<ul style="list-style-type: none">• File and directory names• System output• Code• Text typed at the command line
<i>Monospace, italic font</i>	<ul style="list-style-type: none">• Code variables• Command-line variables
Monospace, bold font	Emphasis of file and directory names, system output, code, and text typed at the command line

 **WARNING!** Indicates that failure to follow directions could result in bodily harm or death.

 **CAUTION:** Indicates that failure to follow directions could result in damage to equipment or data.

 **IMPORTANT:** Provides clarifying information or specific instructions.

 **NOTE:** Provides additional information.

Rack stability

 **WARNING!** To reduce the risk of personal injury or damage to equipment:

- Extend leveling jacks to the floor.
 - Ensure that the full weight of the rack rests on the leveling jacks.
 - Install stabilizing feet on the rack.
 - In multiple-rack installations, secure racks together.
 - Extend only one rack component at a time. Racks may become unstable if more than one component is extended.
-

HP technical support

Telephone numbers for worldwide technical support are listed on the HP support web site:
<http://www.hp.com/support/>.

Collect the following information before calling:

- Technical support registration number (if applicable)
- Product serial numbers
- Product model names and numbers
- Applicable error messages
- Operating system type and revision level
- Detailed, specific questions

For continuous quality improvement, calls may be recorded or monitored.

HP strongly recommends that customers sign up online using the Subscriber's choice web site:
<http://www.hp.com/go/e-updates>.

- Subscribing to this service provides you with e-mail updates on the latest product enhancements, newest versions of drivers, and firmware documentation updates as well as instant access to numerous other product resources.
- After signing up, you can quickly locate your products by selecting **Business support** and then **Storage** under Product Category.

HP-authorized reseller

For the name of your nearest HP-authorized reseller:

- In the United States, call 1-800-282-6672.
- Elsewhere, visit the HP web site: <http://www.hp.com>. Then click **Contact HP** to find locations and telephone numbers.

Helpful web sites

For other product information, see the following HP web sites:

- <http://www.hp.com>
- <http://www.hp.com/go/storage>
- <http://www.hp.com/support/>
- <http://www.docs.hp.com>

1 Introducing XPath OS features

XPath OS is the operating system for the HP StorageWorks MP Router. The MP Router is an open, intelligent switching platform for hosting multi-protocol routing services and other applications within a SAN fabric.

In addition to the full set of standard Fibre Channel switch features, XPath OS supports optional HP multi-protocol routing services and interoperates with the HP Fabric OS. The MP Router can be used to connect HP fabrics to McDATA fabrics.

This chapter contains the following sections:

- [MP Routing Services](#), next
- [Fibre Channel features](#), page 13
- [HP Fabric OS support](#), page 14
- [Daemon overseer service](#), page 15
- [High availability](#), page 17

MP Routing Services

HP multi-protocol routing services provide options for extending SAN benefits over multiple networks to larger SAN sizes, and across longer distances. Using these services, you can interconnect devices between SAN fabrics without merging those fabrics, thereby providing a more secure and flexible storage networking foundation.

These services include:

- HP FC-FC Routing Service for SAN connectivity
- HP Fibre Channel over IP (FCIP) Tunneling Service for SAN extension over distance
- HP iSCSI Gateway Service for sharing Fibre Channel resources with Internet Small Computer Systems Interface (iSCSI) devices

The services are summarized in the following sections.

FC-FC routing service

The FC-FC Routing Service enables devices located in separate SAN fabrics to establish communication without requiring the fabrics to merge into a single large SAN. By using this service, you can interconnect devices without having to redesign or reconfigure the entire environment. You can connect physically separate SANs into logical SANs (LSANs) in a seamless manner. LSANs are ideal for:

- Simplifying SAN design, implementation, and management
- Providing a seamless way to share resources across multiple SANs without the complexity of physically merging those SANs
- Creating a more unified SAN environment with easier interconnection and support for SANs and SAN resources deployed for different purposes
- Allowing routing between HP and McDATA fabrics

For details on the capabilities, terminology, and use of FC-FC routing, see Chapter 4, "[Using the FC-FC Routing Service](#)."

FCIP Tunneling Service

The FCIP Tunneling Service enables you to extend the Fibre Channel SAN over distances that would be impractical or too expensive with native Fibre Channel links. The service employs a proprietary transport protocol that allows the transparent interconnection of geographically distributed SANs through an IP-based network.

XPath OS supports FCIP between two MP Routers only, not between an MP Router and another switch model.

FCIP enables Fibre Channel frames to tunnel through IP networks by dividing frames, encapsulating the result in IP packets when they enter the tunnel, and then reconstructing them as they leave the tunnel.

Figure 1 illustrates a simple FCIP tunneling configuration.

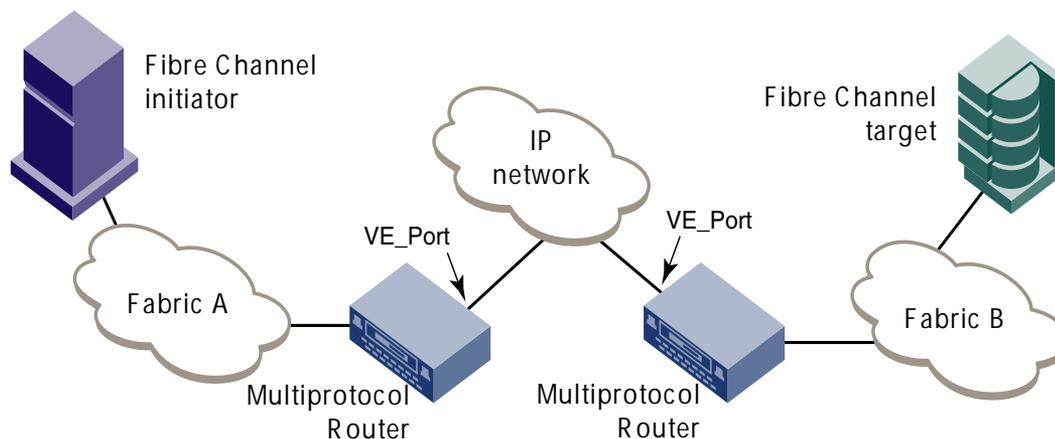


Figure 1 Simple FCIP tunneling configuration

Chapter 5, “Using the FCIP Tunneling Service,” further describes these concepts and provides procedures for configuring FCIP.

iSCSI Gateway Service

The iSCSI is a protocol that defines the processes for transferring block storage over TCP/IP networks by encapsulating SCSI commands into TCP and transporting them over the network via IP. The iSCSI Gateway Service enables organizations to integrate low-cost Ethernet-connected servers into HP StorageWorks Fibre Channel SANs by bridging the iSCSI protocol to the Fibre Channel protocol. This capability allows iSCSI servers to leverage shared SAN resources, improving asset utilization and enabling new applications. This integration greatly reduces the cost of connecting servers to centrally managed storage and helps provide a cost-effective solution to introduce utility computing into the enterprise.

Chapter 6, “Using the iSCSI Gateway Service,” further describes these concepts and provides procedures for configuring an MP Router for use as an iSCSI gateway.

Combining services

Using FCIP tunneling along with FC-FC routing enables you to keep two fabrics separate instead than merging them into a single fabric, which would permit any-to-any connectivity among all devices. For example, you can build a configuration similar to that shown in Figure 2, in which FCIP tunnelling and an IP network provide transport between two MP Routers.

iSCSI cannot run concurrently with FCIP tunneling on the same port, but both can be used concurrently on the same switch without restrictions.

If you combine iSCSI and FC or iSCSI and FCR and FCIP, the iSCSI initiators can access FC storage only within the backbone fabric.

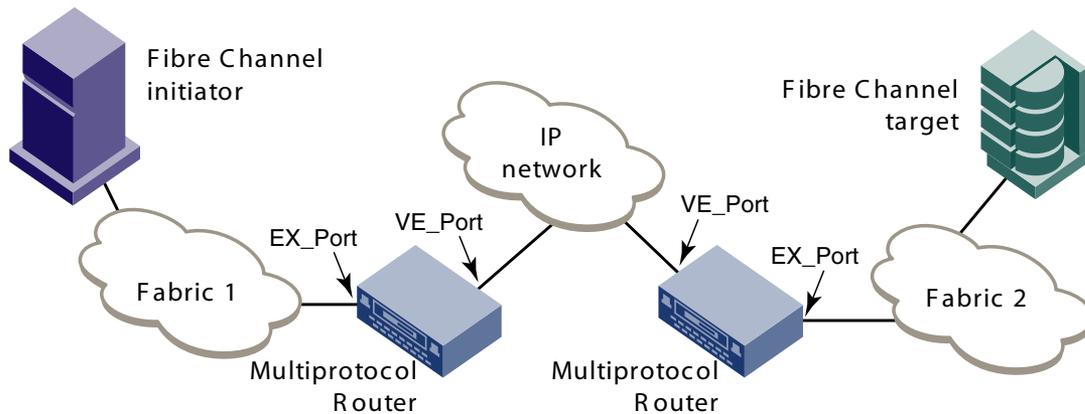


Figure 2 Combining FCIP tunneling and FC-FC routing services

Fibre Channel features

XPath OS provides the Fibre Channel features described in [Table 2](#).

Table 2 XPath OS Fibre Channel switch features

Feature	Description
Arbitrated-loop support	The MP Router supports only one arbitrated-loop physical address (AL_PA) device per port. However, you can use MP Routers in fabrics where multiple loop devices are directly attached to Fabric OS-based switches, such as the HP StorageWorks SAN Switch 2/32.
Gigabit Ethernet (GigE) support	XPath OS supports GigE port connections, which in turn support the optional multi-protocol routing services. Each Gigabit Ethernet port has its own IP address and supports address resolution protocol and ping. Ports must be specifically configured as GigE ports. When configured as GigE, ports are equivalent to 1000Base-SX or 1000Base LX.
Exchange-based trunking	The XPath OS exchange-based trunking feature optimizes load sharing across interswitch links (ISLs). It increases overall bandwidth by grouping all ISLs between the same pair of switches and load-sharing traffic across all equal path links. This functionality works between two MP Routers, as well as from an MP Router to any other Fabric OS-based switch. For more information, see Chapter 8, "Using ISL trunking."
Management access ports	The MP Router provides two 10/100 Mbit/sec Ethernet ports, which are capable of being configured as a single IP address or as two unique IP addresses, and an RS-232 port. If you configure the Ethernet ports with unique IP addresses, when the router is joined to a fabric, the router must have a virtual Ethernet IP identity for routing. You can handle this by setting a virtual management IP address for the router. See the <i>HP StorageWorks XPath OS 7.4.x command reference guide</i> for more information.
Multiple user accounts	XPath OS allows you to create multiple user accounts and assign administrative or user roles.
Name Server support	The MP Router includes an FC-GS3-compliant Name Server implementation. The Name Server supports all commands specified as "Required" in the FC-MI Fibre Channel specification.

Table 2 XPath OS Fibre Channel switch features (continued)

Feature	Description
Point-to-point and loop mode topologies	MP Router ports can be configured to support either point-to-point mode or loop mode using the <code>portCfgTopology</code> command. When in default point-to-point mode, a port automatically detects whether there is another switch on the other end of the connection (where the port automatically becomes an E_Port) or whether there is a host or target on the other end (where the port automatically becomes an F_Port).
Threshold management	XPath OS provides threshold management functions for temperature sensing and power supplies, as well as warnings for fan failures.
User interfaces	XPath OS provides a graphical user interface, Advanced Web Tools, and a command line interface (CLI). The CLI is described in the <i>HP StorageWorks XPath OS 7.4.0 Advanced Web Tools administrator guide</i> . The CLI commands are described in the <i>HP StorageWorks XPath OS 7.4.x command reference guide</i> . In normal operation, the CLI interface is used by running a telnet or SSH session from an administrative workstation networked to an MP Router Ethernet management interface.
Zone server support	The MP Router zone server implementation complies with Fabric OS 2.6.x, 3.x, and 4.x. In multiswitch configurations, the MP Router also complies with the FC-SW2 zone server specification. The maximum number of members allowed in the zone server depends on the size of the zone database; the effective zone database can have a maximum of 4096 zone members and 1024 zones (including fabric assisted zones). The maximum zoning database size is 128 KB.

HP Fabric OS support

The HP StorageWorks MP Router automatically forms a fabric when an enabled port is connected to another switch running a compatible Fabric OS version. For best performance, all fabric switches should be updated to the latest firmware. Firmware updates are available to OEMs, partners, and customers with support service contracts on the HP Fibre Channel switches web site:

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

To download:

1. In the **Storage products** section, click **SAN infrastructure**.
The SAN Infrastructure page opens.
2. In the **SAN Infrastructure products** section, click **Multi-protocol Routers and Gateways**.
The Multi-protocol routers and gateways page opens.
3. Click **B-Series Multi-Protocol Router**.
The HP StorageWorks B-Series Multi-Protocol Router page opens.
4. In the **Product Information** section, click **Software, firmware & drivers**.
The technical support -HP B-series Multi-protocol Router page opens.
5. In the **I would like to** section, click **download drivers and software**.
The download drivers & software - specify product name page opens.
6. In the **select your product** section, click either **HP StorageWorks Multiprotocol Base 8-port Router** or **HP StorageWorks Multiprotocol Full 16-port router**.
The appropriate specify operating system page opens.
7. In the **select operating system** section, click **Cross operating system (BIOS, Firmware, Diagnostics, etc.)**.
The appropriate download drivers and software page opens.
8. Scroll down to the **Firmware** section and click the appropriate **download** button.
9. Follow the prompts in the File Download dialog box.

PID mode requirements

Either core port identifier (PID) mode or extended-edge PID mode must be enabled to successfully execute the XPath OS. In its default configuration, XPath OS has core PID mode enabled.

All switches within a fabric must have the same PID mode enabled. For more information on configuring a fabric PID mode, see ["Verifying the PID mode"](#) on page 20.

For information on configuring the PID mode in Fabric OS, see the *HP StorageWorks Fabric OS 5.x administrator guide*.

 **NOTE:** The PID mode for the MP Router need not match that of the attached fabrics; however, the PID mode set on the EX_Port must match. To set the PID mode, see ["Setting the PID mode"](#) on page 20.

Feature compatibility

XPath OS does not support the following Fabric OS features:

- The `interopMode` command. The MP Router segments from a fabric that has `interopMode` enabled. See the `interopMode` command in the *HP StorageWorks XPath OS 7.4.x command reference guide* for more information.
- Fabric Watch
- Functionality equivalent to the Open E_Port feature of Fabric OS. The Open E_Port feature permits Fabric OS to support certain Fibre Channel E_Port extension devices.

In addition, the following differences exist between Fabric OS and XPath OS:

- XPath OS allows longer connections than the native 10 km distance specified by the Fibre Channel standard. Long-distance connections are configured using the `portCfgLongDistance` command. For details, see ["Configuring a long-distance connection"](#) on page 22.
- XPath OS does not maintain an alias server, but the FC-FC routing service makes use of aliases managed by attached Fabric OS fabrics.
- XPath OS requires that ports be disabled before changing their properties. Ports are disabled by default. If ports have been enabled, actions that return a port to default settings—such as the `configDefault` and `portCfgDefault` commands—disable the port. Before configuring an enabled port for multi-protocol routing services, the port must be stopped. See the `portDisable` and `portStop` commands in the *HP StorageWorks XPath OS 7.4.x command reference guide* for more information.

Daemon overseer service

Many of the applications running on the MP Router are controlled by daemons. An overseer daemon monitors the status of the system daemons.

 **NOTE:** The Full Fabric License, which is shipped as standard with the MP Router, is required to bring up the daemons.

This overseer daemon checks the heartbeat of all the daemons running in the MP Router. If the overseer fails to receive three consecutive heartbeats, it then informs all the daemons who are interested in the failure notification.

Once it receives a response from the appropriate daemon, it takes the action specified: either rebooting the switch or restarting the daemon.

If the notified daemon does not respond within 30 seconds, the overseer then initiates the default action shown.

The overseer daemon monitors the daemons listed in [Table 3](#).

Table 3 Monitored XPath OS daemons

Daemon	Description	Failure	Default
Chassismgr	Chassis manager: Captures all the management functionalities that apply only at the switch level, for example psShow, fanShow, ipAddrSet, and chassisShow.	Base	Reboot
Chassispd	Chassis performance daemon: Periodically monitors certain chassis elements and stores the information in shared memory.	Base	Reboot
Cpmgr	Manages and controls card and port attributes, and validates those attributes coming from the CLI and the Objmgr. Sends and receives messages from lomctlr.	Base	Reboot
Diagd	Diagnostic daemon: Typically indicates hardware failure.	Base	Reboot
Evtmgr	Event manager: Collects all events generated by other management tasks.	Base	Reboot
Fabctl	Fabric controller: Determines E_Ports, assigns unique domain IDs, creates a spanning tree, throttles the trunking process, and distributes domain and aliases lists to all switches in the fabric.	Port	Reboot
Fcipd	FCIP daemon: Tunnels Fibre Channel frames over an IP network.	Port	Reboot
Iomctlr	I/O management controller: Responsible for downloading specific firmware onto the port, depending on the porttype. Manages and controls the card and port attributes, and sends and receives messages to cmpgr and to the port agent.	Base	Reboot
Iswitchd	Iswitch daemon: Provides Fibre Channel routing functionality. Routing license is required.	Port	Reboot
Licensed	License daemon: Manages licensing information on the switch.	Base	Reboot
Msd	Management server daemon: Used with management of a switch using API/SMI Interface.	Base	Reboot
Nsd	Name server daemon: Provides the discovery of devices attached to the SAN based on Fibre channel protocol - FC GS-3.	Base	Reboot
Objmgr	Object manager: Manages access to the database residing in the flash memory. All commands that modify, create, or delete the database go through objmgr, which provides the persistence for the chassis management system.	Base	Reboot
Rcsd	Reliable commit service daemon: Provides the reliable commit service for replication of data for applications like zoning across fabric switches.	Base	Reboot
Rpcd	Remote procedure call daemon: A management daemon server process that facilitates API, FM and SMI-Client for execution of procedure calls on switches across network.	Base	Restart
Sb	Facilitates sending and receiving messages between the various application tasks and daemons.	Base	Reboot

Table 3 Monitored XPath OS daemons (continued)

Daemon	Description	Failure	Default
Snmpd	Simple Network Management Protocol daemon: MANAGEABILITY daemon for monitoring and managing a network device/switch.	Base	Restart
Xmlld	XML-rpc daemon: Handles communication with Web Tools GUI client using xml-rpc.	Base	Restart
Zsd	Zoning daemon: Provides the services of Fibre channel zoning for use with Fabric Zone Merge and zoning configurations.	Base	Reboot

High availability

XPath OS 7.4.x includes switch monitoring capabilities; this means that the MP Router monitors all critical services. If any of these services fail, the router might reboot in order to preserve data and fabric integrity.

If you notice repeated reboots, issue the `errShow` command to determine the reason.

If the switch is unable to move into the `service ready` state within five minutes of rebooting, the MP Router changes to the alternate bank and reboots again. The router is now set to disabled, to preserve the integrity of the fabric. If the router is still not able to activate using the alternate bank, it reboots into the recovery kernel.

To restore the MP Router to service, establish a telnet connection and issue the appropriate CLI command to restore the service.

2 Performing basic configuration

This chapter provides procedures for the basic switch configuration tasks that are frequently performed as part of routine SAN administration:

- [Viewing router information](#), next
- [Verifying the PID mode](#), page 20
- [Creating interswitch links](#), page 21
- [Configuring a long-distance connection](#), page 22
- [Verifying connectivity](#), page 23
- [Synchronizing time with an NTP server](#), page 23
- [Licensed features](#), page 24
- [Changing account passwords](#), page 24
- [Enabling and disabling switches](#), page 25
- [Enabling and disabling ports](#), page 25
- [Activating Ports on Demand: upgrading an 8-port base model to a 16-port full model](#), page 26
- [Setting the domain ID](#), page 27
- [Controlling routing within a fabric](#), page 28
- [Displaying command help](#), page 31

Viewing router information

Use the `switchShow` command to display detailed information about the MP Router.

For example:

```
router:admin> switchshow
Switch Name   : FabricAP
Switch State  : Online
Switch Type   : 38.0
Switch Role   : Subordinate
Switch Domain : 100
Switch ID     : FFFC64
Switch WWN    : 10:00:00:05:1e:13:84:00
Beacon status: OFF
Zoning       : ON (lsan_cfg1)

FC router BB Fabric ID: 1

Port Media Speed State      Info
=====
 0   id   AN   No_Light
 1   id   N2   Online    EX_PORT 10:00:00:60:69:c0:6d:57 "brcd72" (fabric id = 3)
 2   id   N2   Online    EX_PORT 10:00:00:60:69:90:10:98 "brcd73" (fabric id = 4)
 3   id   N2   Online    E_PORT  10:00:00:05:1e:12:f8:00 "mars55" (downstream)
 4   id   AN   No_Light
 5   id   AN   No_Light
 6   id   AN   No_Light
 7   id   N2   Online    F_PORT  21:00:00:e0:8b:08:64:48
 8   id   N2   Online    F_PORT  22:00:00:80:e5:12:66:43
 9   id   AN   No_Light
10   id   AN   No_Light
11   id   AN   No_Light
12   id   AN   No_Light
13   id   AN   No_Light
14   id   AN   No_Light
15   id   AN   No_Light
```

See also “[Viewing hardware status](#)” on page 41.

Verifying the PID mode

XPath OS uses either core PID mode (the default) or extended-edge PID mode. All switches in a fabric must be configured to use the same PID mode.

 **NOTE:** The PID mode for the MP Router need not match that of the attached fabrics; however, the PID mode set on the EX_Port must match.

Viewing the PID mode

1. Log in as admin.
2. Issue the `configShow` command.
3. Verify that the value of following line is either 1 (core PID mode) or 2 (extended-edge PID mode):

```
fabric.ops.mode.pidFormat: 1
```

Setting the PID mode

 **NOTE:** To set the PID mode on an EX_Port, see “[Configuring an interfabric link](#)” on page 52. To set the PID mode of a Fabric OS switch, see the *HP StorageWorks Fabric OS 5.x administrator guide*.

1. Issue the `switchDisable` command to disable the switch.
2. Issue the `configure` command.
3. Press **Enter** after the prompts until the `Switch PID Address Mode` prompt appears.
4. At the `Switch PID Address Mode` prompt, enter 1 to enable core PID mode, or enter 2 to enable extended-edge PID mode.
5. Issue the `switchEnable` command to enable the switch.

Creating interswitch links

When you physically connect cables from an MP Router to other MP Routers or to Fabric OS-based switches, a fabric forms automatically. See the installation guide for your switch for specific ISL connection and cable management information.

The following conditions must be met on Fabric OS switches:

- The Platform Services feature must be disabled (this is the default).
- The switch and ports must be enabled.

 **NOTE:** To avoid inadvertent fabric merging, MP Router ports should either be disabled or configured as `EX_Ports` before physically connecting ISLs.

Use the following procedure to disable the Platform Services feature on a Fabric OS switch. For more information on the commands, see the *HP StorageWorks XPath OS 7.4.x command reference guide*. For more information on Fabric OS procedures, see the *HP StorageWorks Fabric OS 5.x administrator guide*.

Verifying that Platform Services is enabled (Fabric OS or Secure Fabric OS switch only)

1. On a Fabric OS or Secure Fabric OS switch only, log in to the switch as admin.
2. Issue the `msPlatShow` command.

If the platform is disabled, you need not proceed further.

If the Platform Services database is populated, the output looks similar to the following:

```
switch:admin> msplatshow
-----
Platform Name: [9] "first obj"
Platform Type: 5 : GATEWAY
Number of Associated M.A.: 1
[35] "http://java.sun.com/products/plugin"
Number of Associated Node Names: 1
Associated Node Names:
10:00:00:60:69:20:15:71
-----

Platform Name: [10] "second obj"
Platform Type: 7 : HOST_BUS_ADAPTER
Number of Associated M.A.: 1
Associated Management Addresses:
[30] "http://java.sun.com/products/1"
Number of Associated Node Names: 1
Associated Node Names:
10:00:00:60:69:20:15:75
```

Disabling the Management Server feature (Fabric OS or Secure Fabric OS switch only)

1. Issue the `secModeShow` command to see whether the Secure Fabric OS feature is enabled:
 - If it is disabled, go directly to [step 3](#).
 - If it is enabled, log in to the primary FCS switch with secure telnet (`sectelnet`) or secure shell (SSH), and continue with [step 2](#).
For details on `sectelnet` and SSH, see the *HP StorageWorks Secure Fabric OS user guide*.
2. Issue the `secModeDisable` command to disable the Secure Fabric OS feature on all switches in the fabric; enter `y` at the prompt to confirm.

 **NOTE:** The `secModeDisable` command disables security on all switches in the fabric, deleting both the defined and active security databases.

3. Issue the `msplMgmtDeactivate` command and enter `y` when prompted to confirm.
This command can be used from any switch in the fabric.

You can confirm fabric connectivity by logging in to the MP Router as `admin` and issuing the `fabricShow` command to display a summary of all the switches in the fabric:

```
router:admin> fabricshow
Switch ID   Worldwide Name           Enet IP Addr   Name
-----
  1: fffc01 10:00:00:05:1e:12:f8:00 10.33.52.55    "mars55"
  2: fffc02 10:00:00:05:1e:16:1f:12 10.33.52.70    "mars46"
  3: fffc03 10:00:00:60:69:c0:6d:57 10.33.52.72    "brcd72"
100: fffc64 10:00:00:05:1e:13:84:00 10.33.52.66    "FabricAP"
200: fffc08 10:00:00:60:69:90:10:98 10.33.52.73    "brcd73"
The Fabric has 5 switches
```

Enabling the switch and ports

Use the `switchEnable` and `portEnable` commands.

Configuring a long-distance connection

Use the `portCfgLongDistance` command to allocate enough full-size frame buffers on a particular port to support a long-distance link.

[Table 4](#) lists ISL modes. Although the ports configure automatically, for best performance you should configure each port connected by the ISL to the same distance level. When the connection is initiated, the fabric reconfigures.

 **NOTE:** If you are creating a long-distance ISL from the MP Router running XPath OS to a switch running Fabric OS 4.4.x or 3.2.x, you must set VC translation link initialization to 0 on the Fabric OS switch. VC translation link initialization is set to 1 by default in these Fabric OS versions, but it is not supported in XPath OS. For details on setting this option to 0, see the *HP StorageWorks XPath OS 7.4.x command reference guide* for the `portCfgLongDistance` command.

Table 4 ISL modes

Mode	Description	Maximum ISL distance
L0	Level 0 (the default value for switch ports)	10 km at 1 Gbit/sec or 2 Gbit/sec
L0.5	Level 0.5	25 km at 1 Gbit/sec or 2 Gbit/sec, depending on the number of buffers reserved

Table 4 ISL modes (continued)

Mode	Description	Maximum ISL distance
LE	Level E	10 km at 1 Gbit/sec or 2 Gbit/sec
LS	Level S	300 km at 1 Gbit/sec or 2 Gbit/sec

Configuring port 15 for the LS distance level

```

router:admin> portdisable 15
    Port 15 disabled.
router:admin> portcfglongdistance 15 LS
    Distance level is set to LS on port 15.
router:admin> portenable 15
    Port 15 enabled.

```

Verifying connectivity

Verify fabric-wide device connectivity by displaying the fabric-wide device count. The number of devices listed in the Name Server reflects the number of devices that are connected.

1. Log in as admin.
2. Issue the `switchShow` command to verify that local devices are achieving link connectivity.
3. Issue the `nsShow` command to display Name Server information and to verify that local devices have successfully registered with it, as shown in the following example:

```

router:admin> nsshow
{
  Type Pid      COS      PortName                               NodeName
  N   370000;    2,3;10:00:00:00:c9:27:2b:69;20:00:00:00:c9:27:2b:69
      FC4s: FCP [Initiator]
      NodeSymb: Emulex LP952 FV3.91A1 DV5-4.82A16
      Fabric Port Name: 20:00:00:05:1e:16:1f:04
  NL  370101;    2,3;10:00:00:00:c9:34:06:ac;20:00:00:00:c9:34:06:ac
      FC4s: FCP [Initiator]
      Fabric Port Name: 20:01:00:05:1e:16:1f:04
  N   370200;    2,3;10:00:00:00:c9:34:06:aa;20:00:00:00:c9:34:06:aa
      FC4s: FCP [Initiator]
      Fabric Port Name: 20:02:00:05:1e:16:1f:04
  The Local Name Server has 3 entries }

```

4. Issue the `nsAllShow` command to display the 24-bit Fibre Channel addresses of all devices in the fabric:

```

router:admin> nsallshow
0d0ce1 0df001 0df002 0df003 0df004 0df005 0df006 0df007
0df008 0df009 0df00a 0df00b 0df00c 0df00d 0df00e 0df00f
0df010
17 Nx_Port devices present in the fabric

```

Synchronizing time with an NTP server

All MP Routers maintain the current date and time in nonvolatile memory for event logging purposes; proper switch operation does not depend on the correct date and time.

Use the following procedure to synchronize the local time of the MP Router to that of an external NTP server. To synchronize all MP Routers in a fabric, run this procedure on each.

Using an NTP server for time service

1. Log in as admin.
2. Issue the following command:

```
tsclockserver "[ipaddr]"
```

The optional *ipaddr* parameter is the IP address of the NTP server, which must be accessible from the MP Router. The default value is `LOCL` (local).

For example:

```
router:admin> tsclockserver
LOCL
router:admin> tsclockserver "132.163.135.131"
router:admin> tsclockserver
132.163.135.131
```

 **NOTE:** When an NTP server is configured, the `date` command cannot be used to set the date and time; it can only report the time.

Licensed features

The following licenses are included with XPath OS 7.4.x:

- XPath Advanced Web Tools
- XPath Advanced Zoning
- XPath Exchange-Based Trunking
- XPath Fibre Channel Routing Service
- XPath Base Switch or Full Fabric
- XPath FCIP
- XPath iSCSI

Removing a license:

1. Log in as admin.
2. Issue the `licenseRemove` command followed by the license key:

```
router:admin> licenseremove bQebzbRdScRfc0iK
License bQebzbRdScRfc0iK is removed.
```

3. When removing some licenses, you must reboot the system to complete the removal.

Changing account passwords

Accounts can be assigned either the administrative (admin) role or the user role. You can create a very large number of accounts, but because each account consumes system resources, such as file system space for home directories, a practical limit is 10 accounts with the admin role and 10 accounts with the user role.

Accounts with the admin role can reset the passwords to default for all accounts with the admin role or user role. Accounts with the user role can change the password only for themselves; they cannot change the password for other accounts.

Use the `passwd` command to change account passwords or the `userAdd` command to create new accounts. The syntax for the `passwd` command is:

```
passwd [-u name]
```

The *name* parameter is the name of the account for which you want to change the password.

Without *name* specified, you are prompted to change the password for the current account. You must enter the old password and the new password and then re-enter the new password. The maximum length of a password is eight characters.

If *name* is specified and the current account has the admin role, the specified account password is reset to the default value.

Changing the admin password (when logged in as admin)

```
router:admin> passwd
Old password:
New password:
Retype new password:
The password is changed
```

 **NOTE:** As a security measure, passwords are not displayed.

For example, to reset the account user123 password to the default value (when logged in as admin):

```
router:admin> passwd -u user123
Your password:
The password of user123 is set to default
```

Enabling and disabling switches

When a switch boots, all Fibre Channel ports that pass power-on self test (POST) are enabled. If the switch was part of a fabric, it rejoins the fabric.

Enabling a switch

1. Log in as admin.
2. Issue the `switchEnable` command:

```
router:admin> switchenable
Switch enabled
```

Disabling a switch

1. Log in as admin.
2. Issue the `switchDisable` command:

```
router:admin> switchdisable
Switch is being disabled.....
Switch disabled
```

Enabling and disabling ports

Ports are disabled by default. When a switch is disabled, all Fibre Channel ports on the switch are taken offline. If the switch is part of a fabric, the fabric reconfigures. If the port is connected to another switch, the fabric might reconfigure.

When a port is enabled and that port is connected to one or more devices, the devices become available to the fabric.

Enabling a port

1. Log in as admin.
2. Issue the `portEnable` command:

```
router:admin> portenable 1
port 1 enabled
```

Disabling a port

1. Log in as admin.
2. Issue the `portDisable` command:

```
router:admin> portdisable 1
port 1 disabled
```

If the port is connected to another switch, the fabric might reconfigure. If the port is connected to one or more devices, the devices are no longer available to the fabric.

Activating Ports on Demand: upgrading an 8-port base model to a 16-port full model

The MP Router is available with either eight ports, 0 through 7 (base model), or sixteen ports, 0 through 15 (full model) activated. If your MP Router shipped with eight active ports, you can activate the remaining eight ports by purchasing and installing the HP StorageWorks MP Router Upgrade License.

 **NOTE:** Check port status to verify whether the license is pre-installed. For example, use the `portshow` command for ports 8 through 15. If the port status output indicates `Started` and `Licensed`, all 16 ports are activated. See the *HP StorageWorks XPath OS 7.4.x command reference guide* for information on the `portshow` command.

Activating Ports on Demand:

1. If ports 8 through 15 show no license, purchase the HP StorageWorks HP Router Upgrade License from an authorized HP representative.
Your HP representative requires the MP Router's World Wide Name (WWN) to assign a license key.
2. Issue the `switchshow` command to obtain the WWN of your MP Router.
3. Install the HP StorageWorks MP Router Upgrade License (the license key is a string of approximately 16 uppercase and lowercase letters and numerals):
 - a. Log in to the MP Router as admin.
 - b. Issue the `licenseadd` command with the license key enclosed in quotation marks.

 **NOTE:** Enter the license key exactly as issued. If you enter it incorrectly, the license does not function properly.

- c. Issue the `licenseshow` command to determine whether the license is valid.
If a licensed product is not displayed, the license is invalid.
-

 **NOTE:** After entering a license, the licensed product is available immediately; the system does not require a reboot.

4. Configure the inactive ports.
For example, if you are using the ports for routing and connecting to an edge fabric, use the `portcfgexport` command to configure the port as an `EX_Port`.

5. Issue the `portstart` command to start the ports. For example:
`portstart 8-15`
6. Issue the `portenable` command to enable the ports. For example:
`portenable 8-15`
7. Option: Issue the `portshow` command to verify that the newly active ports are Started.

NOTE: If you remove the HP StorageWorks MP Router Upgrade License, ports 8 through 15 no longer work after the next reboot.

Setting the domain ID

A domain ID is assigned dynamically when a switch is enabled. However, the domain ID can be set manually, for example, if you need to control the number to resolve a domain ID conflict when merging fabrics.

Displaying a list of current domain IDs in the fabric

1. Log in as admin.
2. Issue the `fabricshow` command:

```
router:admin> fabricshow
Switch ID   Worldwide Name           Enet IP Addr   Name
-----
74: fffc4a  10:00:00:60:69:50:09:1a  192.168.172.74  "mw74"
75: fffc4b  10:00:00:60:69:50:08:35  192.168.172.75  "mw75"
100: fffc64  10:00:00:05:1e:13:8b:00  192.168.163.103 "apswitch"
103: fffc67  10:00:00:05:1e:13:83:00  192.168.163.102 "ap_102"

The Fabric has 4 switches
```

Fabric information is displayed, including the domain ID (listed in the Switch ID column).

To prevent a domain ID conflict, make sure all domain IDs are unique before connecting an MP Router to the fabric.

Assigning a domain ID manually

1. Log in as admin.
2. Issue the `switchDisable` command to disable the switch.
3. Issue the `configure` command.
4. Enter a unique domain ID at the domain ID prompt.
5. Press **Enter** in response to the remaining prompts.
6. Issue the `switchEnable` command to re-enable the switch.

For example:

```
router:admin> switchdisable
Switch is being disabled..
.....
router:admin> configure
Fabric parameter set. <cr> to skip a parameter
Domain: (1..239 or f(fabric_assign)) [100 unconfigured] 5
BB Credit: [1..32] [16]
R_A_TOV: (4000..120000) [10000]
E_D_TOV: (1000..5000) [2000]
Data field size: (256..2112 multiple of 4) [2112]
Switch PID Address Mode (1..2) [1]
WAN_TOV (0..R_A_TOV/4 ) [0]
MAX_HOP_COUNT (7..19 ) [7]
End-device RSCN Transmission Mode (0..2) [1]
Fabric configuration set
You must run 'switchenable' to put the switch back to running state
router:admin> switchenable
```

Controlling routing within a fabric

There are two major aspects of Fibre Channel routing that you can control:

- Frame delivery order
- Use of dynamic load-sharing

Specifying frame delivery order

In a stable fabric, frames are always delivered in order, even when the traffic between switches is shared among multiple paths. However, when topology changes occur in the fabric (for example, if a link goes down), traffic is rerouted around the failure, and some frames can be delivered out of order. Most destination devices tolerate out-of-order delivery, but some do not.

By default, out-of-order frame-based delivery is allowed to improve speed. You should force in-order frame delivery across topology changes only if the fabric contains destination devices that cannot tolerate occasional out-of-order frame delivery.

Forcing in-order frame delivery across topology changes

1. Connect to the switch and log in as admin.
2. Issue the `iodSet` command.

 **NOTE:** This command can cause a delay in the establishment of a new path when a topology change occurs; it should be used with care.

Restoring out-of-order frame delivery across topology changes

1. Connect to the switch and log in as admin.
2. Issue the `iodReset` command.

Using dynamic load sharing

Routing is generally based on the incoming port and the destination domain. This means that all the traffic coming in from a port (either `E_Port` or `Fx_Port`) directed to the same remote domain is routed through the same output `E_Port`.

To optimize fabric routing, when there are multiple equivalent paths to a remote switch, traffic is shared among all the paths in a process called *dynamic load sharing* (DLS). DLS is recomputed when a switch is booted up or every time a change in the fabric occurs. A change in the fabric is defined as an `E_Port` going up or down or an `Fx_Port` going up or down.

DLS is enabled by default. When DLS is enabled, routing changes can affect working ports. For example, if an Fx_Port goes down, another Fx_Port might be rerouted from one E_Port to a different E_Port. The switch minimizes the number of routing changes, but some are necessary to achieve optimal load sharing.

If DLS is disabled (using the `dlsReset` command), load-sharing route determination is performed only at boot time or when an Fx_Port comes up. Optimal load sharing is rarely achieved with DLS disabled.

Viewing the current DLS setting

1. Log in as admin.
2. Issue the `dlsShow` command to view the current DLS setting.

One of the following messages appears:

- `DLS is set.` The DLS option is turned on. Load sharing is reconfigured with every change in the fabric.
- `DLS is not set.` The DLS option is turned off. Load sharing is reconfigured only when the switch is rebooted or an Fx_Port comes up.

Enabling DLS

1. Log in as admin.
2. Issue the `dlsSet` command to enable dynamic load sharing when a fabric change occurs:

```
router:admin> dlsset  
DLS feature enabled  
router:admin>
```

Disabling DLS

1. Log in as admin.
2. Issue the `dlsReset` command:

```
router:admin> dlsreset  
DLS feature disabled
```

Load-sharing path determination is performed only at boot time or when an FX_Port comes up.

Viewing routing information

You can view routing path information for the fabric, and you can view routing information through a port to a destination domain (called *unicast routing*).

Viewing routing path information

1. Log in as admin.

2. Issue the topologyShow command:

```
router:admin> topologyshow
8 domains in the fabric; Local Domain ID: 100

Domain  Metric  Hops   Out Port  Name
-----
  1      500     1      9         "switch154"
  2     1250    2      5         "switch101"
          8
  3      500     1      5         "switch252"
          8
  4     1250    2      5         "switch165"
          8
 25     1250    2      5         "msit_ref_term"
          8
 210    2750    4      5         "switch210"
          8
 211    2250    3      5         "switch211"
          8

router:admin>
```

The following information is displayed:

- Domain: The destination domain of incoming frame.
- Metric: The cost of reaching destination domain.
- Hops: The maximum number of hops to reach destination domain.
- Out Port: The available ports to which the incoming frame may be forwarded to reach the destination domain.
- Name: The switch name associated with the destination domain ID.

Viewing unicast routing information

1. Log in as admin.
2. Issue the urouteShow command:

```
router:admin> urouteshow
Local Domain ID: 100
In-Port  Domain  Out-Port  Metric  Hops  Flags  Next (Dom,Port)
  0       74      10        1500    2     D      103,10
          75      0         500     1     D      75,0
          103     10        1000    1     D      103,10

 10       74      10        1500    2     D      103,10
          75      0         500     1     D      75,0
          103     10        1000    1     D      103,10
```

The unicast routing information describes how a frame received from a port on the local switch is routed to reach a destination switch. The following information is displayed:

- Local Domain ID: The domain number of the local switch.
- In-Port: The port from which a frame comes in.
- Domain: The destination domain of the incoming frame.
- Out-Port: The port to which an incoming frame is forwarded to reach the destination domain.
- Metric: The cost of reaching the destination domain.
- Hops: The maximum number of hops required to reach the destination domain.
- Flags: Whether this route is dynamic (D) or static (S). A dynamic route is discovered by the FSPF path-selection routing task. A static route is assigned using the urouteConfig command.

- Next (Dom, Port): Domain and port number of the next hop. These are the domain number and the port number of the switch to which Out-Port is connected.

Displaying command help

XPath OS provides a help page for each command, explaining what the command does, its syntax, any operands, and the account role required to run the command.

Displaying help information about a command

1. Log in as admin.
2. Issue the `help` command:

```
help [commandname]
```

where *commandname* is the name of the command for which you would like help.

If you enter the `help` command with no argument, command summaries are displayed in alphabetical order.

3 Performing basic maintenance

This chapter provides procedures for installing software, maintaining configuration files, and checking hardware status. This information is provided in the following sections:

- [Maintaining the router configuration](#), next
- [Maintaining firmware](#), page 35
- [Performing hardware checks](#), page 41

Maintaining the router configuration

It is important that fabric configuration settings be consistent across the fabric, because inconsistent parameters can cause fabric segmentation. You should keep a backup file of the router configuration so that if settings are lost or unintentional changes are made, you can restore the configuration.

You can display configuration settings, back up configurations to a host computer, and restore configurations using the XPath OS commands shown in the following procedures, or you can use Advanced Web Tools, as described in the *HP StorageWorks XPath OS 7.4.x Advanced Web Tools administrator guide*.

Print a copy of important configuration data, including password information, and store it in a secure location. In the event of loss of both configuration information backup file and backup copies, you can use this data to start rebuilding the configuration.

Displaying configuration settings

1. Log in as admin.
2. Issue the `configShow` command:

```
router:admin> configshow
fabric.ops.domain:          100 (unconfigured default)
fabric.ops.BBCredit:        16
fabric.ops.R_A_TOV:         10000
fabric.ops.E_D_TOV:         2000
fabric.ops.dataFieldSize:   2112
fabric.ops.mode.pidFormat:  1
fabric.ops.WAN_TOV:         0
fabric.ops.MAX_HOP_COUNT:   7
switch.rscn_mode:          1

Static route:   In-Port   Domain   Out-Port

route.uicastRoute.Count:0
```

The configuration parameters for your router might vary from those shown in the example.

Backing up a configuration

1. Verify that the FTP service is running on the host computer and that you have an account on that host.
2. Log in as admin.

3. Issue the `configUpload` command:

```
configupload -h hostName -f destFileName -u userName -p password
[-t fileTransferProtocol] [-l configuration]
```

where:

<code>-h <i>hostName</i></code>	Specifies the IP address of the FTP server.
<code>-f <i>destFileName</i></code>	Specifies the destination file name.
<code>-u <i>userName</i></code>	Specifies the account name on the FTP server.
<code>-p <i>password</i></code>	Specifies the password for the account.
<code>-t <i>fileTransferProtocol</i></code>	Specifies the file transfer protocol.
<code>-l <i>configuration</i></code>	Displays the current upload configuration.

The following example uploads a configuration to file `misc/config2.txt` located on the host at IP address 10.7.32.168, using account `JohnDoe` and password `guest`:

```
router:admin> configupload -h 10.7.32.168 -f misc/config2.txt -u JohnDoe -p guest
```

Restoring a configuration

1. Verify that the FTP service is running on the host workstation and that you have an account on that host.
2. Log in as `admin`.
3. Issue the `switchDisable` command to disable the switch.
4. Issue the `configDownload` command:

```
configdownload -h hostName -f sourceFileName -u userName -p password
[-t fileTransferProtocol] [-l configuration]
```

where:

<code>-h <i>hostName</i></code>	Specifies the IP address of the FTP server.
<code>-f <i>sourceFileName</i></code>	Specifies the full path name of a previously stored configuration file.
<code>-u <i>userName</i></code>	Specifies the account name on the FTP server.
<code>-p <i>password</i></code>	Specifies the password for the account.
<code>-t <i>fileTransferProtocol</i></code>	Specifies the file transfer protocol.
<code>-l <i>configuration</i></code>	Displays the current download configuration.

5. Issue the `reboot` command to reboot the MP Router.

The following example downloads a configuration file stored at `misc/config2.txt` on the host at IP address 10.7.32.168, using account `JohnDoe` and password `guest`:

```
router:admin> configdownload -h 10.7.32.168 -f misc/config2.txt -u JohnDoe -p
guest
```

Printing configuration information

1. Log in as `admin`.
2. Enable your telnet logging functionality to create a log file and execute the following commands, or execute the commands and paste their output to a text file:
 - `licenseShow` displays the license keys that are installed.
 - `configShow` displays configuration parameter and setup information, including the port settings.
 - `ipAddrShow` displays the configured IP addresses.
3. Print a copy of the output and store it in a safe location.

 **NOTE:** Depending on your site security procedures, you might want to keep a record of the accounts and passwords for all switches in the fabric. Because this is sensitive information, you should limit access to it.

Maintaining firmware

For the latest Multi-protocol Router firmware updates, visit the MP Router web site:
<http://h18006.www1.hp.com/products/storageworks/mprouter/index.html>.

 **NOTE:** Always back up your configuration as described in “[Backing up a configuration](#)” on page 33 before making any firmware changes.

With the release of XPath OS 7.4.x, flash memory in the switch is divided into banks. Conceptually, one bank is the active bank and the other is inactive. Software (firmware) is always installed to the inactive bank. After all packages are installed to the inactive bank, the switch may be booted to activate the newly installed software.

Several command line interface tools were modified or added to support the dual-bank feature. First, firmware download always installs software packages to the inactive bank and, as an option, boots to the previously inactive bank.

Added commands include `firmwareShow`, which shows the installed software packages in both the inactive and active banks, `altBoot`, which boots to the inactive bank regardless of the last installation, and `firmwareCommit`, which commits (copies) the active bank to the inactive bank. For details about these commands, see the *HP StorageWorks XPath OS 7.4.x command reference guide*.

You can maintain two different versions of the XPath OS firmware on the MP Router (one in each bank) by disabling autocommit (using the `-n` flag) during the firmware download.

If you keep two different versions of firmware on the MP Router, make sure that you verify compatibility when installing optional and third-party components. These must be compatible with the firmware version stored in the inactive partition.

You can upgrade the firmware without affecting licensed features. See the release notes for the latest information on software and licensed-feature compatibility.

Upgrading firmware can provide the necessary software for two distinct personalities on the MP Router. These personalities are activated according to the type of licensing you activate. See “[Licensed features](#)” on page 24 for a discussion of these personalities.

Updating firmware does not change license keys, but new firmware might have new license requirements; for example, the new firmware might include a new feature that requires a new license. Check for changed license requirements when you update firmware.

The firmware is delivered in a compressed file called a *package*.

Make note of these important guidelines:

- Back up the configuration before beginning. See “[Backing up a configuration](#)” on page 33.
- Logs and core files are deleted during the firmware download. If you want to save these files, remove them from the router before beginning.
- Perform only one installation step at a time and use only one telnet session.
- As an alternative, using a serial console to connect during upgrade avoids the disruption that occurs during a telnet session.
- Maintain a separate directory for each version of firmware. Storing multiple packages in the same directory can result in the installation of a different version of the firmware than you intended.

You can also use Advanced Web Tools to install firmware. See the *HP StorageWorks XPath OS 7.4.x Advanced Web Tools administrator guide* for the procedure.

Displaying the installed version

Issue the `version` command to display information about the firmware version that is currently installed:

```
router:admin> version
=====
Installed Packages:
=====
Package Name: xpath_os_v7.4.0_bld17
Install Date: Apr 14, 2005 18:48

router:admin>
```

Installing a package

△ **CAUTION:** Whenever you install a base version of the firmware, any previously installed add-on packages, such as third-party applications, need to be reinstalled.

In preparation for an upgrade to XPath OS 7.4.0, make sure that the current installation is version 7.3.x. The upgrade procedure for previous versions is described at the end of this section.

Upgrading from XPath OS 7.3.x to XPath OS 7.4.0

1. Verify that the FTP service is running on the host and that you have an account on that host.
2. Download the package from the HP MP Router web site <http://h18006.www1.hp.com/products/storageworks/mprouter/index.html> to your host. This is a compressed file with a name similar to `XPath_OS_v7.4.tar`. It contains all of the files required to perform the upgrade.
3. Unzip the package to your FTP directory.
4. Log in to the MP Router (using the serial console) as admin.
5. Use the `firmwareDownload` command to install the package from the host to the MP Router:

```
firmwareDownload hostIPAddr userName pfile password
```

where:

hostIPAddr Specifies the IP address of the FTP server where the package is stored.

userName Specifies your user account name on the FTP server.

pfile Specifies the package file name. Specify a fully-qualified path and file name for the firmware package list, for example, `/xpath_os_v7.4.x/xpath_rk_v1.5.x`. Absolute path names can be specified using forward slashes (`/`).

password Specifies the password for the user account.

The download process takes 10 to 12 minutes. The system does not allow you to enter any commands until the download is finished.

The following is an example of a command to upgrade from XPath OS 7.3.x to 7.4.x:

```
switch:admin> firmwareDownload 10.32.2.25 releaseuser /tmp/xpath_rk_v1.5.2 password
```

6. Reboot the MP Router.

📖 **NOTE:** By default, when the MP Router is rebooted after an upgrade (by omitting the `-n` parameter) the software is synchronized between any partitions on the MP Router. If you do not want the software synchronized, use the `-n` parameter to disable autocommit.

To upgrade from XPath 7.4.x to 7.4.y

Two kinds of software updates are supported for XPath 7.4.0. These include updating the XPath base only and updating the recovery kernel and the base at the same time. These methods both use the `firmwaredownload` command, similar to the upgrade procedure from XPath OS 7.3.x to 7.4.0. Use the following procedure:

1. Log in to the MP Router as admin.
2. Use the `firmwaredownload` command to install the package from the host to the MP Router.
 - To update the base only, specify the `xpath_os_v7.4.x` file.
 - To update the base and recovery kernel, specify the `rkb` file.

The following examples show how the command might be issued when upgrading from XPath OS 7.4.x to 7.4.y.

Upgrading the XPath base only:

```
switch:admin> firmwaredownload 10.32.2.25 releaseuser /tmp/xpath_os_v7.4.x password
```

Upgrading the XPath base and recovery kernel:

```
switch:admin> firmwaredownload 10.32.2.25 releaseuser /tmp/xpath_rkb_1.5.2 password
```

You can also use Advanced Web Tools to perform updates. See the *XPath OS 7.4.x Advanced Web Tools administrator guide* for details.

Upgrading to XPath OS 7.3.0 from a previous version

Because the XPath OS 7.3.0 package included a new recovery kernel (v1.4.0), upgrading from earlier versions requires a special procedure. After you upgrade to XPath OS 7.3.0 using the special procedure, you can use the normal procedure to upgrade to subsequent releases.

Before you begin, back up your configuration as described in “[Backing up a configuration](#)” on page 33.

If the MP Router is running XPath OS 7.1.x or 7.2.x, the following procedure should be used to properly upgrade to XPath OS 7.3.0. With this new procedure, any existing configuration on the MP Router (for example, EX_Port and other port definitions, zoning, and licensing) is retained throughout the installation. However, you must manually transfer from the MP Router before installing any core files or log files that you want to keep.

A server running an FTP service that is standards-compliant is required to perform the procedure. FTP services that come bundled with major operating systems are good candidates. Third-party or shareware programs might not function properly.

The following items are optional:

- Access to the platform through its serial port, using the supplied serial cable and a terminal emulator. This is necessary if you want to follow the progress of the installation process.
- A serial console connection, if you are using the serial console to perform this procedure.

Place the following three files, which come with the XPath OS 7.3.0 release, in the same directory on the FTP server. In general, it is advisable to keep each XPath OS release in its own directory.

- `install-os.rash`
- `xpath_rk_1.4.x`
- `xpath_os_v7.3.x`

Installing XPath OS 7.3.0

1. Using telnet, log in to the Multiprotocol Router as admin.
2. Log into the switch and determine the version of XPath that is currently installed.

At the prompt, issue the `version` command. If the currently installed version is 7.3.x, you will see something similar to the following:

```

# version
RPG file server : 192.168.194.26
Root directory : /dump/74_73_downgrade
FTP username : root
FTP password : *****
Download protocol : ftp
=====
Installed Packages:
=====
Package Name : xpath_os_v7.3.0c
Installed from : bank1
Installed date : Sep 8 15:39
Administrative status : (1)
Primary status : up
Secondary status : installed and running
Disk usage on root fs - Total: 198 Mbytes, Free: 99 Mbytes.

```

- Using the `firmwareDownload` command, point to the `xpath_rk...` file by entering the following command. Using the above file example, if it was placed in `/tmp` on the FTP server with IP address 10.20.30.40 and login name "user" and password "pass", then the command is:

```
admin> firmwaredownload 10.20.30.40 user /tmp/xpath_rk_1.4.1 pass
```

You can also use Advanced Web Tools to perform this installation by pointing to the same file.

- During the above process the telnet session is lost, because this installation process automatically reboots the MP Router several times. During this time, a new recovery kernel is installed, as well as the corresponding Base image. With the example files above, at the end of the installation process, the MP Router will have recovery kernel v1.4.1 installed along with XPath OS 7.3.0.
- The entire process takes approximately 25 minutes to complete. If you want to follow the progress of the installation, connect to the serial port if it is available.

The Multiprotocol Router will now be running XPath OS 7.3.0.

 **NOTE:** You can use a one-step upgrade from XPath OS 7.3.0 to 7.4.0 by using the `firmwareDownload` command.

Place the following files, included with the XPath release, in the same directory on the FTP server. Keep each XPath release in its own directory.

For XPath OS 7.3.0, the recovery kernel version is 1.4.1. For XPath OS 7.4.0, the recovery kernel version is 1.5.2.

- `install-os.osu`
- `xpath_rk_1.4.1` (or 1.5.2)
- `xpath_os_v7.3.0` (or v7.4.0)

In addition, for 7.4.x

- `- xpath_rkb_1.5.2`

During this process, the telnet session is lost because the installation process automatically reboots the MP Router several times. During this time, a new recovery kernel is installed, as well as the corresponding base image.

The entire process takes approximately 25 minutes to complete. To follow the progress of the installation, connect to the serial port, if it is available.

Downgrading from XPath OS 7.3.0 or later to a previous version

This section describes how to downgrade to an earlier release of XPath OS. If the MP Router is running XPath OS 7.3.0 or later and you want to downgrade to release 7.2.x or earlier, follow these steps.

Before you begin, make sure that you have the following required items | :

- Access to the platform through its serial port, using a serial cable and a terminal emulator.
- The correct recovery kernel version:
 - 1.4.1.3 for XPath OS 7.3.x
 - 1.3.0.2 for XPath OS 7.2.x
 - 1.3.0.0 for XPath OS 7.1.x.
- An FTP server that the platform can access (for loading the firmware)
- If you are required to replace the recovery kernel, a TFTP server that the platform can access (the TFTP service can be on the same server as the FTP service)

Downgrading to XPath OS 7.3.0

1. Determine the version of the recovery kernel that is currently installed, log in to the switch and enter the following command:

```
admin> showrecovery
```

2. Downgrade the recovery kernel.

This is required when installing an earlier XPath OS release.

△ **CAUTION:** Downgrading the recovery kernel erases all zoning and other fabric-configuration parameters. To preserve the configuration data, perform a `configUpload` procedure before upgrading the kernel; perform a `configDownload` procedure afterward, to restore the configuration data.

- a. Connect to the switch through the serial console.
- b. Reboot the switch.
- c. Enter the boot loader (PMON) by pressing any key except **Enter** when prompted during the countdown sequence.
- d. Using a TFTP server only (an FTP server will not work), enter the following commands:

```
Multiprotocol Router> set ipaddr ipaddr-of-this-switch  
Multiprotocol Router> set netmask netmask-of-this-switch  
Multiprotocol Router> set gateway gateway-of-this-switch  
Multiprotocol Router> reboot (so that the new IP addresses can take effect)
```

- e. Repeat [step 2c](#) after the boot and then go to [step 2f](#).
- f. Enter the commands shown in the following example. For the `set bankOver` command, use the correct recovery kernel version:
 - 1.4.1.3 for XPath OS 7.3.x
 - 1.3.0.2 for XPath OS 7.2.x
 - 1.3.0.0 for XPath OS 7.1.x.

The following example shows a recovery to XPath OS 7.1.x using the recovery kernel version 1.3.0.0.

```
Multiprotocol Router> set currentdnldproto tftp
Multiprotocol Router> set bootserver ipaddr-of-the-tftp-server
Multiprotocol Router> set basefile XPathRecoverAP7420
Multiprotocol Router> set cfgbank bank0
Multiprotocol Router> set rootdir /
Multiprotocol Router> set bank0ver 1.3.0.0
Multiprotocol Router> netload -p 0xdeeddeed; g
```

 **NOTE:** The value of / is relative to the `/tftpboot` directory on most UNIX systems running the TFTP service; therefore, a value of / implies the `/tftpboot` directory, while a value of `/rk` implies the `/tftpboot/rk` directory.

After following these steps, the switch reboots and displays the recovery kernel prompt.

3. Use the recovery kernel to download the firmware.

 **CAUTION:** Using the recovery kernel to download the firmware permanently removes all zoning and other fabric-configuration parameters. To preserve the configuration data, perform a `configUpload` procedure before updating the firmware; perform a `configDownload` procedure afterward, to restore the configuration data.

- a. Connect to the switch through the serial console.
- b. Reboot the switch.
- c. When prompted during the countdown sequence, enter the boot loader (PMON) by pressing any key except **Enter**.
- d. Enter the following commands:

```
Multiprotocol Router> set cfgbank bank0
Multiprotocol Router> reboot
```

- e. After the reboot, download the XPath OS firmware package. In the following example, `xpath_os_v7.3.0` is installed. You must use an FTP server; a TFTP server will not work due of the size of the download.

```
Recovery Kernel% format all (Answer y to proceed. This will take 1-2 min.)
Recovery Kernel% set basefile xpath_os
Recovery Kernel% set rootdir <ftp dir. path to XPathOS image>
Recovery Kernel% set currentdnldproto ftp
Recovery Kernel% set bootserver <ipaddr-of-the-ftp-server>
Recovery Kernel% set ftpusername <userid> (Your FTP server's user
account.)
Recovery Kernel% set ftppassword <password> (Password for the specified
user.)
Recovery Kernel% set bank1ver 7.3.0 (. This will take 1-2 mins.)
Recovery Kernel% set cfgbank bank1 (This will take about 8
min.) Recovery Kernel% reboot
```

After the reboot, the platform starts up and XPath OS 7.3.0 is loaded.

Performing hardware checks

You can view hardware status and modify the power supply status threshold.

Viewing hardware status

Three components determine the overall status of the hardware:

- Fan speed
- Power supply
- Temperature

If the status of any component is marginal, overall status is `Marginal`. Likewise, if the status of any of the components is down, the overall status is `Down`. In all other cases, overall status is `Healthy`.

Individual rules determine the status of each component. You cannot change these rules from the CLI, but you can edit them manually in the configuration file.

[Table 5](#) lists the default rules of component status.

Table 5 Component status rules

Component type	Number failed	Status
Fan speed	1	Marginal
	>1	Down
Power supply	1	Down
Temperature	2	Marginal
	>2	Down

Viewing hardware status

1. Log in as admin.
2. Issue the `switchStatusShow` command:

```
router:admin> switchstatusshow
Switch overall status: Healthy
Reason:
  All components are in healthy status
Power overall status: Healthy
Fan overall status: Healthy
Temp overall status: Healthy
```

Viewing port status

1. Log in as admin.
2. Issue the `portShow` command, specifying the number of the port for which you want to see status.

This command displays operational status and traffic statistics, for example:

```
router:admin> portshow 12
      port 12 info
      Configuration      Current
Name :      port 12
State:      STARTED      UP
Type :      FC            FC
Link Status:  ENABLED     UP
Topology:    P-P          P-P
Speed:      AN            N2
LinkCost:   AUTO          500
Distance:   L0.5 (LM)    L0.5 (LM)
WWN:       20:0c:00:05:1e:12:f8:00

Licensed      : YES

Diag result   : PASSED

inFrames:    188
outFrames:    188
inOctets:    11408
outOctets:    8368
discards:    0
```

Viewing fan status

1. Log in as admin.
2. Issue the `fanShow` command:

```
router:admin> fanshow
Fan 1  Status: OK   Set_Speed: NORMAL   Actual_speed: 7010
RPM
Fan 2  Status: OK   Set_Speed: NORMAL   Actual_speed: 7180
RPM
Fan 3  Status: OK   Set_Speed: NORMAL   Actual_speed: 7068
RPM
Fan 4  Status: OK   Set_Speed: NORMAL   Actual_speed: 7116
RPM
Fan 5  Status: OK   Set_Speed: NORMAL   Actual_speed: 7155
RPM
Fan 6  Status: OK   Set_Speed: NORMAL   Actual_speed: 7001
RPM
```

The possible values for fan status are:

- OK: Fan is present and functioning correctly.
- NOT PRESENT: Fan is not present.
- FAIL: Fan is present, but faulty.

 **NOTE:** The number of fans and valid range for RPMs varies, depending on the type of switch. For more information, see the installation guide for your switch.

Viewing temperature status

Use the `tempShow` command to display readings of all temperature sensors in the MP Router:

1. Log in as admin.
2. Issue the `tempShow` command:

```
router:admin> tempshow
```

Index	Status	Centigrade	Fahrenheit
1	OK	21	70
2	OK	22	72
3	OK	29	84
4	OK	24	75
5	OK	25	77

The possible values for temperature status are:

- OK: Temperature is within the acceptable range.
- MARGINAL: Temperature is outside the acceptable range. Damage might occur to the switch.

The number of sensors might vary on different MP Routers. Every sensor is indexed by a sequential number.

For more information on the location of the sensors and the range of temperatures for safe operation, see the installation guide for your switch.

Viewing power supply status

1. Log in as admin.
2. Issue the `psShow` command:

```
router:admin> psshow
POWER SUPPLY 1 Serial no:0000101 Part no:60-0000754-01 Rev: A Status: OK
POWER SUPPLY 2 Serial no:0000096 Part no:60-0000754-01 Rev: A Status: OK
```

The possible values for power supply status are:

- OK: Power supply is present and functioning correctly.
- NOT PRESENT: Power supply is not present.
- FAIL: Power supply is present, but faulty.

A power supply identification line might be shown. If present, this line contains the serial number and other identification information.

Modifying the power supply status threshold

By default, a switch running on a single power supply is in a marginal status. In some cases, this is a valid configuration and you can modify the power supply status threshold to reflect this:

1. Log in as admin.
2. Issue the `configUpload` command to back up the current configuration.

See ["Backing up a configuration"](#) on page 33.

3. Open the configuration file in a text editor and find the following line:

```
switch.status.policy.PowerSupplies.marginal:1
```

Change the 1 to 0, as follows:

```
switch.status.policy.PowerSupplies.marginal:0
```

4. Issue the `configDownload` command to download the modified configuration to the switch. See ["Restoring a configuration"](#) on page 34.
5. Reboot the MP Router.

4 Using the FC-FC Routing Service

The FC-FC Routing Service is an optional, fee-based license that provides Fibre Channel routing between two or more fabrics without merging those fabrics.

The MP Router can be used simultaneously as a Fibre Channel router and as an FCIP tunnel.

This contains the following sections:

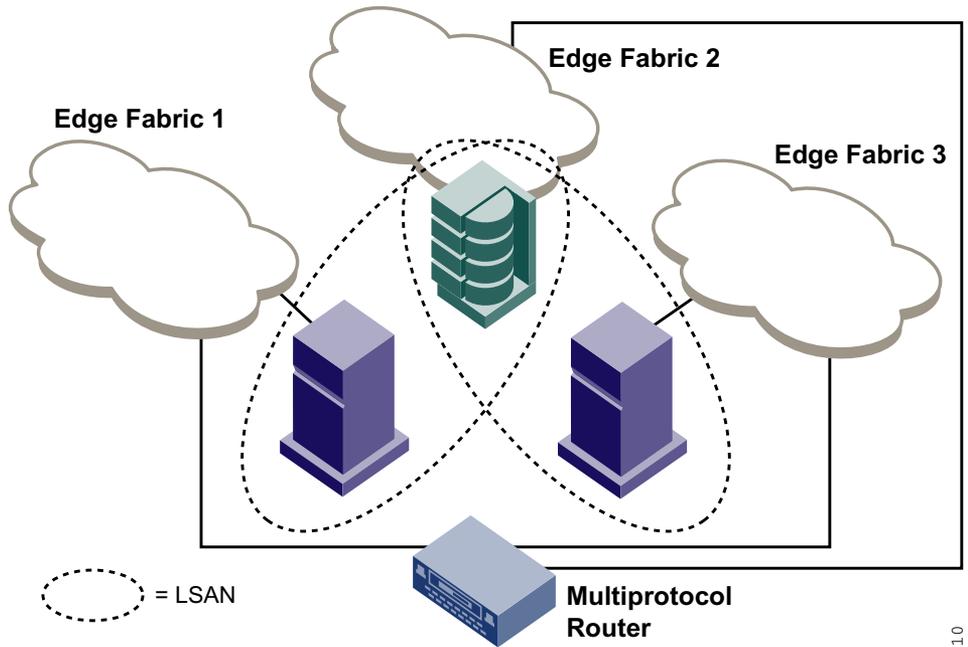
- [About Fibre Channel routing](#), next
- [Configuring an interfabric link](#), page 52
- [XPath OS and Secure Fabric OS](#), page 54
- [Setting a proxy PID](#), page 57
- [Monitoring resources](#), page 58
- [Routing ECHO](#), page 58
- [Connecting to McDATA SANs](#), page 59
- [Configuring the fabrics for interconnectivity](#), page 61
- [LSAN zoning](#), page 67

About Fibre Channel routing

FC-FC routing introduces the following concepts (see [Figure 3](#), [Figure 4](#) on page 46, and [Figure 5](#) on page 47):

- You can create logical storage area networks (LSANs) that can span fabrics. These LSANs allow Fibre Channel zones to cross physical SAN boundaries without merging the fabrics and while maintaining the access controls of zones. (See [Figure 3](#).)
- A special type of port, called an *EX_Port*, functions somewhat like an *E_Port* but terminates at the router and does not propagate fabric services or routing topology information from one edge fabric to another. The link between an *E_Port* and an *EX_Port* is called an *interfabric link* (IFL). You can configure multiple IFLs from one MP Router, from additional MP Routers, or from both. (See [Figure 4](#) on page 46.)
- One property of every *EX_Port* is its *fabric ID* (FID). Multiple *EX_Ports* attached to the same fabric have the same FID. *EX_Ports* attached to different fabrics have different FIDs.
- A standard Fibre Channel SAN with Fibre Channel targets and initiators connected through an MP Router to another Fibre Channel SAN is called an *edge SAN*.
- The edge SAN fabric is called an *edge fabric*.
- MP Routers can also connect edge fabrics using a *backbone fabric*—similar to a TCP/IP backbone network. A backbone fabric consists of at least one MP Router, and possibly a number of Fabric OS-based Fibre Channel switches. It provides transport but is otherwise transparent to the hosts, targets, and LSANs. A fabric cannot be both a backbone fabric and an edge fabric. (See [Figure 5](#) on page 47.)
- The term *meta-SAN* is used for the collection of all SANs interconnected with Fibre Channel routers. Simple meta-SANs can be constructed using a single MP Router. Additional MP Routers and router ports can be used to increase the available bandwidth between fabrics and for redundancy.

[Figure 3](#) shows a simple meta-SAN consisting of one MP Router connecting hosts in Edge Fabric 1 and Edge Fabric 3 with storage in Edge Fabric 2 through the use of LSANs.



10

Figure 3 Simple meta-SAN

Figure 4 shows a meta-SAN consisting of two edge fabrics connected through an MP Router with interfabric links.

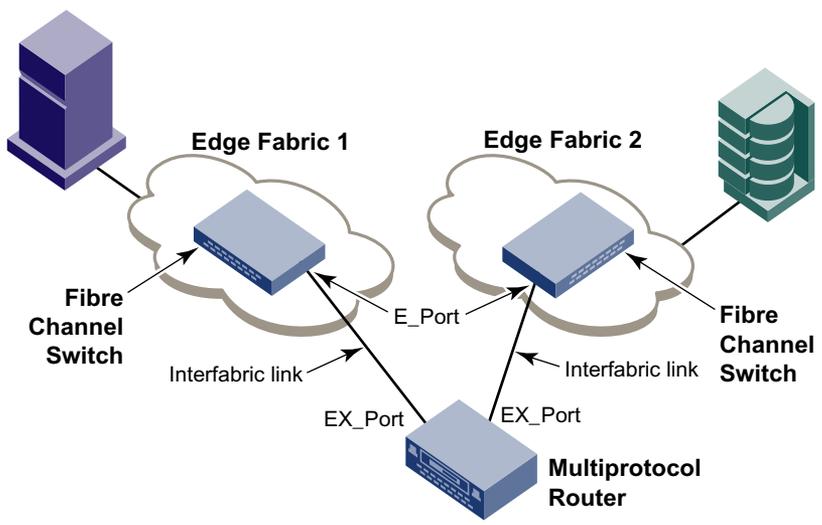


Figure 4 Meta-SAN with interfabric links

Figure 5 shows another meta-SAN consisting of a host in Edge SAN 1 connecting to storage in Edge SAN 2 through a backbone fabric connecting two MP Routers.

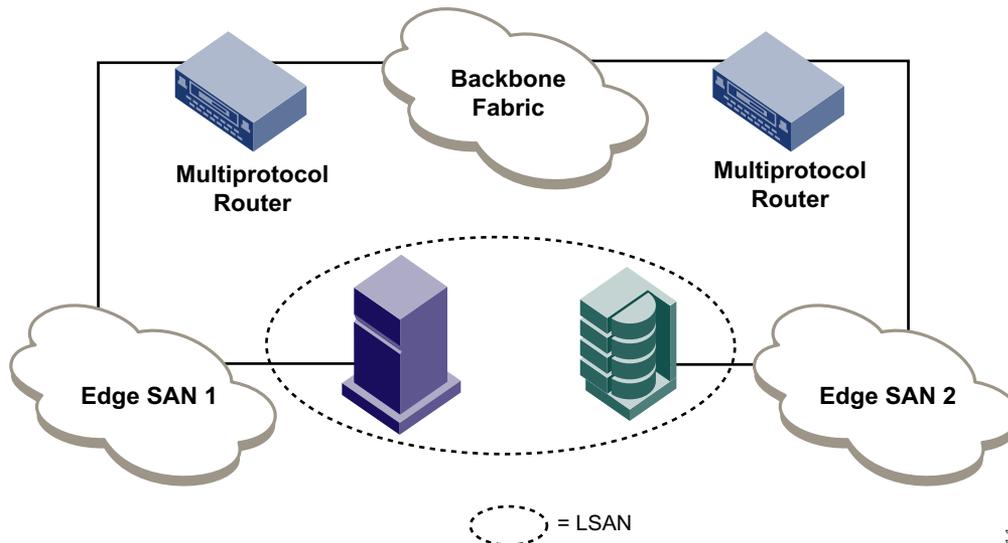


Figure 5 Edge SANs connected through a backbone fabric

Proxy devices

In an isolated SAN, the physical topology of the interconnections between nodes and switches is closely modeled by the logical topology of the connections between PIDs; this is not so in a meta-SAN. With an MP Router in a meta-SAN, a node is projected into the logical topology as a *proxy device*. This is a *proxy topology*.

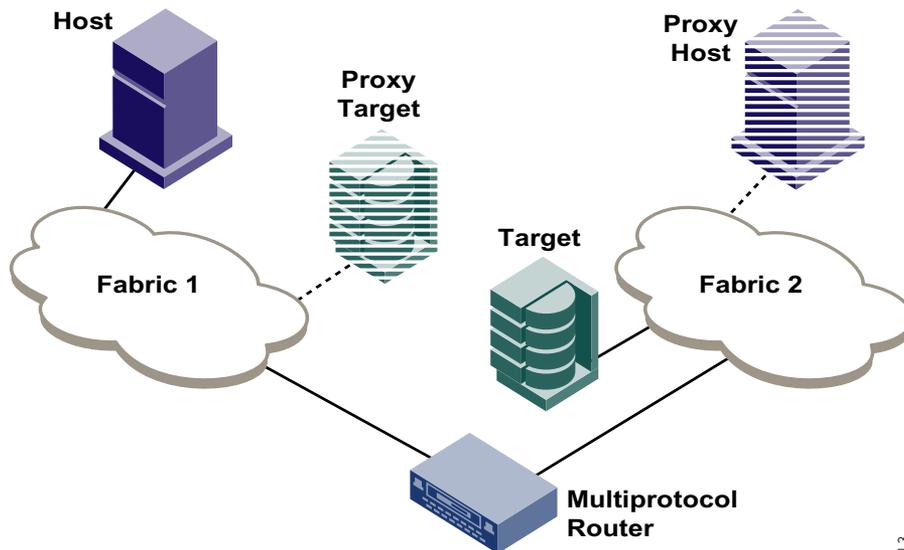
An MP Router achieves interfabric device connectivity by creating proxy devices (hosts and targets) in attached fabrics that represent real devices in other fabrics. For example, a host in Fabric 1 can communicate with a target in Fabric 2 as follows:

- A proxy target in Fabric 1 represents the real target in Fabric 2.
- Likewise, a proxy host in Fabric 2 represents the real host in Fabric 1.

The host discovers and sends Fibre Channel frames to the proxy target. The MP Router receives these frames, translates them appropriately, and then delivers them to the destination fabric for delivery to the target.

The target responds by sending frames to the proxy host. Hosts and targets are exported from the edge SAN to which they are attached and, correspondingly, imported into the edge SAN reached through Fibre Channel routing.

Figure 6 illustrates this concept.



13

Figure 6 Proxy topology

LSANs and zoning

An LSAN is defined by a zone in an edge fabric. You can define and manage LSANs using HP Advanced Zoning or HP Fabric Manager.

Zones are locally defined. Names and memberships, with the exception of hosts and targets exported from one fabric to another, do not need to be coordinated between the edge fabrics. For example, in [Figure 5](#) on page 47, when the zones for Edge SAN 1 are defined, you do not need to consider the zones in Edge SAN 2, and vice versa.

The exception is hosts and targets that are shared between the two fabrics. Although an LSAN zone is managed using the same tools as any other zone on the edge fabric, two characteristics distinguish an LSAN zone from a conventional zone:

- A required name convention. The name of an LSAN zone begins with `LSAN_`. This is case insensitive; for example, `lsan_` is equivalent to `LSAN_`, `Lsan_`, and so on.
- Members must be identified by their port WWN, because PIDs are not necessarily unique across fabrics.

Because zoning is enforced by all involved edge fabrics, any communication from one edge fabric to another must be allowed by the zoning setup on both edge fabrics. If the SANs are under separate administrative control, the separate administrators maintain access control.

You can perform LSAN zoning for both HP and for McDATA fabrics using HP LSAN zoning.

The following example procedure illustrates how LSAN zones control the devices that can communicate with one another. The example shows the creation of two LSAN zones (called `lsan_zone_fabric1` and `lsan_zone_fabric2`), which involve the following devices:

- Switch1 and the host are in fabric1.
- Switch2 and targets A and B are in fabric2.
- Switch1 is connected to the MP Router using an EX_Port.
- Switch2 is connected to the MP Router using another EX_Port.
- Host is at WWN 10:00:00:00:c9:2b:6a:2c (connected to switch1).
- Target A is at WWN 22:00:00:20:37:c3:11:71 (connected to switch2).
- Target B is at WWN 22:00:00:20:37:c3:1a:8a (connected to switch2).

1. Connect to switch1, log in as admin, and create and enable the first LSAN zone:
 - a. Use the `nsShow` command to list the WWN of the host (10:00:00:00:c9:2b:6a:2c).
 - b. Use the `zoneCreate` command to create the LSAN zone called `lsan_zone_fabric1`, which includes the host.
 - c. Use the `zoneAdd` command to add Target A to the LSAN zone.
 - d. Use the `cfgCreate` and `cfgEnable` commands to create and enable the LSAN zone configuration.

```
switch1:admin> nsshow
{
  Type Pid      COS      PortName      NodeName
  TTL(sec)
  N      030000;    2,3;10:00:00:00:c9:2b:6a:2c;20:00:00:00:c9:2b:6a:2c; na
  FC4s: FCP
  NodeSymb: [34] "Emulex LP952 FV3.81A3 DV5-4.82A4 "
  Fabric Port Name: 20:00:00:60:69:c0:05:89
  The Local Name Server has 1 entry }
switch1:admin> zonecreate "lsan_zone_fabric1", "10:00:00:00:c9:2b:6a:2c"
switch1:admin> zoneadd "lsan_zone_fabric1", "22:00:00:20:37:c3:11:71"
switch1:admin> cfgcreate "zone_cfg", "lsan_zone_fabric1"
switch1:admin> cfgenable "zone_cfg"
Starting the Commit operation...
cfgEnable successfully completed
```

2. Connect to switch2, log in as admin, and create and enable the second LSAN zone:
 - a. Use the `nsShow` command to list Target A (22:00:00:20:37:c3:11:71) and Target B (20:00:00:20:37:c3:1a:8a).
 - b. Use the `zoneCreate` command to create the LSAN zone called `lsan_zone_fabric2`, which includes the host (10:00:00:00:c9:2b:6a:2c), Target A, and Target B.
 - c. Use the `cfgCreate` and `cfgEnable` commands to create and enable the LSAN zone configuration.

```
switch2:admin> nsshow
{
  Type Pid      COS      PortName      NodeName      TTL(sec)
  NL      6100e8;    3;22:00:00:20:37:c3:11:71;20:00:00:20:37:c3:11:71; na
  FC4s: FCP [SEAGATE ST318304FC 0003]
  Fabric Port Name: 20:00:00:60:69:c0:21:26
  NL      6100ef;    3;22:00:00:20:37:c3:1a:8a;20:00:00:20:37:c3:1a:8a; na
  FC4s: FCP [SEAGATE ST318304FC 0003]
  Fabric Port Name: 20:00:00:60:69:c0:21:26

  The Local Name Server has 2 entries }
switch2:admin> zonecreate "lsan_zone_fabric2", "10:00:00:00:c9:2b:6a:2c;
22:00:00:20:37:c3:11:71;22:00:00:20:37:c3:1a:8a"
switch2:admin> cfgcreate "zone_cfg", "lsan_zone_fabric2"
switch2:admin> cfgenable "zone_cfg"
Starting the Commit operation...
cfgEnable successfully completed
```

3. Connect to the MP Router, log in as admin, and use the following commands to display information about the LSAN zones:
 - a. Use `lsanZoneShow -s` to show the LSAN zones.
 - b. Use `fcrPhyDevShow` to show the physical devices in the LSAN zones.
 - c. Use `fcrProxyDevShow` to show the proxy devices in the LSAN zones.

On the MP Router, the host and Target A are imported, because both are defined by `lsan_zone_fabric1` and `lsan_zone_fabric2`. However, Target B, which is defined by `lsan_zone_fabric2`, is not imported because `lsan_zone_fabric1` does not allow Target B to be imported.

```

router:admin> lsanzoneshow -s
Fabric ID: 1 Zone Name: lsan_zone_fabric1
    10:00:00:00:c9:2b:6a:2c Exist
    22:00:00:20:37:c3:11:71 Imported
Fabric ID: 2 Zone Name: lsan_zone_fabric2
    10:00:00:00:c9:2b:6a:2c Imported
    22:00:00:20:37:c3:11:71 Exist
    22:00:00:20:37:c3:1a:8a Exist
router:admin> fcrphydevshow
Device          WWN          Physical
Exists          PID
in Fabric
-----
    1    10:00:00:00:c9:2b:6a:2c  030000
    2    22:00:00:20:37:c3:11:71  6100e8
    2    22:00:00:20:37:c3:1a:8a  6100ef
router:admin> fcrproxydevshow
Proxy          WWN          Proxy          Device          Physical          State
Created       PID          Exists         Device          PID
in Fabric     in Fabric
-----
    1    22:00:00:20:37:c3:11:71  02f001         2             6100e8  Imported
    2    10:00:00:00:c9:2b:6a:2c  01f001         1             030000  Imported

```

When a PLOGI, PDISC or ADISC arrives at the MP Router, the SID and DID of the frame are checked. If they are LSAN-zoned at both SID and DID edge fabrics, the frame is forwarded to the DID. If they are not zoned, the unauthorized frames are dropped.

Fibre Channel NAT and phantom domains

Within an edge fabric (or across a backbone fabric), the standard Fibre Channel FSPF protocol determines how frames are routed from the host or target node through the fabric to the destination node. When frames leave the fabric through an MP Router, the frames are routed to an EX_Port. Fibre Channel fabrics require that all ports be identified by a unique PID. Within a single fabric, fabric formation guarantees that domain IDs are unique, and so a PID formed by a domain ID and area number is unique within a fabric. However, the domain IDs and PIDs in one fabric might be duplicated within another fabric, just as IP addresses unique to one private network are likely to be duplicated within another private network.

In an IP network, a network router can maintain network address translation (NAT) tables to replace private network addresses with public addresses when a packet is routed out of the private network, and to replace public addresses with private addresses when a packet is routed from the public network to the private network. The Fibre Channel routing equivalent to this IP network address translation is *Fibre Channel network address translation* (FC-NAT). Using FC-NAT, the proxy devices in a fabric can have different PIDs than the real devices that they represent, allowing the proxy devices to have appropriate PIDs for the address space of their corresponding fabric.

Proxy devices are presented to the fabric as being topologically attached to *phantom domains* created by the FC-FC Routing Service. The MP Router creates two types of phantom domains for each edge fabric accessed:

- Each EX_Port projects a unique *front phantom domain* (front domain).
- Each EX_Port also projects one *translate phantom domain* (xlate domain) for each edge fabric accessed through it. All EX_Ports connected to an edge fabric use the same xlate domain number for a remote edge fabric; this value persists across switch reboots and fabric reconfigurations. Xlate domains are presented as being connected topologically behind one or more front domains. This allows redundant EX_Ports or MP Routers with redundant paths to remote fabrics to present redundant paths to proxy devices to an edge fabric.

Phantom domains are like logical switches that appear to be connected to an edge fabric through the EX_Ports. The combination of front domains and xlate domains allows routing around path failures, including path failures through the routers. The multiple paths to an xlate domain provide additional bandwidth.

Connecting multiple EX_Ports to an edge fabric

You can connect multiple EX_Ports to the same edge fabric. The EX_Ports can all be on the same MP Router, or they can be on multiple routers. Multiple EX_Ports create multiple paths for frame routing. Multiple paths can be used in two different, but compatible, ways:

- Failing over from one path to another
- Using multiple paths in parallel to increase effective data transmission rates

Routing failover is automatic, but it can result in frames arriving out of order when frames take different routes. The MP Router can force in-order delivery, although frame delivery is delayed immediately after the path failover.

You can control whether in-order delivery is required; see [“Specifying frame delivery order”](#) on page 28.

Source EX_Ports can balance loads across multiple destination EX_Ports attached to the same edge fabric using exchange IDs from the routed frames as keys to distribute the traffic. This feature is enabled automatically when the exchange-based trunking feature is enabled. See Chapter 8, [“Using ISL trunking,”](#) for details on enabling this feature.

Matching fabric parameters

By default, the MP Router detects, autonegotiates, and configures the fabric parameters without user intervention.

As an option, you can configure these parameters manually. Use the `configure` command on a switch in the edge fabric to change the fabric parameters of a switch in the edge fabric. Then use the `portCfgEXPort` command to change the fabric parameters of an EX_Port on the MP Router.

For details see [“Supported modes”](#) on page 59.

Fabric parameter settings must be the same on EX_Ports and on the fabrics to which they are connected: E_D_TOV (error-detect timeout value), R_A_TOV (resource-allocation timeout value), and PID format. You can set the PID format on an EX_Port when you configure an interfabric link.

The default values for E_D_TOV and R_A_TOV for an EX_Port match those values on HP StorageWorks switches. You need to adjust these parameters for an EX_Port only if you have adjusted them for the fabric.

The default values for R_A_TOV and E_D_TOV are the recommended values for all but very large fabrics (ones requiring four or more hops) or high-latency fabrics (such as ones using long-distance FCIP links).

SAN scalability

Fabrics have scalability limits; for example, the maximum number of Name Server entries allowed limits the number of devices that can join a fabric. When you try to merge fabrics, this limit might be exceeded. But using Fibre Channel routing, you need not merge the fabrics. Instead, you can selectively import only those specific hosts or targets to be shared between the fabrics, and thus use Name Server entries more efficiently.

For example, suppose the maximum number of Name Server entries is 1024. Consider Fabric A with 700 devices and Fabric B with 600 devices. If you try to merge Fabrics A and B, the result requires 1300 Name Server entries, which exceeds the maximum of 1024. Using Fibre Channel routing, you can perform either of the following:

- Import up to 324 hosts or targets to Fabric A from Fabric B ($700 + 324 = 1024$)
- Import up to 424 hosts to Fabric B from Fabric A ($600 + 424 = 1024$).

Scaling of SANs that comprise multiple vendor fabrics is limited by the native fabric scalability.

Configuring an interfabric link

Before you configure an IFL, consider the following rules:

- Devices, whether Fibre Channel initiators, Fibre Channel targets, or others, connected to an edge fabric cannot be in an LSN with devices attached to a backbone fabric.
- You cannot configure both IFLs (EX_Ports) and ISLs (E_Ports) from one MP Router to the same edge fabric.

Configuring an interfabric link involves stopping ports and cabling them to other fabrics, configuring those ports for their intended use, and then starting the ports. The following procedure demonstrates how to configure for both edge and backbone connections:

1. Make sure that the FC-FC Routing Service license is installed.

See “[Licensed features](#)” on page 24.

2. Issue the `portStop` command to stop all ports that are to be configured.

The ports must be stopped to prevent any backbone fabrics from merging with edge fabrics when their ports are cabled into a fabric. For example, to stop ports 1 through 3:

```
router:admin> portstop 1-3
port 1 stopped.
port 2 stopped.
port 3 stopped.
```

3. Cable the ports to the edge and backbone fabrics.
4. Configure each port that connects to an edge fabric as an EX_Port.
5. Issue the `portCfgExPort` command to:
 - Enable the EX_Port mode.
 - Set the fabric ID (avoid using fabric ID 1, which is the default for backbone connections).
 - Set the PID format, if necessary.
 - Start routing between HP and McDATA fabrics.

For example:

```
portcfgexport port -a 1 -f fabricID -p pidformat
```

where:

- a Sets the EX_Port to enabled (1) or disabled (2).
- f Sets the fabric ID (1 to 128). Each edge fabric must have a unique ID, and EX_Ports connected to the same edge fabric must have the same ID. The default value is the port number plus 2.
- m Sets the connectivity mode (0 = Brocade Native, 1 = McDATA Open)
- p Port ID format (1 = core, 2 = extended-edge, and 3 = native). The value must match the edge fabric setting. The default value is 1.

The fabric ID settings for all EX_Ports attached to the same fabric must match. EX_Ports attached to more than one edge fabric must configure a different fabric ID for each edge fabric.

The PID mode for the MP Router (the backbone fabric PID mode) and the edge fabric PID mode do not need to match, but the PID mode for the EX_port and the edge fabric to which it is attached must match. The various edge fabrics may have different PID modes.

6. Configure each port that connects to a backbone fabric as an E_Port.
7. Assign the backbone fabric ID:
 - a. Issue the `switchDisable` command to disable the switch.
 - b. Issue the `fcrConfigure` command. At the prompt, enter the fabric ID, or press **Enter** to specify the default fabric ID (1). Make sure the fabric ID is different from that set for edge fabrics, and avoid using fabric IDs 2 through 17, which are the defaults for edge fabrics.
Multiple MP Routers attached to the same backbone fabric must have the same backbone fabric ID.
 - c. Issue the `switchEnable` command to enable the switch. For example:

```
router:admin> switchdisable
Switch is being disabled.....
Switch disabled
router:admin> fcrconfigure
FC Router parameter set. <cr> to skip a parameter
Backbone fabric ID: (1-128)[1]
You must run switchenable to return the switch back to online state.
router:admin> switchenable
Switch enabled
```

8. Establish LSAN zones; see Chapter 7, “[Creating and maintaining zones](#),” for details.
9. Issue the `portStart` command to start the ports:

```
router:admin> portstart 1-3
port 1 started
port 2 started
port 3 started
```

10. Issue the `portShow` command to verify that each port is configured correctly:

```
router:admin> portshow 2
port 2 info
Configuration Current
State: ENABLED UP
Type : FC FC
Link Status: ENABLED UP
Topology: LOOP LOOP
Speed: 1G 1G
LinkCost: 1000 (STATIC)
WWN: 20:01:00:05:1e:00:30:00

Licensed : YES

Diag result : PASSED

inFrames:14
outFrames:12
inOctets:1472
outOctets:1168
```

XPath OS and Secure Fabric OS

Beginning with XPath OS 7.4.x, the MP Router supports routing between secure fabric employing HP Secure Fabric OS with non-secured fabrics through Challenge-Handshake Authentication Protocol (DH-CHAP). Secure Fabric OS is an optional, licensed product that provides customizable security restrictions through local and remote management channels on an HP StorageWorks fabric.

Secure Fabric OS uses digital certificates based on PKI or Diffie-Hellman with DH-CHAP shared secrets to provide switch-to-switch authentication.

For details about Secure Fabric OS, see the *HP StorageWorks Secure Fabric OS administrator guide*.

To determine that an EX_Port is connected to a Secure Fabric OS fabric, issue the `portShow` or `portCfgExPort` command, as described in the *HP StorageWorks XPath OS 7.4.x command reference guide*.

Configuring a secure XPath OS DH-CHAP secret

While Secure Fabric OS supports the SLAP, FCAP, and DH-CHAP authentication protocols to communicate with each switch, XPath OS 7.4.x supports only DH-CHAP.

The MP Router does not initiate DH-CHAP authentication requests; it responds to DH-CHAP requests only from the edge switch to which it is connected—in this case, the Secure Fabric OS switch.

DH-CHAP is set on the Fabric OS side of the configuration, rather than the XPath OS side. As soon as you connect the MP Router to a Secure Fabric OS switch, DH-CHAP authentication is initiated.

The DH-CHAP secrets are configured both on the Secure Fabric OS switch and the MP Router. Each entry specifies the WWN of the peer to which it is connected. For example, on the MP Router, specify the WWN of the Secure Fabric OS switch and the secrets. On the Secure Fabric OS switch, specify the WWN of the front domain (EX_Port) and the secrets.

If a Switch Connection Controls (SCC) policy is defined, the WWN of the front domain (EX_Port) that is connected to the Secure Fabric OS switch should be present in the SCC list. See the *HP StorageWorks Secure Fabric OS user guide* for details about the SCC and other Secure Fabric OS features.

Configuring a DH-CHAP secret word on the MP Router

When configuring the DH-CHAP secret on the MP Router, you must know the WWN for the Fabric OS switch to set as the peer entry. To find out the WWN of this switch, log in to the switch as admin and issue the `switchShow` command. For example:

```
switch:admin> switchshow
switchName:      fcr_mojo_14
switchType:      16.2
switchState:     Online
switchMode:      Native
switchRole:      Principal
switchDomain:     99
switchId:        fffc63
switchWwn:       10:00:00:60:69:80:05:14
switchBeacon:    OFF
Zoning:          ON (cfg1)
port 0: id N2 No_Light
port 1: -- N2 No_Module
port 2: -- N2 No_Module
port 3: -- N2 No_Module
port 4: -- N2 No_Module
port 5: id N2 No_Light
port 6: -- N2 No_Module
port 7: -- N2 No_Module
value = 8 = 0x8

switch:admin>
```

When you have the necessary information, configure the secret words on the MP Router.

1. Log in to the MP Router with administrative privileges.

2. Issue the `secAuthSecret` command:

```
secAuthSecret --set
```

The secret must consist of 8 to 40 characters.

Setting up secret keys does not initiate DH-CHAP authentication. DH-CHAP authentication is performed whenever a port or a switch is enabled.

3. Follow the instructions provided on screen, as shown in the following example:

a. Enter the port or switch WWN.

b. Enter and confirm the peer secret.

This is the secret that authenticates at the peer.

c. Enter and confirm the local secret.

4. After you have added all the DH-CHAP secret information, press **Enter** to indicate that you have completed the secret key setup.

5. When prompted, enter `y`.

The DH-CHAP secret is now stored in the secret word database and is ready for use. For example:

```
router:admin> secAuthSecret --set
```

This command sets up secret keys for the DH-CHAP authentication.

The minimum length of a secret key is 8 characters and maximum 40

characters. Setting up secret keys does not initiate DH-CHAP

authentication. It is performed whenever a port or a switch is enabled.

Following inputs should be specified for each entry.

1. WWN for which secret is being set up.

2. Peer secret: The secret of the peer that authenticates to peer.

3. Local secret: The local secret that authenticates peer.

Press Enter to start setting up shared secrets >

Enter WWN, Domain, or switch name (Leave blank when done): **10:00:00:60:69:80:05:14**

Enter peer secret:

Re-enter peer secret:

Enter local secret:

Re-enter local secret:

Enter WWN, Domain, or switch name (Leave blank when done):

Are you done? (yes, y, no, n): [no] **y**

Saving data to key store... Done.

Configuring a DH-CHAP secret on the Fabric OS switch

You must know the front domain WWN of the MP Router to use as the peer entry when setting the secret word on the Fabric OS switch. To discover the correct WWN, log in to the router and issue the `portCfgExport` command as shown in the following example:

```
router:admin> portcfgexport 10
      Port  10  info
Admin:                enabled
State:                NOT OK
Pid format:           core(N)
Operate mode:         Brocade Native
Edge Fabric ID:       17
Preferred Domain ID: 35
Front domain WWN:     50:00:51:e1:57:b0:0e:0a
Fabric Parameters:    Auto Negotiate
R_A_TOV:              Not Applicable
E_D_TOV:              Not Applicable

Authentication Type: None
DH Group: N/A
Hash Algorithm: N/A
Edge fabric's primary WWN: N/A
Edge fabric's version stamp: N/A
```

When you have the necessary information, configure the secret word on the Fabric OS switch.

1. Log in to the switch with admin privileges.
2. Issue the `secAuthSecret` command:

```
secAuthSecret --set
```

The secret must be between 8 and 40 characters.

Setting up secret keys does not initiate DH-CHAP authentication. DH-CHAP authentication is performed whenever a port or a switch is enabled.

3. Follow the instructions provided on screen, as shown in the following example:
 - a. Enter the port or switch WWN.
 - b. Enter and confirm the peer secret.
This is the secret that authenticates at the peer
 - c. Enter and confirm the local secret.
4. After you have added all the DH-CHAP secret information, press **Enter** to indicate that you have completed the secret key setup.
5. When prompted, enter `y`.

The DH-CHAP secret is now stored in the secret word database and is ready for use. For example:

```
router:admin> secAuthSecret --set  
This command sets up secret keys for the DH-CHAP authentication.  
The minimum length of a secret key is 8 characters and maximum 40  
characters. Setting up secret keys does not initiate DH-CHAP  
authentication. It is performed whenever a port or a switch is enabled.  
  
Following inputs should be specified for each entry.  
1. WWN for which secret is being set up.  
2. Peer secret: The secret of the peer that authenticates to peer.  
3. Local secret: The local secret that authenticates peer.  
Press Enter to start setting up shared secrets >  
Enter WWN, Domain, or switch name (Leave blank when done):  
50:00:51:e1:57:b0:0e:0a  
Enter peer secret:  
Re-enter peer secret:  
Enter local secret:  
Re-enter local secret:  
Enter WWN, Domain, or switch name (Leave blank when done):  
Are you done? (yes, y, no, n): [no] y  
Saving data to key store... Done.
```

Viewing a DH-CHAP secret word database

1. Log in to the MP Router as admin.
2. At the prompt, issue the following command:

```
secauthsecret --show
```

You should see output similar to the following:

```
admin:switch> secauthsecret --show  
  
WWN                               DId  Name  
-----  
10:00:00:60:69:80:05:14 1     switch
```

For details about the `setAuthSecret` command, see the *HP StorageWorks Secure Fabric OS user guide*.

Setting a proxy PID

When an MP Router is first configured, the PIDs for the proxy devices are automatically assigned. Proxy PIDs (as well as phantom domain IDs) persist across reboots.

The most common situation in which you would set a proxy PID is when you replace MP Router hardware. You can minimize disruption to the edge fabrics by setting the proxy PIDs to the same values used with the old hardware.

The `fcProxyConfig` command displays or sets the persistent configuration of proxy devices. Used with the `-s` (slot) option, it can also influence the assignment of the `xlate` domain port number (which is used to determine the area field of the PID) and the `AL_PA` field. Like the PIDs in a fabric, a proxy PID must be unique. If the `slot` argument results in a duplicate PID, it is ignored. See the `fcProxyConfig` command in the *HP StorageWorks XPath OS 7.4.x command reference guide* for more details.

Use the `fcXlateConfig` command to display or assign a preferred domain ID to a proxy device. See the `fcXlateConfig` command in the *HP StorageWorks XPath OS 7.4.x command reference guide* for more details.

Monitoring resources

It is possible to exhaust resources, such as proxy PIDs. Whenever a resource is exhausted, the MP Router generates an error message. Messages are described in the *HP StorageWorks XPath OS 7.4.x system error messages reference guide*.

You can monitor MP Router resources with the `fcResourceShow` command:

```
router:admin> fcresourceshow
Daemon Limits:
                Max Allowed          Currently Used
-----
LSAN Zones:           1000             9
LSAN Devices:         10000            75
Proxy Device Slots:   10000             6

                WWN Pool Size        Allocated
-----
Phantom Node WWN:     4096              2
Phantom Port WWN:    16384             6

Port Limits:
Max proxy devices:  2000
Max NR_Ports:      1000
Currently Used(row 1: proxy, row2: NR_Ports):
  0   1   2   3   4   5   6   7   8   9  10  11  12  13  14  15
-----
  0   0   6   6   0   0   0   0   0   0   0   0   0   0   0   0
  0   0   4   4   0   0   0   0   0   0   0   0   0   0   0   0
router:admin>
```

See the *HP StorageWorks XPath OS 7.4.x command reference guide* for details about the command.

Routing ECHO

The MP Router and Fibre Channel Routing let you route the ECHO generated when an `fcPing` command is issued on a Fabric OS-based switch, providing `fcPing` capability between two devices in different fabrics across the MP Router.

 **NOTE:** Each MP Router in a backbone fabric routing to McDATA fabrics must have XPath OS 7.4.x or later installed. Fabrics that consist entirely of HP devices work with MP Routers using XPath OS 7.1.2, 7.3.x, or 7.4.x.

Checking for Fibre Channel connectivity problems

1. On the Fabric OS switch, issue the `fcPing` command, which:

- Checks the zoning configuration for the two ports specified
- Generates an Extended Link Service (ELS) frame ECHO request to the source port specified and validates the response
- Generates an ELS ECHO request to the destination port specified and validates the response

Regardless of the device's zoning, the `fcPing` command sends the ELS frame to the destination port. A device can take any one of the following actions:

- Send an ELS Accept to the ELS request.
- Send an ELS Reject to the ELS request.
- Ignore the ELS request.

There are some devices that do not support the ELS ECHO request. In these cases, the device either does not respond to the request or sends an ELS reject. When a device does not respond to the ELS request, further debugging is required; however, do not assume that the device is not connected to the Fibre Channel.

For details about the `fcPing` command, see the *HP StorageWorks Fabric OS command reference guide*.

Connecting to McDATA SANs

XPath OS 7.4.x furnishes the MP Router with the ability to connect to McDATA fabrics in McDATA Open mode. Consult the following sections:

- [Connectivity features](#), next
- [Supported modes](#), page 59
- [Configuring the fabrics for interconnectivity](#), page 61

Connectivity features

XPath OS 7.4.x lets you connect an HP fabric to a McDATA fabric. Because of the high degree of connectivity HP provides, the devices across the remote fabrics can be shared.

Interconnectivity benefits

Connected SANs provide additional functionality not possible with segregated SANs. Some of these include:

- Island consolidation—Use the XPath OS 7.4.x MP Router to connect isolated McDATA and HP Fabric and share devices.
- Backup consolidation—Consolidate backup solutions across HP and McDATA fabrics.
- Manageable large scale storage network—Use the XPath OS MP Router to localize traffic while connecting devices in the meta-SAN. This provides a large number of fabrics with a large number of devices.
- Sharing across a FCIP link—Share devices between HP and McDATA fabric over campus Ethernet or over long-distance IP links beyond 1000 km.
- Sharing across a long-distance FC link—Share devices between HP fabrics over long-distance FC links as far as 300 km.
- LUN sharing—Use a high-end RAID array connected to a McDATA director to share targets with an HP fabric; just connect one McDATA director port to an MP Router EX_Port and the one EX_Port to the HP fabric.

Connectivity

Check for the latest list of tested and approved equipment, see the *HP StorageWorks SAN Design Guide* at <http://h18006.www1.hp.com/products/storageworks/san/documentation.html>.

Scalability

The connectivity limitations of a meta-SAN comprising HP and McDATA fabrics are defined on each side by the scalability of the individual fabric.

Every EX_Port connected to a McDATA fabric uses one domain ID per link. This can potentially cause problems because McDATA supports a maximum of 31 domains in a fabric (in McDATA Fabric 1.0 mode). One translation domain per edge fabric is also created.

Check for the latest list of tested and approved equipment, see the *HP StorageWorks SAN Design Guide* at <http://h18006.www1.hp.com/products/storageworks/san/documentation.html>.

See the McDATA fabric documentation for McDATA scalability limitations.

Supported modes

XPath OS 7.4.x supports connectivity with McDATA Open mode. To allow interconnectivity with McDATA SANs, the CLI command `portCfgExport` uses the `-m` parameter to indicate the connectivity mode. Valid parameters are described in [Table 6](#).

Table 6 portCfgExPort -m interop parameters

Value	Description
0	HP Native (default)
1	McDATA Open mode

See the *HP StorageWorks XPath OS 7.4.x command reference guide* for details about the portCfgExPort command and other XPath OS commands.

Once the port is properly configured and connected, log in to the MP Router and issue the switchShow command. For example:

```
router:admin> switchshow
Switch Name   :  routerA
Switch State  :  Online
Switch Type   :  38.2
Switch Role   :  Subordinate
Switch Domain:  1
Switch ID     :  FFFC01
Switch WWN    :  10:00:00:05:1e:15:02:00
beacon status:  OFF
zoning        :  OFF

FC router BB Fabric ID: 1

Port Media Speed State      Info
=====
0   id   N1   Online  EX_PORT 10:00:08:00:88:40:50:ac "" (fabric id = 10)
1   id   N1   Online  EX_PORT 10:00:08:00:88:40:50:ac "" (fabric id = 10)
2   id   N2   Online  E_PORT 10:00:00:05:1e:13:83:00 "routerB" (upstream)

(output truncated)
13  --   AN   No_Module
14  id   N2   Online EX_PORT 10:00:00:60:69:80:04:98 "sw12kA21" (fabric id = 11)
15  id   N2   Online EX_PORT 10:00:00:60:69:80:04:98 "sw12kA21" (fabric id = 11)
router:admin>
```

The McDATA switch connected to the MP Router's EX_Port is visible. On the McDATA side, the show command also shows the front domain.

If the LSAN is configured and the proxy devices are created, the proxy device shows up in the Name Server of the edge fabric. The xlate domain also shows in the fabric of the edge fabric.

Configuring the fabrics for interconnectivity

When connecting an HP fabric with a McDATA fabric using the MP Router, you must configure the switch on both fabrics as well as the MP Router. This configuration is described in the following sections:

- [Configuring the MP Router](#), next
- [Preparing the HP StorageWorks switch for connectivity](#), page 63
- [Configuring the McDATA fabric for interconnection](#), page 64
- [Completing the configuration](#), page 68

Configuring the MP Router

When configuring your HP StorageWorks fabric to connect to a McDATA fabric, you must perform some preparation on the MP Router:

1. Using the `version` command, make sure that XPath OS 7.4.x is installed on the MP Router:

```
router:admin> version
=====
Installed Packages:
=====
Package Name: xpath_os_v7.4.0_prealpha1_bld17
Install Date: Apr 14, 2005 18:48

router:admin>
```

2. On the MP Router, stop the EX_Port being configured (the one connected to the HP StorageWorks switch) by issuing the `portStop` command.

You can verify that the port has been stopped by issuing the `portShow` command for the port.

In the following example, port 14 is used on the HP side to connect the HP fabric to a McDATA fabric. The McDATA fabric is connected from port 1.

```
router:admin> portshow 14
      port   14  info
      Configuration   Current
Name   :           port 14
State:   STARTED      DOWN
Type   :           FC
Link Status: DISABLED  DOWN
Topology: P-P         P-P
Speed:   AN           AN
LinkCost: AUTO
Distance: L0          L0
WWN:    20:01:00:05:1e:16:1f:08
Licensed : NO
Diag result : PASSED
inFrames: 0
outFrames: 0
inOctets: 0
outOctets: 0
discards: 0

router:admin>
```

3. Using the `portCfgExPort` command, configure the port as an EX_Port and provide a Fabric ID (FID). If no FID is specified, multiple links to the same fabric are assigned different FIDs. This results in a front domain oversubscription.

For the HP StorageWorks fabric, use the `-p` flag of the `portCfgExPort` command to match the PID setting for the fabric to which you are connecting. This port will connect to the HP StorageWorks switch.

The following example sets port 14 to admin-enabled, assigns a Fabric ID of 10, and sets the port to Core PID. For complete information about any XPath OS command, see the *HP StorageWorks XPath OS 7.4.x command reference guide*.

```
switch:admin> portcfgexport 14 -a 1 -f 10 -p 1
```

- Restart the port by issuing the `portStart` command.
- Still on the MP Router, use the `portStop` command to stop the EX_Port that is to be used to connect to the McDATA switch.

```
router:admin> portstop 1
port 1 stopped.
router:admin>
```

- Issue the `portCfgExport` command to configure the port as an EX_Port with a different FID within the McDATA Fabric PID mode.

This port can now connect to a McDATA switch in McDATA Open mode.

The following example sets port 1 to admin-enabled, assigns a Fabric ID of 11, and sets the McDATA connection mode to McDATA Fabric. For complete information about XPath OS commands, see the *HP StorageWorks XPath OS 7.4.x command reference guide*.

```
router:admin> portcfgexport 1 -a 1 -f 11 -m 2
```

- Restart the port by issuing the `portStart` command.
- You can now physically attach your ISLs from the MP Router to the switches. ISLs apply only to HP switches that are not connected as an edge fabric (IFLs). When a McDATA switch is present, it is assumed that you are creating an edge fabric. For details about interswitch linking, see Chapter 8, “[Using ISL trunking](#).”
- Capture a SAN profile of the McDATA and HP SANs, identifying the number of devices in each SAN. By projecting the total number of devices and switches expected in each fabric when the LSANs are active, you can quickly determine the status of the SAN by issuing the commands `nsAllShow` and `fabricShow` on the HP fabric and using SAN Pilot (see [Figure 7](#)) or EFCM to gather similar information for the McDATA fabric. For example, on the HP side:

```
router:admin> fabricshow
Switch ID   Worldwide Name           Enet IP Addr   Name
-----
   1: fffc01 10:00:00:05:1e:15:02:00 192.168.11.25  "routerA"
 100: fffc64 10:00:00:05:1e:13:83:00 192.168.12.25 >"routerB"

The Fabric has 2 switches

router:admin>
```

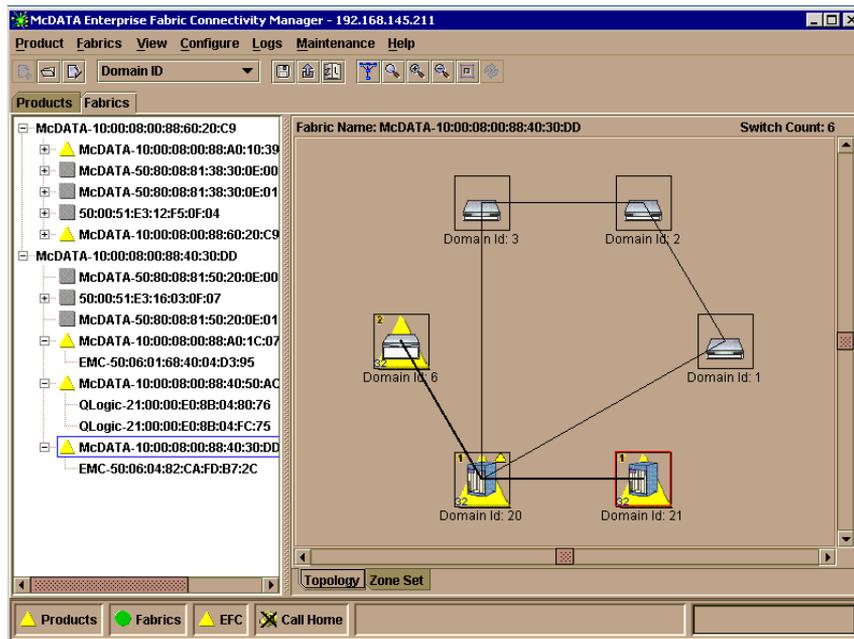


Figure 7 SAN Pilot

Preparing the HP StorageWorks switch for connectivity

Now that the MP Router is configured to connect to a McDATA fabric, you must create your LSA and zones for the SAN. Either of the following procedures may be used.

1. Create a telnet connection to the HP StorageWorks switch.
You can also use Advanced Web Tools to perform this procedure.
2. Configure the LSA, using the LSA_n_xxxx naming schema.
In the following example, two LSAs are created for two devices: LSA_host_1 and LSA_Host_2.

```
switch:admin > alicreate "lsan_host_1", "10:00:00:00:c9:08:05:00"
switch:admin > alicreate "lsan_host_2", "10:00:00:00:c9:08:04:00"
```

3. Create the zone.
In the following example, the zone LSA_1 is created, and the LSAs created in the previous step are added to it.:

```
switch:admin > zonecreate "LSAN_1", "lsan_host_1; lsan_host_2"
```

4. Issue the `cfgShow` command to verify that the zones are correct:

```
switch:admin> cfgshow
Defined configuration:
zone: LSA_1 lsan_host_1; lsan_host_2
alias: lsan_host_1
      10:00:00:00:c9:08:05:00
alias: lsan_host_2
      10:00:00:00:c9:08:04:00

Effective configuration:
no configuration in effect
```

5. Issue the `cfgCreate` command to create the domain and add the zone:

```
switch:admin> cfgcreate "Domain1", "LSAN_1"
```

6. Issue the `cfgEnable` command to enable the zone configuration.:

```
switch:admin> cfgenable "Domain1"

zone config "Domain1" is in effect
Updating flash ...
```

7. Optional: Reissue the `cfgShow` command to verify that the zoning is correct.:

```
switch:admin> cfgshow
Defined configuration:
cfg:   Domain1 LSAN_1
zone:  LSAN_1  lsan_host_1; lsan_host_2
alias: lsan_host_1
      10:00:00:00:c9:08:05:00
alias: lsan_host_2
      10:00:00:00:c9:08:04:00

Effective configuration:
cfg:   Domain1
zone:  LSAN_1  10:00:00:00:c9:08:05:00
      10:00:00:00:c9:08:04:00
```

Alternately, use the following procedure:

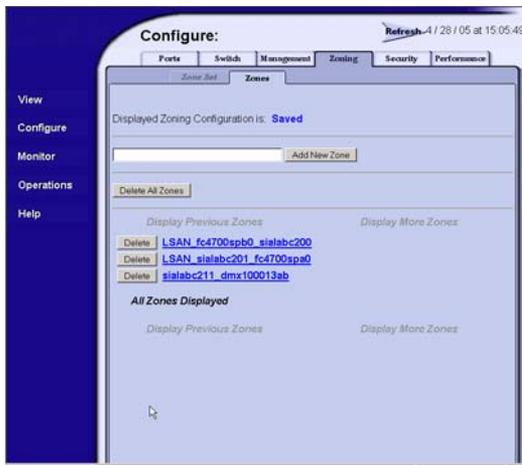
1. Create a telnet or Advanced Web Tools connection to the HP StorageWorks switch.
2. Configure the LSAN, using the LSAN_xxxx naming schema, and append the new LSAN to the active zoneset.
3. Enable zone configuration that now includes the new LSAN.

Configuring the McDATA fabric for interconnection

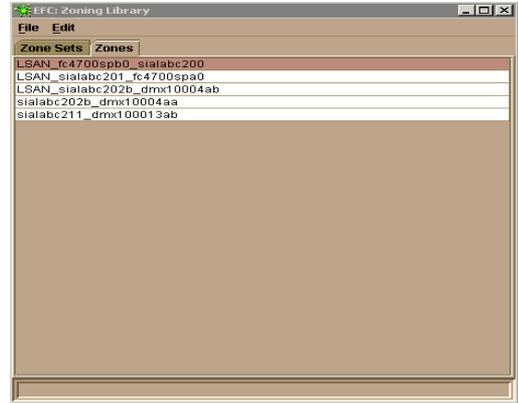
To ensure connectivity with the HP fabric, the McDATA fabric must be prepared for use with the MP Router. Either of the following procedures may be used.

 **NOTE:** The procedures described in this section were current when the document was written, but may have changed slightly since then. For the most up-to-date information, see the McDATA documentation about zone configuration on McDATA fabrics.

1. Log in to SAN Pilot or EFCM (see [Figure 8](#)).
ED5000 uses only EFCM.
2. From the Configure menu, select **Zones (Zoning Library** on EFCM).
3. Tab to **Zones**.



SAN Pilot



EFCM

Figure 8 SAN Pilot and EFCM zones

4. Enter the desired name in the **Zone Name** box using the LSAN_xxxx naming schema.
 - In EFCM, move to the list of ports and nodes, and highlight the devices to include in the LSAN.
 - In SAN Pilot, click **Add New Zone** and then select the new zone to display the **Modify Zone** tab (see [Figure 9](#)). Add the desired devices to the zone.

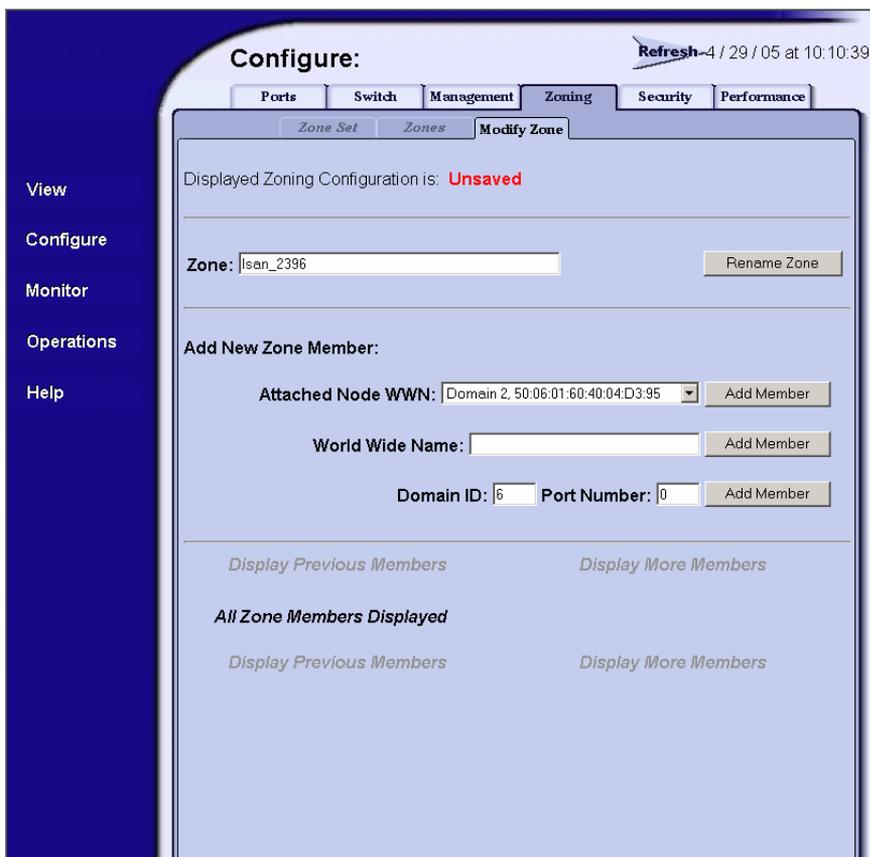


Figure 9 Modify Zone tab

5. To add devices that are connected to the HP fabric, enter the WWN into the **World Wide Name** field and then click **Add Member** (see [Figure 10](#)).

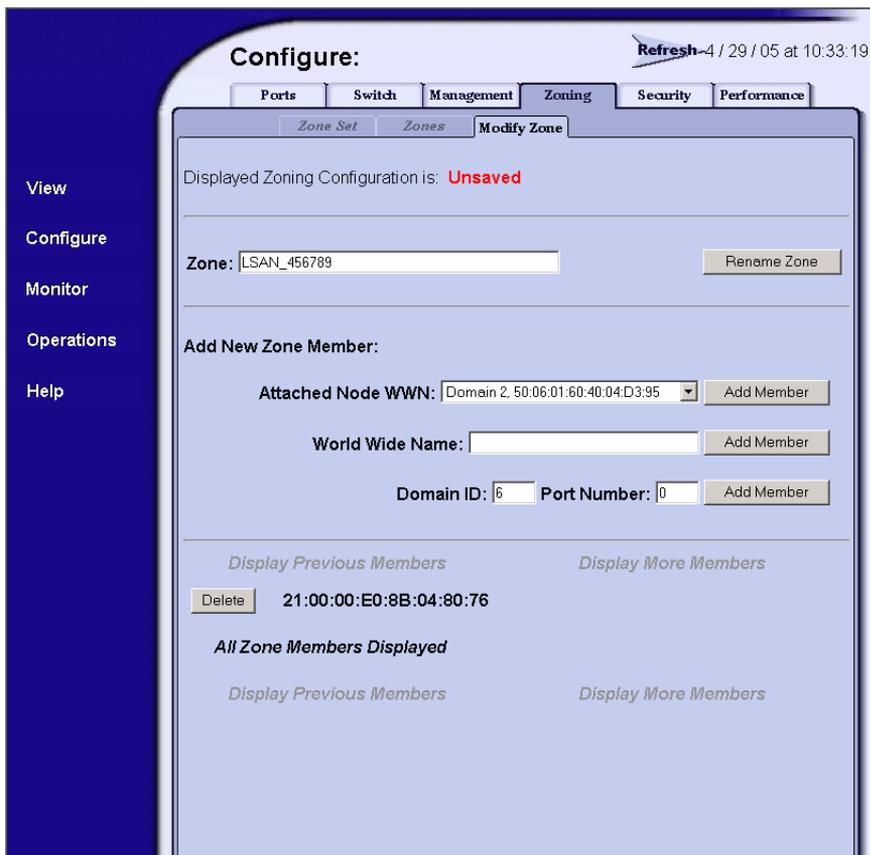


Figure 10 World Wide Name box

If you are using EFCM, use **Add Detached Node** to enter the WWN port name (see Figure 11).

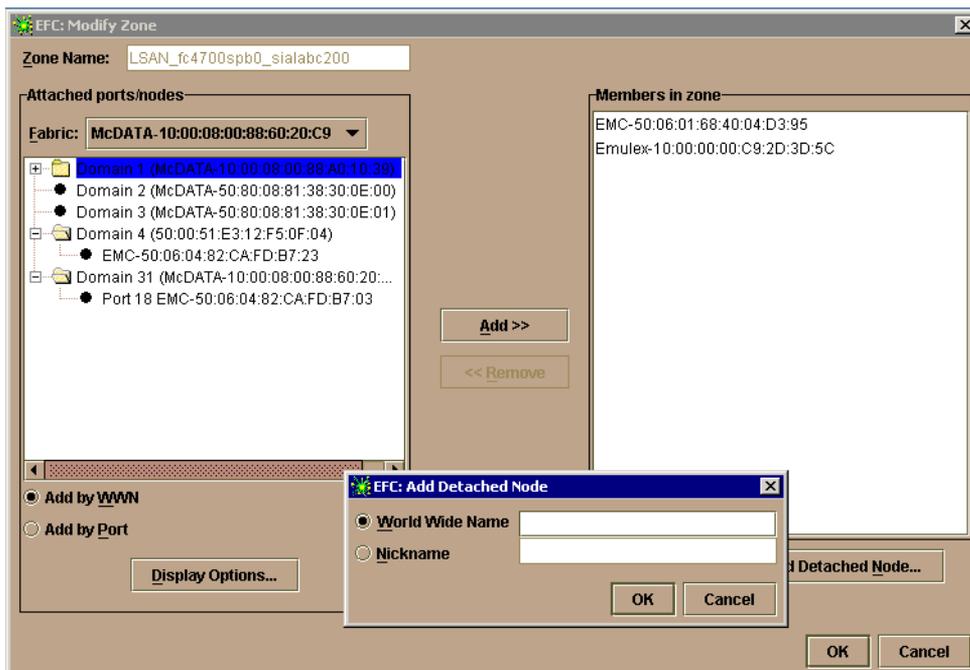


Figure 11 Modify zone window

6. Move to the **Zone Set** tab in SAN Pilot.

If you are using EFCM, or the Zoneset Library window, tab to **Zone Sets** and select **File > New**.

7. Enter a name for the zoneset in the **Zone Set Name** box.
8. Select the zone to include in this zoneset and click **Add Zone Set**.
The steps for EFCM are similar.
9. In SAN Pilot, click **Save and Activate Zoning Configuration**.
In EFCM, return to the main window and select **Configure**, and then select **Activate Zone Set** to launch the zoneset activation window (Figure 12).

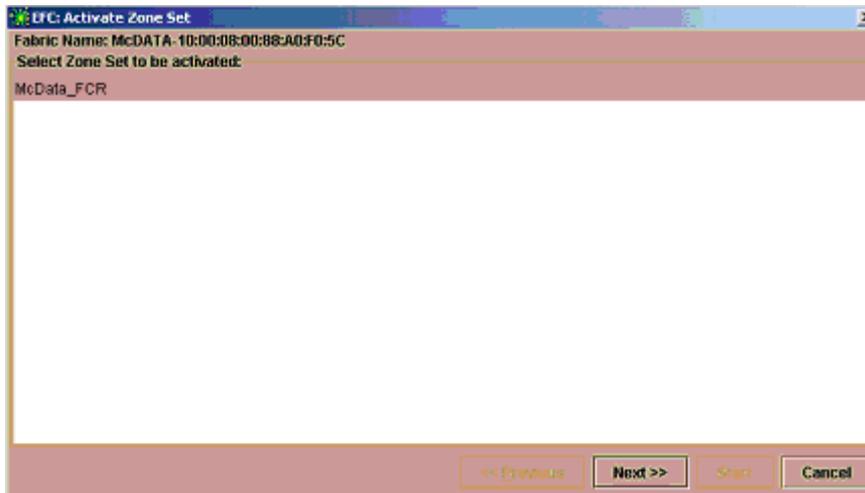


Figure 12 Activate zone set

10. Highlight the zoneset to be activated and then click **Next**.
11. Click **Next** again, and then click **Start** to activate your zoneset.
12. Regardless of the method used, verify that the new zoneset containing your LSAN has been added.

Alternately, use the following procedure:

1. Create the LSAN, using the LSAN_xxxx naming scheme.
2. Append the newly created zoneset to a currently active zoneset.
3. Activate the updated zoneset.

LSAN zoning

An LSAN is defined by a zone in an edge fabric. When zoning an LSAN containing multiple manufacturers' fabrics (such as an HP-McDATA SAN), you must use port WWNs. Because port IDs are not necessarily unique across fabrics, you cannot use the *domain,port* method of identification.

For more details about LSAN zoning, see "[LSANs and zoning](#)" on page 48.

If the LSAN devices appear only in one of the fabrics within a multiple-fabric SAN, use the following procedure to correct the problem:

1. Log in to each fabric and verify that all the devices are physically logged in.
2. Verify that the devices are properly configured in the LSAN zone in both edge fabrics.
3. Issue the `fabricShow` command on the HP fabric.
4. Use McDATA's EFCM or SAN Pilot to verify the McDATA fabric, including the front and translation domains.

5. Move back to the MP Router and issue the `fcrProxyDevShow` command to verify that the devices are configured and exported. For example:

```

router:admin> fcrProxyDevShow
Proxy          WWN          Proxy    Device    Physical    State
Created       in Fabric   PID      Exists   in Fabric
-----
   10    20:00:00:01:73:00:59:dd  05f001    12      610902  Imported
   10    21:00:00:e0:8b:04:80:76  02f002    11      340713  Imported
   10    50:06:01:68:40:04:d3:95  02f001    11      660713  Imported
   11    10:00:00:00:c9:2d:3d:5c  020001    10      011500  Imported
   11    50:06:01:60:40:04:d3:95  020002    10      011400  Imported
router:admin>

```

6. Create a telnet connection and configure the connection to capture text.
7. Issue the `supportShow` command, and save the output.
8. If the fabric does not appear:
 - a. Disable the EX_Port on the connected fabric.
 - b. Issue the `portLogClear` command for the port.
 - c. Enable the port on the MP Router.
 - d. Issue the `portLogDump` command for the port, capturing the output.
 - e. Use the `portLogDump` tool to troubleshoot the problem, using the command output.

 **NOTE:** If an EX_Port connecting an MP Router and an edge fabric is disabled due to an error, the error causing that port's most recent disabled state appears in the `switchShow` command output. This error appears until that port comes back online, even after the cables have been detached from the port. To remove the error listing in the `switchShow` output, reboot the MP Router. Possible errors displayed under this condition include:

```

ISW_PORT_ERR_TYPE_UNKNOWNISW_PORT_ERR_TYPE_LICENSE
ISW_PORT_ERR_TYPE_FTAG_OVERISW_PORT_ERR_TYPE_FTAG_CONFLICT
ISW_PORT_ERR_TYPE_FOWNER_CONFLICTISW_PORT_ERR_TYPE_ZONE_RESOURCE
ISW_PORT_ERR_TYPE_PORT_STATE_TOISW_PORT_ERR_TYPE_AUTHN_REJECT
ISW_PORT_ERR_TYPE_SEC_FCS_LISTISW_PORT_ERR_TYPE_SEC_FAILURE
ISW_PORT_ERR_TYPE_INCOMPATIBLE_MODEISW_PORT_ERR_TYPE_SEC_SCC_LIST

```

Completing the configuration

After the MP Router, HP StorageWorks switch, and McDATA switch have been prepared for use, complete the configuration using the following procedure:

1. Physically connect the EX_Port that you configured for the HP switch to the MP Router.
2. Log into the HP switch as admin.
3. Issue the `fabricShow` command.

New domains should be visible for each ISL (front domain) that connects the HP switch to the MP Router, and one domain should be visible for the translate domain.

4. Physically connect the configured MP Router's EX_Port to the McDATA switch.
5. Start SAN Pilot (or EFCM) and select the fabric for the McDATA switch.
6. View the fabric topology.

New domains should be visible for each ISL that connects the McDATA switch to the MP Router, and one domain should be visible for the translate domain.

In EFCM, the McDATA switch should appear green, and the front domains (as well as the translate domain) are grayed out and inaccessible: EFCM cannot manage them. Tab to **Zone** and verify that the zoneset configuration is correct: A blue icon beside each entry indicates that the devices are logged into the fabric.

7. Log in to the HP StorageWorks switch and issue the `nsAllShow` command.
All the devices from both LSANs should appear in the output. If they do not, issue the `cfgShow` command to verify your zone configuration.

5 Using the FCIP Tunneling Service

The optional FCIP Tunneling Service enables Fibre Channel frames to *tunnel* through IP networks by dividing frames, encapsulating the result in IP packets upon entering the tunnel, and then reconstructing them as they leave the tunnel. The FCIP Tunneling Service is discussed in the following sections:

- [Synchronizing time](#), page 71
- [Configuring an FCIP interswitch link](#), page 72
- [Disabling and enabling an FCIP interswitch link](#), page 76

XPath OS supports eight FCIP interswitch links between two MP Routers. All 16 links can be used as FCIP links with the remaining links going to other routers.

Ports configured for FCIP are called *virtual E_Ports* (VE_Ports). After you configure the VE_Ports on the two MP Routers, an FCIP connection is established between them.

Fibre Channel frame encapsulation on one VE_Port and the reconstruction of Fibre Channel frames on the other VE_Port are transparent to the initiator and target, but the administration of VE_Ports is different than the administration of other Fibre Channel port types.

[Figure 13](#) shows a portion of a Fibre Channel network using FCIP. The FCIP interswitch link (VE_Ports connected through the IP network) joins the two SANs (Fabric 1 and Fabric 2) into parts of a larger SAN.

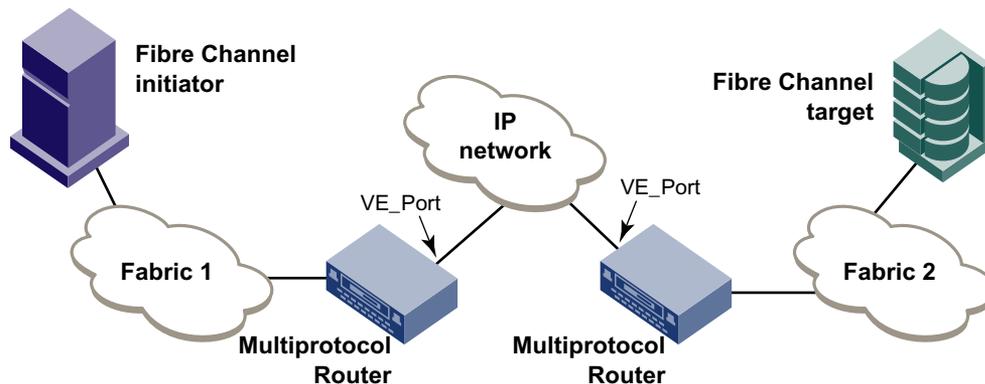


Figure 13 Network using FCIP

Synchronizing time

Fibre Channel framing and signaling standards require that a Fibre Channel fabric define and enforce a frame lifetime. The standard defined timeout values are R_A_TOV and E_D_TOV.

The FCIP and Fibre Channel frame encapsulation RFCs add wide-area-network (WAN) timeout value (WAN_TOV). When a Fibre Channel frame is encapsulated, the encapsulation header includes a timestamp. The current time, WAN_TOV, and the timestamp are checked when the frame is reconstructed upon leaving the tunnel. If the lifetime of the frame exceeds the WAN_TOV, the frame is discarded.

If the time settings for the MP Routers at the two ends of an FCIP interswitch link are not synchronized, frame lifetime calculations are not meaningful. If the receiving MP Router time lags the sending router time, frames that should be discarded are not discarded; if the receiving router time leads the sending router time, frames are discarded that should not be discarded.

You can avoid this problem by synchronizing both ends to the same NTP server. Use the `tsClockServer` command to specify an NTP server for time service. See "[Synchronizing time with an NTP server](#)" on page 23 for a procedure. Use the `configure` command to set WAN_TOV to a value higher than 1000 ms. Enable WAN_TOV enforcement on an FCIP port with the `portCfgFcip` command.

Configuring an FCIP interswitch link

You must configure both the local and the remote MP Routers to enable an FCIP ISL. This configuration requires the use of the `portCfgFcip` command (not the `portCfgExPort` command).

If the two MP Routers have a direct connection between them (no router in between), the IP addresses for both routers must be in the same subnet, and the default gateway setting is not required.

If two MP Routers are connected through a router, the IP addresses for both MP Routers must be in different subnets and the default gateway must be set on both the MP Routers. (Typically, the default gateway is the router's IP address.)

The basic steps for configuring an FCIP ISL between two MP Routers are:

1. Configure the local MP Router:
 - a. Issue the `portStop` command to stop a port.
 - b. Issue the `portType` command to configure the port for GigE.
 - c. Issue the `portCfgGige` command to configure the IP network parameters and ensure that the port is set as an FCIP GigE port.
 - d. Issue the `portCfgFcip` command to enable the VE_Port.
2. Configure the remote MP Router:
 - a. Issue the `portStop` command to stop a port.
 - b. Issue the `portType` command to configure the port to be a GigE port.
 - c. Issue the `portCfgGige` command to configure the IP network parameters.
 - d. Issue the `portCfgFcip` command to enable the VE_Port and configure the VE_Port to initiate the setup of the FCIP ISL.

NOTE: For all Continuous Access EVA configurations using FCIP, you must set trunking to `SID_DID` or none. See the FCIP `portcfgfcip` command.

NOTE: All ports on an MP Router are optical ports. To connect a VE_Port to a copper network, use an optical-to-copper converter.

Figure 14 illustrates the topology and port values used in the configuration example in the following two procedures.

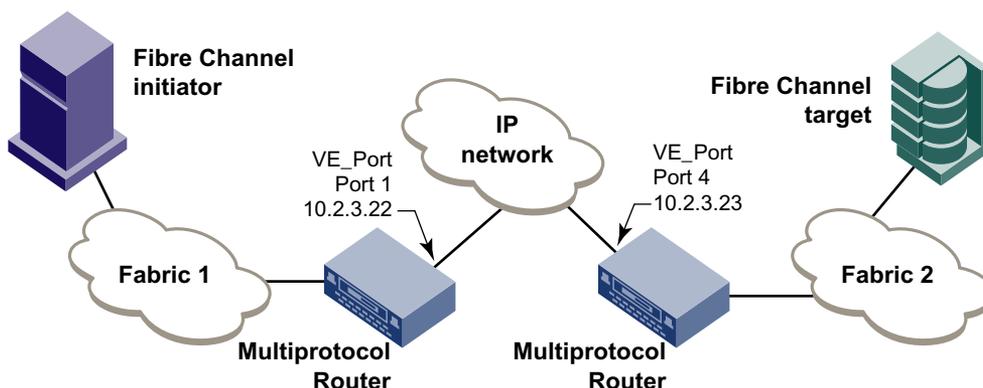


Figure 14 Example for configuration procedures

The configuration example uses the following values:

- Port 1 (IP address 10.2.3.22) on the local MP Router
- Port 4 (IP address 10.2.3.23) on the remote MP Router
- IP network subnet mask 255.255.255.0

The example sets the port type to GigE, sets the GigE port parameters, and configures and enables the FCIP ISL.

See the `portType` command in the *HP StorageWorks XPath OS 7.4.0 command reference guide* for additional information on configuring GigE ports and the `portCfgGige` command for additional information on configuring GigE parameters for the ports.

A given port can be configured for only one FCIP ISL.

Configuring the local MP Router

1. Stop the port:

```
router:admin> portstop 1
port 1 stopped.
```

2. Configure Port 1 to be a GigE port:

```
router:admin> porttype 1 g
port 1 set to type GIGE
```

3. Set the GigE port parameters.

The following example is for local and remote ports on the same subnet:

```
router:admin> portcfggige 1 -i 10.2.3.22 -n 255.255.255.0 -v 1 -p fcip
port 1 proto set to: fcip
port 1 proto ver set to: 1
```

If the local and remote ports are on different subnets, use the `-g` option to specify the gateway:

```
router:admin> portcfggige 1 -i 10.2.3.22 -n 255.255.255.0 -g 10.2.3.1
-v 1 -p fcip
```

4. Start the port:

```
router:admin> portstart 1
port 1 started
```

5. Disable the FCIP tunnel:

```
router:admin> portcfgfcip 1 -a 2
port 1 admin status set to : 2
```

6. Set the local port to listen for the remote port:

```
router:admin> portcfgfcip 1 -i 0.0.0.0
```

7. Re-enable the FCIP tunnel:

```
router:admin> portcfgfcip 1 -a 1
port 1 admin status set to : 1
```

Configuring the remote MP Router

1. Stop the port:

```
router:admin> portstop 4
port 4 stopped.
```

2. Configure the port as a GigE port:

```
router:admin> porttype 4 g
port 4 set to type GIGE
```

3. Set the GigE port parameters.

The following example is for local and remote ports on the same subnet:

```
router:admin> portcfggige 4 -i 10.2.3.23 -n 255.255.255.0 -v 1 -p fcip
port 4 proto set to: fcip
port 4 proto ver set to: 1
port 4 ipaddress set to: 10.2.3.23
port 4 net mask set to: 255.255.255.0
```

If the local and remote ports are on different subnets, use the `-g` option to specify the gateway:

```
router:admin> portcfggige 4 -i 10.2.4.23 -n 255.255.255.0 -g 10.2.4.1
-v 1 -p fcip
```

4. Start the port:

```
router:admin> portstart 4
port 4 started
```

5. Disable the FCIP ISL:

```
router:admin> portcfgfcip 4 -a 2
port 1 admin status set to : 2
```

6. Set the remote port to send to the local port:

```
router:admin> portcfgfcip 4 -i 10.2.3.22
```

7. Re-enable the FCIP ISL:

```
router:admin> portcfgfcip 4 -a 1
port 1 admin status set to : 1
```

Verifying that the connection is up

Use the `rnPing` command. The following example pings the remote port at IP address 10.2.3.23, using a packet length of 1200 bytes. The example assumes the MP Routers are on the same subnet. You can issue the `portShow` command to see the port configuration and the `fcipShow` command to see FCIP protocol information:

```

router:admin> rnping 1 10.2.3.23 -1 1200
Pinging 10.2.3.23
Reply from 10.2.3.23: bytes=1200 time<14ms TTL=255
The rnping is completed
router:admin> portshow 1
      port 1 info
          Configuration  Current
Name :      port_1
State:      STARTED      UP
Type :      GIGE         GIGE
Link Status:  ENABLED     UP
IP addr:     10.2.3.23    10.2.3.23
Net mask:    255.255.255.0 255.255.255.0
Default route: 0.0.0.0 0.0.0.0
Mac address: 00:05:1e:31:30:51

Protocol:    fcip ver 1    fcip ver 1
Base OS version : 7.1.0.0.BS25
IPS version  : 7.1.0.0.BS25
Diag result  : PASSED
router:admin> fcipshow 1
----- fcip protocol info(port 1 1) -----
          Configured          Current
State:      UP                UP
Local IP addr: 10.2.3.23      10.2.3.23
Remote IP addr: 0.0.0.0      0.0.0.0
WAN_TOV timeout enabled          enabled
Remote WWN:    00:00:00:00:00:00:00:00 00:00:00:00:00:00:00:00
Time sync state          synchronized

in_frame_ip:      15370
in_frame_fc:      15372
out_frame_ip:     15372
out_frame_fc:     15370
in_octet_ip:      1296008
in_octet_fc:      693260
out_octet_ip:     693260
out_octet_fc:     742688
error_frame_ip:   0
error_frame_fc:   0
error_resync:     0
drop_frame_fc:    0
drop_frame_ip:    0
frame_timeout:    0
authen_failure:   0

```

Disabling and enabling an FCIP interswitch link

You can disable a configured FCIP ISL by changing the admin state to disabled on both the local and the remote MP Router, using the following syntax:

```
portcfgfcip portnumber -a 2
```

You can re-enable a disabled FCIP ISL by changing the admin state to enabled on both the local and remote MP Router, using the following syntax:

```
portcfgfcip portnumber -a 1
```

6 Using the iSCSI Gateway Service

The HP iSCSI Gateway Service facilitates communication between TCP/IP networks and Fibre Channel SANs. It displays iSCSI gateway configuration information across multiple MP Routers. The iSCSI Gateway Service is discussed in the following sections:

- [Summary of configuration steps](#), next
- [Configuring an iSCSI portal](#), page 78
- [Configuring iSCSI gateway zones](#), page 79
- [Configuring CHAP](#), page 80
- [Administering iSCSI configurations](#), page 80
- [Working with the WWN mapping table](#), page 82
- [Enabling and disabling failover](#), page 82

XPath OS supports eight iSCSI sessions per port and 96 sessions per MP Router.

Figure 15 shows a simple application of the iSCSI Gateway Service, in which an iSCSI initiator communicates with a Fibre Channel target through an iSCSI gateway. The gateway projects the initiator onto the Fibre Channel SAN, and projects the target onto the TCP/IP network. The iSCSI gateway converts iSCSI to FCP FC-4 protocol and maps iSCSI fully qualified names (IQNs) from TCP/IP to Fibre Channel in a process called *IQN-to-WWN mapping*.

The projection of an iSCSI initiator into the Fibre Channel SAN creates a *proxy Fibre Channel initiator*, the projection of a Fibre Channel target to the TCP/IP network creates a *proxy iSCSI target*. The proxy initiator fully participates in access controls, such as zoning, and the proxy target responds like a regular iSCSI target.

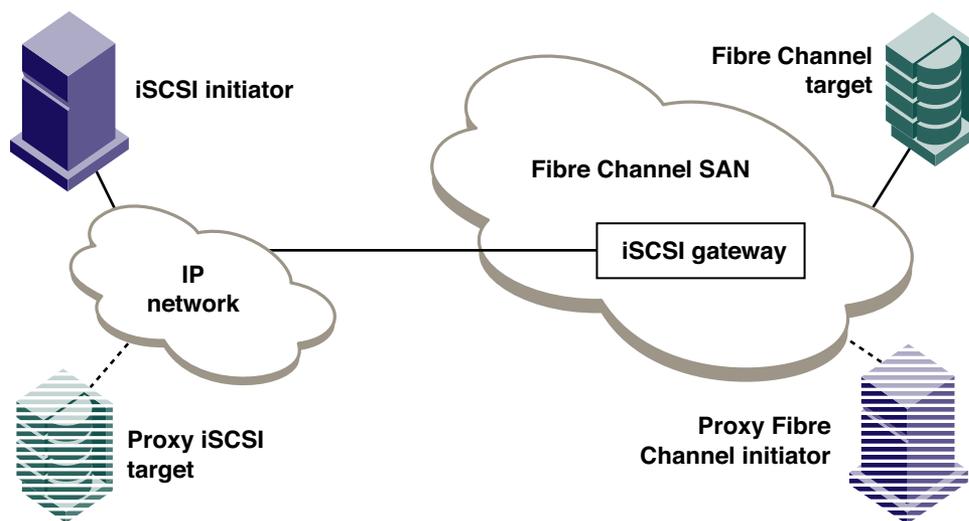


Figure 15 Simple application of the iSCSI Gateway Service

TCP/IP servers can use iSCSI arrays that are directly connected to the IP network; Fibre Channel servers connected to the SAN can use Fibre Channel arrays.

An iSCSI portal cannot be used for any other purpose, but other ports on the same MP Router can be used for any other HP multi-protocol routing service.

You can achieve access control by configuring iSCSI initiators and the iSCSI portal to use iSCSI challenge handshaking authentication protocol (CHAP). CHAP allows remote servers and clients to exchange authentication credentials securely.

If you define iSCSI portals on multiple MP Routers, iSCSI drivers that support failover can fail over from one portal to another if the connection to the first portal is lost.

If there are multiple iSCSI portals on one MP Router, configuration information, such as IQN-to-WWN mapping and CHAP secrets, is automatically shared among the portals. If there are iSCSI portals on more than one MP Router, you can use the IP fabric configuration server (iFCS) to control the sharing of iSCSI gateway configuration information across multiple MP Routers.

Summary of configuration steps

The configuration of an iSCSI gateway is illustrated in [Figure 16](#).

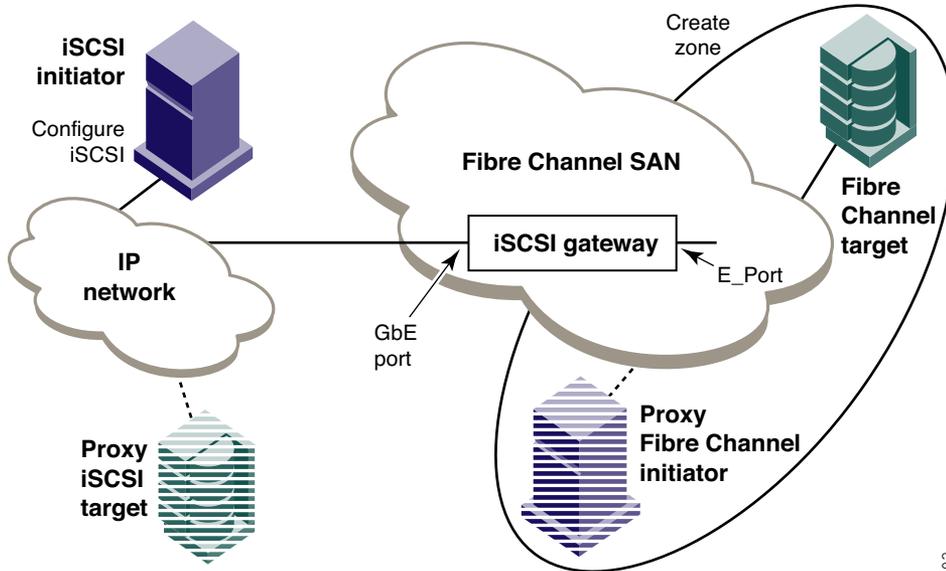


Figure 16 iSCSI gateway configuration

The steps required to configure an iSCSI gateway are summarized below and detailed in the referenced sections.

1. Configure an iSCSI portal on the MP Router.
Use the procedure in "[Configuring an iSCSI portal](#)" on page 78.
2. Optional: Create zones with the iSCSI initiators and the Fibre Channel targets as zone members.
See "[Configuring iSCSI gateway zones](#)" on page 79 for details.
3. Optional: Configure CHAP access control for the Fibre Channel targets.
Use the procedure in "[Configuring CHAP](#)" on page 80 to configure the shared secret for the Fibre Channel targets.
4. Install and configure the iSCSI initiator driver.
You must assign a unique IQN to the driver. If you configured CHAP for the target in [step 3](#), you must configure matching CHAP shared secrets. See the iSCSI driver documentation for details on how to configure the driver.

NOTE: All ports on an MP Router are optical ports. To connect a port to a copper network, use an optical-to-copper converter.

Configuring an iSCSI portal

Use the following procedure to configure a port as an iSCSI portal:

1. Issue the `portStop` command to stop the port.
2. Issue the `portType` command to set the port for GigE.
3. Issue the `portCfgGige` to set the GigE port parameters, ensuring that the port is set as an iSCSI GigE port.
4. Issue the `portStart` command to start the port.

5. Issue the `portShow` command to verify the configuration.

For example:

```
router:admin> portstop 1
port 1 stopped.
router:admin> porttype 1 g
port 1 set to type GIGE
router:admin> portcfggige 1 -i 192.168.0.10 -n 255.255.255.0 -g
192.168.0.1 -p iscsi
port 1 proto set to: iscsi
port 1 proto ver set to: 1
router:admin> portstart 1
port 1 started
router:admin> portshow 1
      port 1  info
                Configuration      Current
Name :         port_1
State:         STARTED              UP
Type :         GIGE                  GIGE
Link Status:   ENABLED              UP
IP addr:       192.168.0.10          192.168.0.10
Net mask:      255.255.255.0        255.255.255.0
Default route: 0.0.0.0              0.0.0.0
Mac address:   00:05:1e:31:7e:18

Protocol:      iscsi ver 13         iscsi ver 13

Licensed      : YES

Diag result   : PASSED
```

Displaying iSCSI portal information

1. Issue the `iscsiShow` command to display a summary of all iSCSI sessions and port counters.
2. Issue the `iscsiPortShow` command to display iSCSI sessions and port counters on a specific portal.

Configuring iSCSI gateway zones

Create and manage iSCSI gateway zones as you do any zones, but you must use iSCSI fully qualified names (IQN) to distinguish the iSCSI initiator members from other Fibre Channel zone members. The IQN format is:

```
iqn.iscsi_driver:initiator
```

where:

iscsi_driver is the name of the iSCSI driver

initiator is the name you assign to the initiator

For example, the following IQN indicates that an initiator called `initiator1` is using the iSCSI driver called `1991-05.com.microsoft`:

```
iqn.1991-05.com.microsoft:initiator1
```

For more information on creating zones, see Chapter 7, “[Creating and maintaining zones.](#)”

Configuring CHAP

The iSCSI standard supports access control with CHAP. You can configure the iSCSI gateway to use one-way authentication, where the target authenticates the initiator, or two-way authentication, where first the target authenticates the initiator and then the initiator authenticates the target.

 **NOTE:** The MP Router supports DH-CHAP only. It consists of the CHAP protocol combined with the Diffie-Hellman exchange.

Configuring one-way authentication

Issue the following command to configure one-way authentication for the target:

```
iscsiauthcfg -i targetiqn -c secret1
```

where *targetiqn* is the IQN for the Fibre Channel target and *secret1* is the shared secret.

Remember the shared secret; you need it to configure the iSCSI driver on the iSCSI initiator.

For example:

```
router:admin> iscsiauthcfg -i iqn.2002-12.com.brocade:21000004cf4c54e9 -c
1234abcd1234
Create [iqn.2002-12.com.brocade:21000004cf4c54e9, *****] successful.
```

Configuring two-way authentication

1. Configure one-way authentication (*secret1*) as described in “[Configuring one-way authentication.](#)”
2. Issue the following command to configure authentication for the initiator:

```
iscsiauthcfg -i iscsihostiqn -c secret2
```

where *iscsihostiqn* is the IQN for the iSCSI initiator and *secret2* is the shared secret, which must be different from *secret1*.

Remember the shared secret; you need it to configure the iSCSI driver on the iSCSI initiator.

For example:

```
router:admin> iscsiauthcfg -i iqn.1991-05.com.microsoft:isi154110 -c 5678abcd5678
Create [iqn.1991-05.com.microsoft:isi154110, *****] successful.
```

Removing a CHAP secret

Issue the following command:

```
iscsiauthcfg -d iqname
```

where *iqname* is the IQN of the initiator or the target.

To add or remove a DH-CHAP secret in Secure Fabric OS, see “[Configuring a secure XPath OS DH-CHAP secret](#)” on page 54.

Administering iSCSI configurations

When multiple iSCSI portals are defined on one MP Router, they share information, such as IQN-to-WWN mapping and CHAP secrets. If you have iSCSI portals defined on more than one MP Router, you can use the IP fabric configuration server (iFCS) to perform out-of-band sharing of iSCSI gateway configuration information across the MP Routers. The iFCS function distributes the IQN-to-WWN mapping of each iSCSI host and their shared CHAP secret configuration to all IP-aware switches in the fabric. This distribution enables iSCSI hosts to move from one switch to another switch within a fabric.

The MP Router on which you enable iFCS—the *primary iFCS router*—distributes and synchronizes the information to the other MP Routers, and continues to do so whenever there is a change in any IP storage configuration.

A secondary switch automatically becomes the primary if the current primary is removed from the fabric. When this occurs, the primary selection is based on the second and third least-significant bytes of the switch WWN. The secondary switch in the fabric with the larger value of those two bytes becomes the primary. For example, a secondary with a WWN of 10:00:00:05:1e:15:84:00 becomes primary over another secondary with a WWN of 10:00:00:05:1e:12:de:00 because 0x1584 is larger than 0x12de.

Enabling and disabling iFCS

The MP Router on which you enable iFCS becomes the primary iFCS router. The iFCS function is disabled by default.

Use the following commands to enable or disable iFCS:

Issue the `ifcsEnable` command to enable iFCS.

On the primary iFCS router, issue the `ifcsDisable` command to disable iFCS.

Displaying iFCS information

Once the primary iFCS router has distributed and synchronized the configuration information across MP Routers, you can display information about the entire iSCSI configuration from any MP Router in the fabric.

The types of information you can view are described in [Table 7](#).

Table 7 iFCS information

To display	Use the command	Parameter
iFCS information	<code>ifcsShow</code>	none
IQN-to-WWN mapping table	<code>iscsiWwnAlloc</code>	none
Zoning information	<code>iscsiWwnAlloc</code>	<code>-v</code>
CHAP information	<code>iscsiAuthCfg</code>	none

iFCS behavior during configuration download

The status of iFCS is indicated by the `iscsi.ifcsEnableStatus` entry in the configuration file.

If iFCS is enabled, the entry reads `iscsi.ifcsEnableStatus:YES`

If iFCS is disabled, the entry reads `iscsi.ifcsEnableStatus:NO`

[Table 8](#) summarizes the resulting iFCS status for a fabric when a configuration file is downloaded to that fabric.

Table 8 Effect of downloading a configuration file with iFCS enabled or disabled

iFCS status in configuration file	iFCS enabled in destination file	Result
<code>iscsi.ifcsEnableStatus:YES</code> (enabled)	YES	The switch joins the fabric as a secondary iFCS switch.
<code>iscsi.ifcsEnableStatus:NO</code> (disabled)	NO	The switch joins the fabric with iFCS disabled.
<code>iscsi.ifcsEnableStatus:YES</code> (enabled)	NO	The switch joins the fabric as the primary iFCS switch.
<code>iscsi.ifcsEnableStatus:NO</code> (disabled)	YES	The switch joins the fabric as a secondary iFCS switch.

Working with the WWN mapping table

The WWN is an HP Organizational Unique Identifier (OUI).

You can use the `iscsiWwnAlloc` command to assign WWNs.

Displaying the WWN list

Issue the `iscsiWwnAlloc` command with no arguments.

You can use the `aliShow`, `zoneShow`, and `cfgShow` commands with the `-i` option to display IQNs.

Adding an iSCSI initiator to the WWN mapping table

Issue the following command on the primary iFCS router, where `IQN_name` is the IQN of the initiator:

```
iscsiwwnalloc -i IQN_name
```

Removing an iSCSI initiator from the WWN mapping table

Issue the following command:

```
iscsiwwnalloc -d IQN_name
```

If you cannot log in to the MP Router from an iSCSI initiator, there might be an error in iSCSI-name-to-port-WWN mapping. Log in as admin. Issue the `iscsiWwnAlloc` command to display the WWN allocation map. If there are two WWN entries for the initiator, perform the following steps:

1. If the IQN is used in zones or aliases, remove it from them.
2. Delete the IQN using the following command, where `IQN_name` is the IQN of the initiator:

```
iscsiwwnalloc -d IQN_name
```
3. Re-create the IQN by issuing the following command on the primary iFCS router, where `IQN_name` is the IQN of the initiator:

```
iscsiwwnalloc -i IQN_name
```
4. If necessary, re-create the IQN in zones or aliases using the `zoneCreate` or `aliCreate` command.

Enabling and disabling failover

You can take advantage of the high availability (HA) failover feature by defining multiple iSCSI portals from a TCP/IP network to a fabric. In this configuration, iSCSI drivers that support failover can fail over from one portal to another if the connection to the first portal is lost.

For example, [Figure 17](#) shows three iSCSI portals connecting a TCP/IP network to a fabric. IP1, the primary, is on one MP Router (iSCSI Gateway 1); IP2 and IP3 are on a second MP Router (iSCSI Gateway 2). When HA failover is enabled, if the path including IP1 fails, the iSCSI initiator tries IP2, and if not successful, IP3.

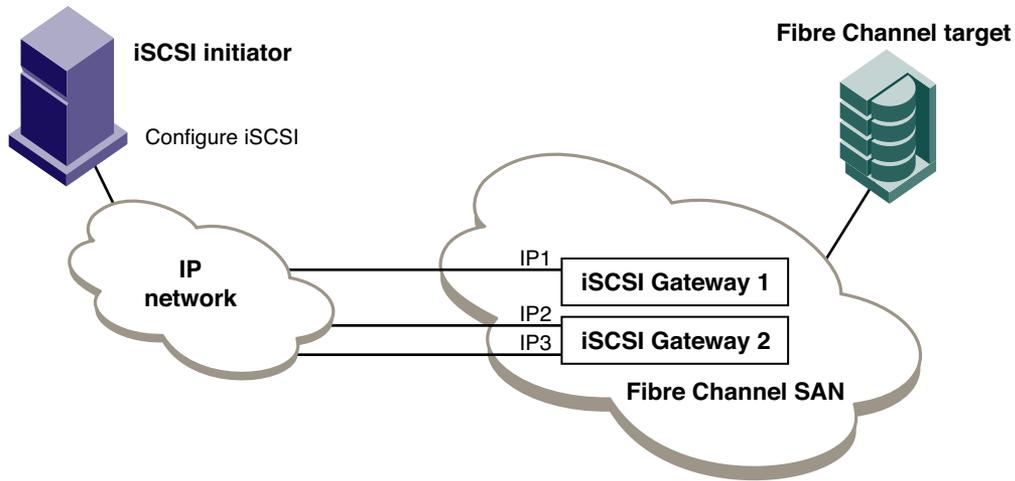


Figure 17 iSCSI high availability configuration

Enabling failover

1. Issue the `ifcsEnable` command to enable iFCS.
2. Issue the following command, where *WWN* is the WWN of the failover switch:

```
iscsifailoveradd WWN
```

With reference to [Figure 17](#), the following example shows how to specify that iSCSI Gateway 1 should fail over to iSCSI Gateway 2, whose WWN is 10:00:00:05:1e:15:a9:00.

```
router:admin> iscsifailoveradd 10:00:00:05:1e:15:a9:00
The failover switch is added
```

Disabling failover

On the primary iFCS router, issue the `iscsiFailoverDelete` command:

```
router:admin> iscsifailoverdelete
The failover switch is deleted
```


7 Creating and maintaining zones

This chapter provides procedures for using XPath OS zoning in the following sections:

- [Zoning terminology](#), page 86
- [Zoning enforcement](#), page 86
- [Configuring zones](#), page 87

You can use zones to create logical subsets of the fabric to accommodate environments such as closed user groups or functional areas within the fabric. Any zone object connected to the fabric can be included in one or more zones. Zone objects can communicate only with other objects in the same zone. For example, consider [Figure 18](#), which shows that:

- Three zones are configured, named Red, Green, and Blue.
- Loop 1 is a member of both the Red zone and the Green zone. Both Server 1 and Server 3 have access to Loop 1. Server 1 and Server 3 do not have access to each other.
- Loop 2 is not a member of any zone and is therefore inaccessible to other devices in the fabric.
- Server 2 can access the RAID 1 and RAID 2 devices, but not Loop 1, Server 1, or Server 3.
- The RAID 1 device is a member of the Blue zone and the Green zone. Server 2 and Server 3 have access to it, but Server 1 does not.
- Only Server 2 has access to the RAID 2 device.

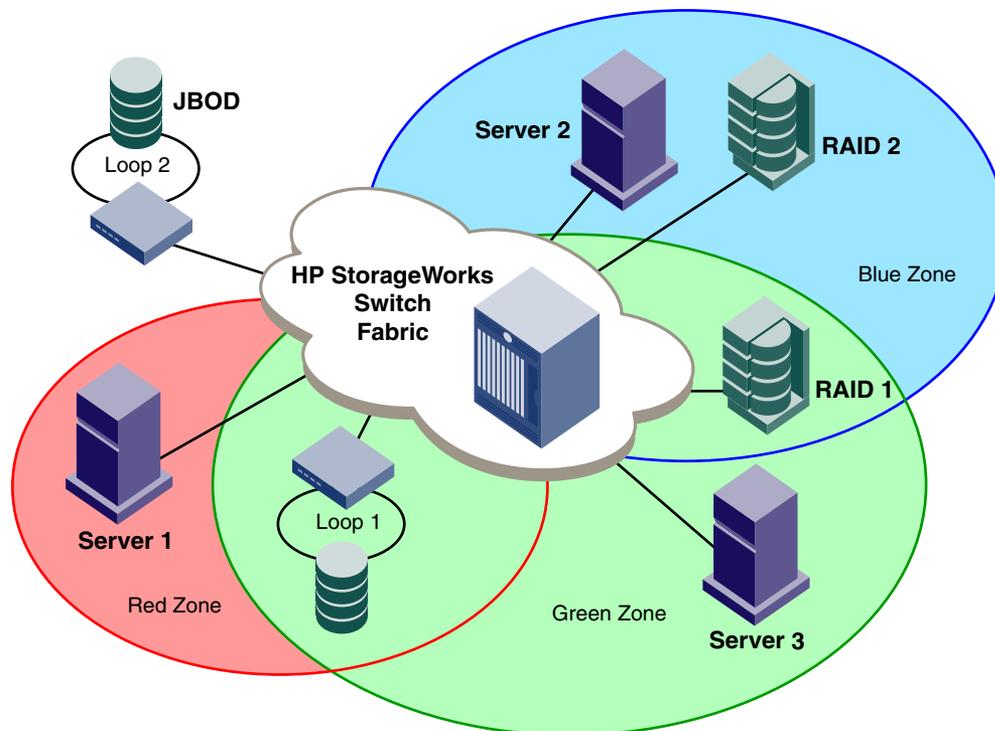


Figure 18 Sample zone configuration

A zone member can be specified by using a *domain,port* pair, node WWN, or the port WWN. The mechanisms needed to enforce zoning are discussed next.

NOTE: Zones used with an MP Router have additional member naming requirements. Use port WWNs for Fibre Channel attached nodes with MP Router LSAN zones.

Zoning terminology

A *zone* is a specified group of fabric-connected devices, also called *zone objects*. Any device or zone object connected to the fabric can be included in one or more zones. Zone objects within a zone can communicate only with other zone objects in the same zone.

After zoning is enabled, if a device is not explicitly defined in a zone, that device is isolated and is inaccessible to other devices in the fabric.

Zone objects are grouped into zones, and zones are grouped into a *zone configuration*. Zones can overlap; that is, a zone object can belong to more than one zone and a fabric can have multiple zones. You can create multiple zone configurations; however, only one configuration can be enabled in the fabric at a time.

Zoning enforcement

To design, configure, and administer Fibre Channel zoning on your fabric, you need to understand the difference between a *zone member specification* and the mechanism used for *zoning enforcement*. XPath OS supports two basic zoning enforcement mechanisms:

- Software-enforced zoning (also called *soft zoning*)
- Hardware-enforced zoning (also called *hard zoning*)

Soft zoning

In soft zoning, the Name Server, getting its information from the zone server, limits the targets returned to a Fibre Channel initiator only to those that Fibre Channel initiator can communicate with (as specified by the zoning configuration). This is the method required by the Fibre Channel standards FC-SW-2, FC-GS-3, and FC-GS-4. If the Fibre Channel host bus adapters (HBAs) are standards-compliant, this method is reliable. The standard requires that HBAs purge their address tables and query the Name Server upon receipt of a Fibre Channel registered state change notification (RSCN). This ensures that all entries in the HBA address table correspond to members of the same zone of which the HBA is a member.

With soft zoning, an initiator that is a member of a zone has access to the names and addresses of other members of the zone only, and so cannot address frames to devices outside of the zone.

XPath OS soft zoning is industry standards-based. It is compatible with hard zoning, with Fibre Channel switch zoning implementations, and with WWN and *domain,port* member types.

You can specify soft zones with the `configureZoning` command.

Hard zoning

Hard zoning enforcement must be used in those situations in which the HBAs are not standards-compliant, for example, when your hardware predates the standards. Since hard zoning is not standards-based, different implementations are possible, but hard zoning is frequently implemented by switch ports maintaining lists of allowed source and destination ID pairs. If a frame S_ID/D_ID pair is not on a port's list of allowed pairs, the frame is not transmitted.

In XPath OS hard zoning, the frame source and destination addresses are compared to allowed addresses at the MP Router ingress F_Port or FL_Port (for devices directly attached to the MP Router) or at the egress F_Port or FL_Port (for devices directly attached to a Fabric OS-based switch). The details of hard zoning enforcement are not needed to configure zones or for routine administration of the MP Router; however, see Appendix A, "[Hard zoning background](#)," if you would like to learn more about hard zoning enforcement.

The MP Router hard zoning implementation is compatible with other HP StorageWorks switches.

Zone server compatibility

The zone server maintains the zone configuration information used for both soft and hard zoning, so a fabric can simultaneously use both hard and soft zoning. The zone server distributes the soft zoning information to the Name Server and the allowed S_ID/D_ID pairs to the appropriate ports for hard zoning. XPath OS hard zoning can coexist with other implementations of hard zoning, although XPath OS hard zoning cannot necessarily share zoning information with other hard zoning implementations.

The MP Router zone server implementation complies with Fabric OS 2.6.x, 3.x, and 4.x. In multiswitch configurations, the MP Router also complies with the FC-SW2 zone server specification. The maximum number of members allowed in the zone server depends on the size of the zone database; the effective zone database can have a maximum of 4096 zone members and 1024 zones (including fabric assisted zones). The maximum zoning database size is 128 KB.

Standards and zoning compatibility

In a multiswitch fabric configuration with HP switches, a change to a zone definition is not reflected in other switches until the configuration that contains that zone is enabled. This feature complies with FC-SW2 standards.

Although software enforcement is standards-based and hardware enforcement is proprietary, there is no standards-based method of exchanging zone information between conforming and nonconforming methods. For example, HP Fabric OS-based switches provide hardware enforcement by maintaining, at the target egress port, a list of switch ports permitted to communicate with the target.

Configuring zones

The following procedure outlines basic zone configuration and includes an example of configuring zones on multiple fabrics for an LSAN zone. [Table 9](#) on page 89 describes zoning commands.

The MP Router can be included in zones using either *domain,port* or WWN naming. Fabric OS switches and XPath OS MP Routers enforce hard zoning for both WWN and *domain,port* zone object naming. However, if a zone configuration on one switch uses WWN naming and an identical zone configuration on another switch uses *domain,port* naming, the configurations do not merge, and the fabric segments. If you use *domain,port* naming, HP recommends that you use the `configure` command to set the domain ID so that it persists across reboots, power cycles, failovers, and fabric reconfiguration.

The zone configuration is stored in the configuration file. This enables you to upload and download the zoning configuration as a text file and to export a zone configuration to another switch easily.

 **NOTE:** Zones used with an MP Router have additional member-naming requirements. Use port WWNs for Fibre Channel attached nodes with MP Router LSAN zones. Each zone object defined—alias, zone, zone configuration—must have a unique name. Include a space after each comma in zoning commands that contain commas.

Implementing zoning

1. Review the existing zoning configurations using the `cfgShow` command.
2. Create an alias:

```
router:admin> alicreate "host1", "1,6"  
Alias Create Successful
```

3. Create zones:

```
router:admin> zonecreate "zone1", "host1"  
Zone Create Successful  
  
router:admin> zonecreate "zone2", "1,7"  
Zone Create Successful
```

4. Create a zone configuration:

```
router:admin> cfgcreate "backupcfg", "zone1;zone2"  
Cfg Create Successful
```

5. Check that the zone configuration is correct:

```
router:admin> zoneshow
Defined configurations:
  Cfg: backupcfg      zone1;zone2
  Zone:               zone1
                    host1
  Zone:               zone2
                    1,7
  Alias:              host1
                    1,6
Effective configuration:
  No configuration enabled
```

6. Enable the zone configuration:

```
router:admin> cfgenable "backupcfg"
Cfg Enable Successful
```

7. Save the zone configuration:

```
router:admin> cfgsave
Cfg save Successful
```

8. Review the enabled configuration to verify that the enabled configuration is correct:

```
router:admin> cfgshow
Defined configurations:
  Cfg: backupcfg      zone1;zone2
  Zone:               zone1
                    host1
  Zone:               zone2
                    1,7
  Alias:              host1
                    1,6

Effective configuration:
  cfg: backupcfg
  Zone:               zone1
                    1,6
  Zone:               zone2
                    1,7
router:admin>
```

Mapping iSCSI names

The zone-related commands and alias-related commands convert iSCSI names (which are preceded with the characters `iqn.`) into WWNs from the preallocated WWN pool.

Specific zone-related commands are `zoneCreate`, `zoneAdd`, and `zoneRemove`.

Specific alias-related commands are `aliCreate`, `aliAdd`, and `aliRemove`.

Zone information with an iSCSI name can be retrieved with the `aliShow`, `zoneShow`, and `cfgShow` commands using the `-i` option.

For more information about iSCSI name mapping with WWN, see ["Working with the WWN mapping table"](#) on page 82.

Implementing an iSCSI name in zoning

1. Create an alias:

```
router:admin> alicreate "ALIAS_2", "iqn.2001-04.com.example:arraysa86"  
Alias Create Successful
```

2. Create a zone:

```
router:admin> zonecreate "ZONE_C", "iqn.2002-04.com.example:arraysa86"  
Zone Create Successful
```

3. Check the zone configuration:

```
router:admin> cfgshow -i  
Defined configurations:  
Cfg:    cfg_iscsi          ZONE_C  
zone:   ZONE_C  
        iqn.2002-04.com.example:arraysa86  
alias:  ALIAS_2  
        iqn.2001-04.com.example:arraysa86  
Effective configuration:  
cfg:    cfg_iscsi  
Zone:   ZONE_C  
        iqn.2002-04.com.example:arraysa86
```

Zoning commands

Use the commands in [Table 9](#) to configure and manage your zone configuration. See the *HP StorageWorks XPath OS 7.4.x command reference guide* for detailed command information.

Table 9 Zoning commands

Command	Description
Zone management commands	
configureZoning	Configures zoning
configZoningShow	Displays zone configurations
Alias commands	
aliAdd	Adds a member to a zone alias
aliCreate	Creates a zone alias
aliDelete	Deletes a zone alias
aliRemove	Removes a member from a zone alias
aliShow	Displays alias information
Zone commands	
zoneAdd	Adds a member to a zone
zoneCreate	Creates a zone
zoneDelete	Deletes a zone
zoneRemove	Removes a member from a zone.
zoneShow	Displays zone information
Configuration commands	

Table 9 Zoning commands (continued)

Command	Description
cfgAdd	Adds a zone to a zone configuration
cfgActvShow	Prints the effective zone configuration
cfgClear	Clears all zone configurations
cfgCreate	Creates a zone configuration
cfgDelete	Deletes a zone configuration
cfgDisable	Disables a zone configuration
cfgEnable	Enables a zone configuration
cfgRemove	Removes a zone from a zone configuration
cfgSave	Saves zone configurations in flash memory
cfgShow	Displays zone configurations in flash memory
cfgSize	Displays the size of the configurations in flash memory
Qloop commands	
qloopAdd	Adds a member to a Qloop
qloopCreate	Creates a Qloop
qloopDelete	Deletes a Qloop
qloopRemove	Removes a member from a Qloop
qloopShow	Prints Qloop information
Fabric assist zone commands	
fazoneAdd	Adds a member to a fabric assist zone
fazoneCreate	Creates a fabric assist zone
fazoneDelete	Deletes a fabric assist zone
fazoneRemove	Removes a member from a fabric assist zone
fazoneShow	Prints fabric assist zone information

8 Using ISL trunking

This chapter provides information on HP Interswitch Link (ISL) trunking, and consists of the following sections:

- [How exchange-based trunking works](#), next
- [Enabling trunking](#), page 92
- [Managing trunking](#), page 92
- [Trunking commands](#), page 93

How exchange-based trunking works

The MP Router exchange-based trunking feature increases overall bandwidth by distributing network traffic across ISLs connecting pairs of switches. More precisely, load balancing occurs across all equal-cost paths to a destination domain.

This functionality works in two situations:

- Between two MP Routers
- From an MP Router to any other Fabric OS-based switch

Exchange-based trunking also works across multiple EX_Ports when MP Router ports are configured as Fibre Channel routers.

Exchange-based trunking is supported at both 1 Gbit/sec and 2 Gbit/sec link speeds. There is no restriction on the port location/numbers that can be part of the trunk.

Exchange-based trunking is activated when you install the appropriate HP licensed feature bundle. To create a trunk, use the `trunkSet` command.

 **NOTE:** MP Router exchange-based trunking is not supported with Continuous Access EVA. For all Continuous Access EVA configurations using FCIP, you must set trunking to `SID_DID` or `none`. See the `portcfgfcip` command.

 **NOTE:** Fibre Channel routing link-cost calculation uses a simple formula based only on link speed.

When exchange-based trunking is enabled, all equal-cost paths are used for sending traffic to the given destination domain. Each ingress port contains a list of egress ports that can be used to reach a destination domain. On receiving a frame at the ingress port, the following fields in the Fibre Channel frame header are used to determine which egress port to use:

- Source ID (S_ID)
- Destination ID (D_ID)
- Originator Exchange ID (OX_ID)

This scheme guarantees in-order delivery within a given exchange.

Exchange-based trunking does not depend on having exchange-based trunking operating on the destination switch; the routing decision for any given frame of data is made at the origin switch. Thus, the choice of which ISL to take always occurs when an MP Router is the source switch and any other switch is the destination switch. If the other switch is not an MP Router, the return traffic does not benefit from exchange-based trunking load balancing.

MP Routers are compatible with other HP StorageWorks switches. However, they are not compatible with *trunk groups* (groups of four ports) of HP ISL trunking in Fabric OS 4.2 and earlier. When an MP Router has multiple ISLs to a different model HP StorageWorks switch, traffic from the MP Router to the other model is load balanced using exchange-based trunking. Traffic from the other model to the MP Router is load balanced at a granularity of initiator device, target device pair. That is, rather than at a

data-frame-by-data-frame level, load is instead rebalanced each time an initiator or target device's connection or disconnection impacts the routing behavior of a given ISL path.

All ingress ports share all available routes to a destination domain. There is a maximum of 16 routes to a destination domain.

In the fabric illustrated in [Figure 19](#), Domain 100 egress ports 1, 2, 3, and 4 can be used to reach destination Domain 200. All these egress ports have an equal cost of 1000 (routes through Domain 100 ports 1, 2, and 3 are direct 1-Gbit/sec links to Domain 200; the route through Domain 100, Port 4 to Domain 300 and then to Domain 200 has two 2-Gbit/sec links). The host port (Domain 100, port 6) has all these egress ports available for routing frames to Domain 200, Port 11.

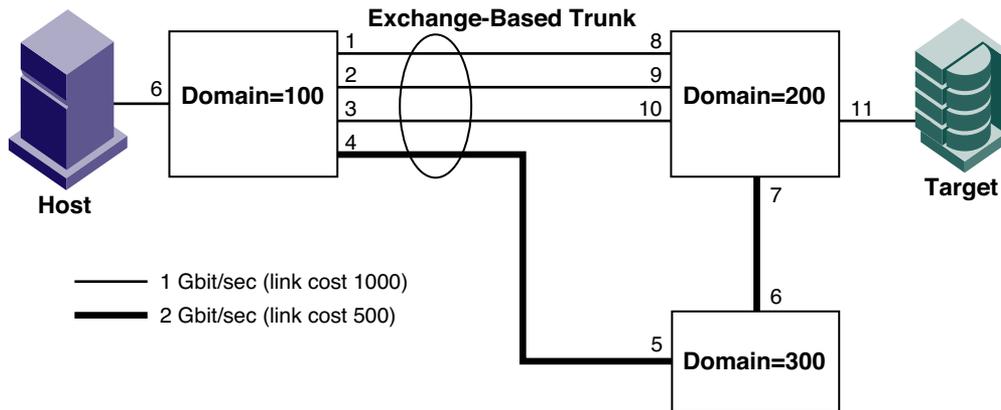


Figure 19 Sample exchange-based trunking configuration

The following are characteristics of exchange-based trunking:

- Any E_Port can be part of the trunk (and can be connected to any port).
- No interaction between the switches is required to enable the feature.
- Any mixture of link speed E_Ports is supported; all links must have an equal cost to the destination domain.
- Load balancing distribution might not be exactly equal; in a four-port trunk, for example, the load sharing might not be exactly 25% each.
- Load balancing is ineffective if the exchange ID does not vary (for example, same OX_ID usage). However, this is not a problem with typical FC-4 protocols like FCP, IP over Fibre Channel, and so on.

Enabling trunking

Exchange-based trunking is activated when the appropriate license bundle is installed. You can activate it by entering a license key, available from your switch supplier, and using the `trunkSet` command. When exchange-based trunking is activated, trunking is automatically implemented for any eligible ISLs. Trunking is enabled or disabled separately on each switch. You are not required to have exchange-based trunking enabled or disabled uniformly across a fabric or on multiple FCRs.

You can verify and activate licenses through the command line interface (CLI) or through Advanced Web Tools. For instructions on activating a license, see [“Licensed features”](#) on page 24 or the *HP StorageWorks XPath OS 7.4.x Advanced Web Tools administrator guide*.

ISL trunking is supported in HP fabrics, but not in McDATA or mixed fabrics.

Managing trunking

This section provides information on managing exchange-based trunking on a switch.

NOTE: The trunking feature is disabled by default and can be enabled only after the trunking license is installed.

Determining whether the trunking feature is enabled on your switch

1. Log in as admin.
2. Issue the `trunkShow` command:

```
router:admin> trunkshow  
Trunking is disabled
```

Enabling trunking on a switch

1. Log in as admin.
2. Issue the `trunkSet` command:

```
router:admin> trunkset  
Trunk feature enabled
```

All Fibre Channel ports that passed POST are enabled for trunking. If the switch was part of a fabric, the fabric reconfigures.

Disabling trunking on a switch

1. Log in as admin.
2. Issue the `trunkReset` command:

```
router:admin> trunkreset  
Trunk feature disabled
```

All trunking on the switch is disabled. If the switch was part of a fabric, the fabric reconfigures.

Checking for traffic distribution from ingress port to destination domain

1. Connect fabrics with multiple paths from one switch to another.
2. Connect hosts on one switch to targets on another switch.
3. Issue the `topologyShow` command and determine the total number of equal-cost ISLs from the ingress port to the remote domain.
4. Check routes assigned to an ingress port to reach a destination domain, using the `urouteShow` command.
The number of routes from an ingress port to a remote domain should be equal to the total number of equal-cost ISLs to that domain, as determined from the `topologyShow` command output in [step 3](#).
5. Run traffic tests from the hosts to the targets.
6. Review the load-balancing configuration using the `portPerfShow` command.

Trunking commands

[Table 10](#) lists commands used to configure and manage exchange-based trunking. For detailed information on these commands, see the *HP StorageWorks XPath OS 7.4.x command reference guide*.

Table 10 Trunking commands

Command	Description
<code>iodReset</code>	Disables in-order delivery
<code>iodSet</code>	Enables in-order delivery
<code>iodShow</code>	Displays current in-order delivery setting
<code>portCfgSpeed</code>	Displays or sets the configured port speed
<code>portPerfShow</code>	Displays the performance of all ports on a switch at selected intervals

Table 10 Trunking commands (continued)

Command	Description
<code>topologyShow</code>	Displays the fabric topology, as seen by the local switch
<code>trunkReset</code>	Disables trunking on a switch
<code>trunkSet</code>	Enables trunking on a switch
<code>trunkShow</code>	Displays whether trunking is enabled or disabled on the switch
<code>urouteShow</code>	Displays the unicast routing information for a port

9 Monitoring system logs

This chapter discusses the following topics:

- [System error log](#), next
- [Port log](#), page 97
- [Using the syslog daemon](#), page 98

There are three log file systems in the XPath OS:

- The *system error log* displays system daemon errors, in addition to all events from the event log.
- The *port log* displays port information.
- The *event log* displays events only.

Log entries for all three logs are described in the *HP StorageWorks XPath OS 7.4.x system error messages reference guide*.

For most MP Router troubleshooting tasks, you can use the `errShow` command to refer to the system error log. The system error log is formatted to match the Fabric OS error message format; it provides a complete set of system messages. This chapter includes general information on the format and severity of error log entries and provides procedures for viewing and clearing the event log.

The port log displays port events, such as port down, link events, and fabric segmentation. This chapter provides information on viewing the port log.

The event log provides a limited set of event messages, but does not display errors from the firmware subsystems.

A limited amount of storage is shared by the three logs on the MP Router. You can configure the XPath OS syslog daemon to copy log entries to a server.

See the *HP StorageWorks XPath OS 7.4.x system error messages reference guide* for information on the meaning of specific entries and information on using the log.

System error log

View this log using the `errShow` command. The error log displays system errors, such as faulty or failing daemon processes in the firmware, as well all the events listed in the event log. [Table 11](#) describes the commands that manage the system error log.

System error messages are saved in the `/var/log/messages` directory. When the message file reaches 512 KB, the XPath OS archives the log with the file name `messages.0.gz`, and creates a new system error log. There can be a total of five archived files on the system, named `messages.0.gz` through `messages.4.gz`. As each archive file is added, XPath OS rearranges the files so that `messages.0.gz` is the newest archive and `messages.4.gz` is the oldest. After the maximum of five archive files is reached, the oldest archive file is replaced with the next oldest. The `errShow` command displays the contents of the current system error log only; it does not display archived messages.

Table 11 Managing the system error log

Command	Description
<code>errShow</code>	Displays the contents of the system error log.
<code>errClear</code>	Clears the contents of the system error log.

Message severity levels

There are six severity levels for messages, ranging from Panic to Debug. The definitions in [Table 12](#) can be used as general guidelines for troubleshooting. See the *HP StorageWorks XPath OS 7.4.x system error messages reference guide* for complete error message descriptions; review the descriptions thoroughly before taking action.

Table 12 Message severity levels

Event level	Description
0 = Panic	Panic messages indicate that a specific software subsystem has detected a fatal or unrecoverable error condition. Examples are memory allocation failure, system call failure, and software detection of problems with the ASIC or with hardware subsystems. These errors usually indicate partial or complete failure of a subsystem.
1 = Critical	Critical messages indicate that the software has detected serious problems that eventually cause a partial or complete failure of a subsystem if not corrected immediately. A power supply failure, for example, or a rise in temperature must receive immediate attention. Some of the critical errors might overlap in severity with the panic messages.
2 = Error	Error messages indicate error conditions that do not significantly affect overall system functionality. For example, error messages might indicate time-outs on certain operations, failures of certain operations after retries, invalid parameters, or failures to perform requested operations.
3 = Warning	Warning messages highlight current operating conditions that should be checked before they cause failures. For example, a power supply failure in a redundant system relays a warning that the system is no longer operating in redundant mode, and that the failed power supply should be replaced or fixed.
4 = Info	Info messages report the current status of the system components other than error status. For example, detecting on and off line status of a fabric port.
5 = Debug	Debug messages are for debugging use only. They are produced by code inserted by the vendor to inform the user that a suspected problem has occurred.

Viewing the system error log

The system error log is a collection of all daemon-related errors (such as memory allocation errors), as well as all events from the event log.

1. Log in as user or admin.
2. Issue the `errshow` command:

```
router:admin> errshow
```

The system error log content is displayed, with page breaks.

Alternatively, you can use the `-a` option to display the system error log content without page breaks:

```
router:admin> errshow -a
```

Sample system error log message

The following is a sample message from the system error log:

```
router:admin> errshow

Error 239
-----
301 (EvtMgr): Jan 26 13:10:34
Info CM-EVENT_USER_LOGIN_SUCCESS, 4, User login: admin
```

The fields in the event message are described in [Table 13](#).

Table 13 System error log message field descriptions

Example	Variable name	Description
Error 239	Error log buffer number	Displays a rotating number that describes the position the message holds in the buffer. This number is not permanently associated with the error itself and should not be used when contacting your service provider; provide the error code name instead.
301 (EvtMgr)	Reporting process ID	Displays the process ID and name of the module reporting the error.
Jan 26 13:10:34	Date and time stamp	Displays the date and time the error message occurred.
Info	Severity level	Displays the severity of the message: Panic, Critical, Error, Warning, Info, or Debug.
CM-EVENT_USER_LOGIN_SUCCESS	Error code name	Displays the code name for the error.
4	Severity level	Displays the severity of the error, in a numbered format: 0 = Panic 1 = Critical 2 = Error 3 = Warning 4 = Info 5 = Debug
User login: admin	Error description	Displays error-specific data, such as the error reason.

Clearing the system error log

1. Log in as admin.
2. Issue the `errClear` command:

```
router:admin> errclear  
The error log is cleared.
```

Port log

View this log using the `portLogShow` command. The port log displays port events, such as port down, link events, and fabric segmentations. Events from this log are for information only.

The XPath OS maintains an internal port log of all port activity. The port log stores entries for each port as a circular buffer. Each port has space to store 2048 log entries. If the port log is disabled, an error message is displayed. Once the log is full, the newest log entries overwrite the oldest log entries. Port logs are not persistent and are lost over power-cycles and reboots.

Port log management

Use the commands in [Table 14](#) to view and manage port logs.

Table 14 Port log management commands

Command	Description
portLogClear	Clears port logs for all or specified ports
portLogDisable	Disables port logs for all or specified ports
portLogDump	Displays port logs for all or specified ports, without page breaks
portLogEnable	Enables port logs for all or specified ports
portLogShow	Displays port logs for all or specified ports, with page breaks

See the *HP StorageWorks XPath OS 7.4.x command reference guide* for detailed information on these commands.

Sample port log

Following is a sample usage of the `portLogShow` command:

```
router:admin> portlogshow 8
Total records present      = 12
Number of records displayed = 12

Time          Module  Event   Port  Len  Log info
-----
18:36:52.036 fabctl  PrtSCN  08    0    st=1, Topo=2, Spd=0
18:36:52.361 WKA     Rx       08   140  22ffffffe,00000000,01a6ffff,04000000
18:36:52.362 fabctl  PrtSCN  08    0    st=2, Topo=2, Spd=2
18:36:52.365 fabctl  Debug   08    0    Loading routes
18:36:52.379 fabctl  Tx       08   140  23640800,00ffffffe,01a60001,02000000
18:36:52.379 WKA     Rx       08   140  22ffffffc,00640800,02ceffff,03000000
18:36:52.382 nsd     Tx       08   140  23640800,00ffffffc,02ceffff,02000000
18:36:52.382 WKA     Rx       08    32  22ffffffd,00640800,02cdffff,62000000
18:36:52.383 fabctl  Tx       08    28  23640800,00ffffffd,02cd0001,02000000
18:36:52.383 WKA     Ct_in   08    52  02ffffffc,00640800,02d1ffff,01000000
18:36:52.384 nsd     Tx       08    40  03640800,00ffffffc,02d1ffff,01000000
18:36:52.384 WKA     Ct_in   08    84  02ffffffc,00640800,02d0ffff,01000000
```

See the *HP StorageWorks XPath OS 7.4.x system error messages reference guide* for interpretation of the fields.

Using the syslog daemon

XPath OS can be configured to send error log messages to a UNIX® or Linux® host system or a host running any other operating system that supports standard syslogd functionality. This host system can be configured to receive error messages from the MP Router and store them in files on the computer hard drive, which overcomes the size limitations of the internal log buffers.

The syslogd process reads and logs messages to the system console, log files, other machines, and users, as specified by its configuration file. See the manual pages and related documentation for your particular UNIX host system for more information on the syslogd process and its capabilities.

XPath OS syslogd CLI commands

Table 15 shows the commands related to the syslogd configuration. See the *HP StorageWorks XPath OS 7.4.x command reference guide* for more details.

Table 15 syslogd configuration commands

Command	Purpose
syslogdIpAdd	Adds the IP address of the remote syslogd host to the MP Router
syslogdIpRemove	Removes the IP address of the remote syslogd host from the MP Router
syslogdIpShow	Shows the list of configured syslogd IP addresses on the MP Router
eventShow	Displays messages from the event log on the MP Router
errShow	Displays messages from the system error log on the MP Router
errClear	Clears messages from the system error log on the MP Router

Enabling syslogd

This procedure explains how to configure the MP Router to dispatch error log messages to a remote syslogd host:

1. Log in as admin.
2. Issue the `syslogdIpAdd` command with the following syntax:
`syslogdipadd IP_address_of_the_remote_syslogd_host`
3. Verify that the IP address was entered correctly by issuing the `syslogdIpShow` command.

The following example shows how to configure the MP Router to dispatch error log messages to a remote syslogd host at IP address 10.10.10.1 and then verify that configuration:

```
router:admin> syslogdipadd 10.10.10.1  
  
syslog.IP.address 10.10.10.1 is added  
router:admin> syslogdipshow  
syslog.IP.address.1 10.10.10.1
```

Disabling syslogd

To disable logging of error messages to a previously enabled remote syslogd host:

1. Log in as admin.
2. Issue the `syslogdIpRemove` command with the following syntax:
`syslogdipremove IP_address_of_the_remote_syslogd_host`
3. Verify that the IP address was deleted by issuing the `syslogdIpShow` command.

The following example shows how to disable sending of error log messages to a previously configured remote syslogd host whose IP address is 10.10.10.1.

```
router:admin> syslogdipremove 10.10.10.1  
syslog.IP.address 10.10.10.1 is removed
```


A Hard zoning background

In XPath OS hard zoning, the frame source and destination addresses are compared to permitted addresses at the MP Router ingress F_Port or FL_Port (for devices directly attached to the MP Router) or at the egress F_Port or FL_Port (for devices directly attached to a non-HP switch). You do not need to know the details of hard zoning to configure zones or for routine administration of the MP Router.

Each MP Router maintains a zone server. The zone server distributes to the ports allowed combinations of source address/destination address pairs. Hard zoning supports both WWN and *domain, port* member types.

For hosts and targets directly connected to an MP Router, the source/destination pair for each frame entering the MP Router through an F_Port or FL_Port is checked for zoning compliance. (See [Figure 20](#) for an picture of this configuration.) The frame is discarded if the source/destination communication is not allowed by the zone rules.

For frames originating from the host, the source/destination combination is checked against the zone membership list at MP Router 1, Port M. For frames originating at the target, the source/destination combination is checked against the zone membership list at MP Router 1, Port N.

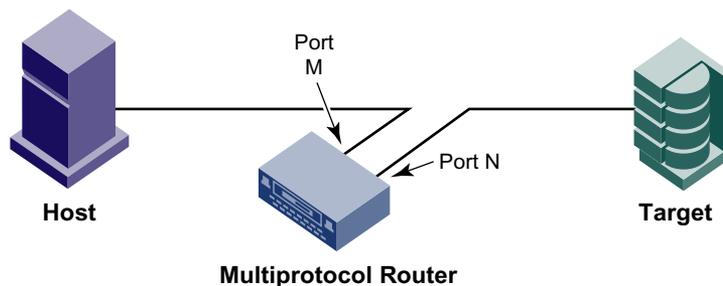


Figure 20 Host and target directly connected to the MP Router

For hosts and targets directly connected to a homogeneous fabric with multiple MP Routers, the source/destination pair for each frame entering the MP Router through an F_Port or FL_Port is checked for zoning compliance. The frame is discarded if source/destination communication is not allowed by the zone rules. [Figure 21](#) illustrates this combination.

For frames originating from the host, the source/destination combination is checked against the zone membership list at MP Router 1, Port W. For frames originating at the target, the source/destination combination is checked against the zone membership list at MP Router 2, Port Z.

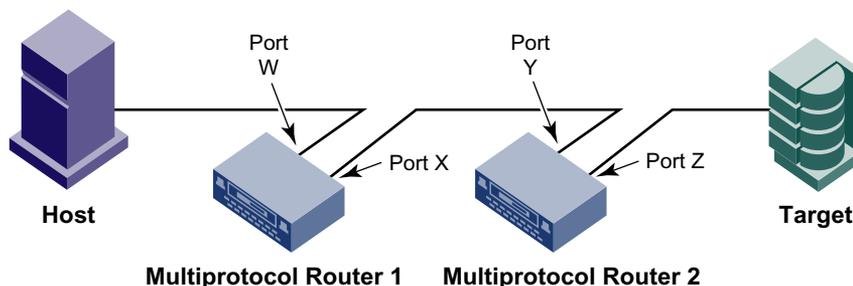


Figure 21 Host and target in a homogeneous fabric

For frames entering an MP Router through an E_Port (Port M in [Figure 22](#)), the source domain is checked at the E_Port; if the frame source domain is not an MP Router, the frame is tagged and passed through to the proper egress port, Port N in [Figure 22](#).

If the originating source domain is an MP Router, the zoning check has already been performed; no further zone enforcement is required. This is the situation shown in [Figure 21](#). At the egress port, the S_ID/D_ID pair is checked.

If the destination ID is directly attached to an MP Router port (as in [Figure 22](#)), the S_ID/D_ID pair is checked against the list of allowed combinations and is dropped if the pair is not permitted.

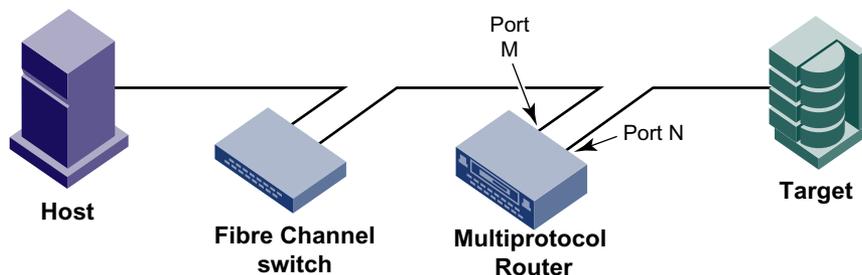


Figure 22 Host and target in a heterogeneous fabric, combination 1

If the destination ID is not directly attached to an MP Router port (as in [Figure 23](#)), the combination is not checked.

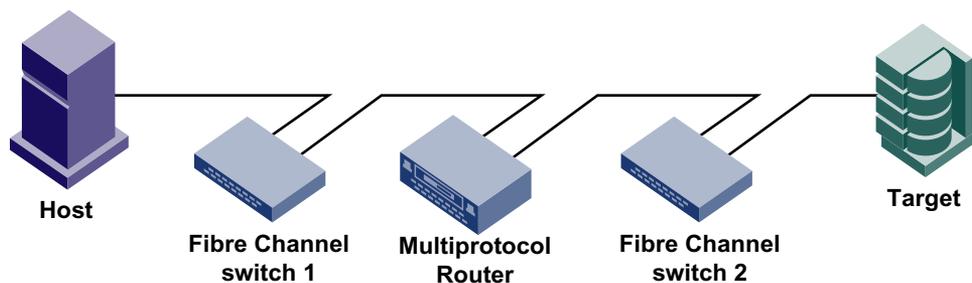


Figure 23 Host and target in a heterogeneous fabric, combination 2

B Recovery kernel for XPath OS 7.4.x

The recovery kernel (RK) is a part of the XPath OS that is stored in a reserved portion of memory, protected from erasure. This kernel provides the mechanics of the boot process: It allows the hardware to locate and boot the proper bank. The kernel finds either bank1 or bank2 as bootable. This bank can then be used to recover firmware.

The NVRAM value, `cfgBank`, dictates the bank from which to boot.

When entering NVRAM values, use lowercase only. Mixed case is shown in this document for readability only.

Software installation support environment

Software installation is supported by an FTP/TFTP server that is accessible to (that is, it resides on the same network as) the MP Router. The server must be configured to support both FTP and TFTP services: TFTP is used to download the recovery kernel and FTP is used to download the base firmware and other packages.

Most commonly, within the FTP/TFTP server the `/tftpboot` directory is used to store the uploaded kernel, as well as the base firmware and third-party application packages.

Unlike these packages, which are provided in an RPM format, the recovery kernel is a unique, proprietary file format.

Using the recovery kernel

The recovery kernel may be installed using the low level process monitor (PMON). Once installed, it can be used to install the base router firmware by issuing the `firmwareDownload` command.

Installing the recovery kernel using the process monitor

1. Reboot the router.
2. During the boot-up process, the MP Router displays a user `boot delay` prompt, showing the message:

```
Press Return to boot or any key to enter PMON
```

3. Press the space bar to cause the MP Router enter PMON.
4. Enter the following commands, which are necessary to define the correct recovery kernel version and boot server IP address:

```
router:PMON> set bankover 1.5.x.y
router:PMON> set rootdir /
router:PMON> set basefile XPathRecoverAP7420
router:PMON> set bootserver xxx.xxx.xxx.xxx
router:PMON> set currentdnldproto tftp
router:PMON> set cfgbank bank0
router:PMON> netload -p 0xdeeddeed;g
```

The boot server IP address is specific to a local site; your network administrator can provide this information. The recovery kernel version is provided with each software release.

When you issue the `netLoad` command, the system begins downloading the recovery kernel, writes it into NVRAM, and then reboots the router. After the reboot completes, the system presents a recovery kernel prompt:

```
Recovery Kernel%
```

5. After you are presented with the recovery kernel prompt, install the base package firmware as described in [“Installing the XPath OS 7.4.0 base firmware from the recovery kernel.”](#)

Installing the XPath OS 7.4.0 base firmware from the recovery kernel

1. Clear the storage of old files by entering:

```
format all
```

at the recovery kernel prompt and then pressing **Enter** at the recovery kernel prompt.

2. After you are returned to a prompt, issue the `firmwareDownload` command with the necessary parameters.

In the following example, the boot server is located at the IP address `192.168.25.100`, the FTP-access user name is `user`, and the file `xpath_os_v7.4.0` is stored in the `/tftpboot` directory.

```
recovery kernel% firmwaredownload 192.168.25.100 user  
/tftpboot/xpath_os_v7.4.x
```

3. When prompted, enter the password for the user name you entered in the previous step.
See "[Upgrading the XPath base and recovery kernel:](#)" on page 37 for details, or see the *HP StorageWorks XPath OS 7.4.x command reference guide* for details about the `firmwareDownload` command.
4. After the software packages are installed, you must reboot the MP Router.
The system is rebooted from bank1 and the banks are synchronized before the system returns the login prompt.

Glossary

AL_PA

Arbitrated-loop physical address. A unique 8-bit value assigned during loop initialization to a port in an arbitrated loop. Also called *arbitrated-loop parameters*.

alias

A logical grouping of elements in a fabric. An alias is a collection of port numbers and connected devices used to simplify the entry of port numbers and WWNs when creating zones.

alias server

A fabric software facility that supports multicast group management.

arbitrated loop

A shared 100-Mb/sec Fibre Channel transport structured as a loop. Can support up to 126 devices and one fabric attachment. See also [topology](#).

arbitration

A method of gaining orderly access to a shared-loop topology.

ARP

Address Resolution Protocol. A TCP/IP function for associating an IP address with a link-level address.

backbone fabric

An optional capability that enables scalable meta-SANs by allowing the networking of multiple FC routers, which connect to the backbone fabric via EB_Port interfaces.

BB_Credit

Buffer-to-buffer credit. The number of frames that can be transmitted to a directly connected recipient or within an arbitrated loop. Determined by the number of receive buffers available. See also [buffer-to-buffer flow control](#).

BB fabric

A backbone fabric that connects FC routers. The FC routers communicate over the backbone fabric using Fibre Channel Router Protocol (FCRP).

buffer-to-buffer flow control

Management of the frame transmission rate in either a point-to-point topology or in an arbitrated loop. See also [buffer-to-buffer flow control](#).

configuration

1. A set of parameters that can be modified to fine-tune the operation of a switch. Issue the `configShow` command to view the current configuration of your switch.
2. In HP StorageWorks zoning, a zoning element that contains a set of zones. The configuration is the highest-level zoning element and is used to enable or disable a set of zones on the fabric. See also [zone configuration](#).

congestion

The realization of the potential of oversubscription. A congested link is one on which multiple devices are contending for bandwidth.

core PID

Core switch port identifier. The core PID must be set for Fabric OS 3.1 and earlier switches included in a fabric of Fabric OS 4.1 switches. This parameter is located in the `configure` command of Fabric OS 3.1 and earlier. All Fabric OS 4.1 switches and later use the core PID format by default; this parameter is not present in the `configure` command for these switches.

credit

In Fibre Channel technology, the number of receive buffers available to transmit frames between ports. See also [buffer-to-buffer flow control](#).

D_ID

Destination identifier. A three-byte field in the frame header that indicates the address identifier of the N_Port to which the frame is headed.

defined zone configuration

The set of all zone objects defined in the fabric. Can include multiple zone configurations. See also [enabled zone configuration](#), [zone configuration](#).

DLS

Dynamic load-sharing. Dynamic distribution of traffic over available paths. Allows for recomputing of routes when an Fx_Port or E_Port changes status. See also [Fx_Port](#), [E_Port](#).

domain ID

A unique identifier for all switches in a fabric, used in routing frames. Usually assigned by the principal switch, but can be assigned manually. The domain ID for an HP StorageWorks switch can be any integer from 1 through 239.

E_D_TOV

Error-detect timeout value. The minimum time a target waits for a sequence to complete before initiating recovery. Can also be defined as the maximum time allowed for a round-trip transmission before an error is declared. See also [R_A_TOV](#).

E_Port

Expansion port. A standard Fibre Channel mechanism that enables switches to network with each other, creating an ISL. See also [ISL](#).

edge fabric

A Fibre Channel fabric connected to an FC router through an EX_Port (where hosts and storage are attached in a meta-SAN). See also [EX_Port](#).

enabled zone configuration

The currently enabled configuration of zones. Only one configuration can be enabled at a time. See also [defined zone configuration](#), [zone configuration](#).

EX_Port

A type of E_Port that connects an FC router to an edge fabric. EX_Ports limit the scope of fabric services, but provide device connectivity using FC-NAT.

exchange

The highest-level Fibre Channel mechanism used for communication between N_Ports. Composed of one or more related sequences, it can work in either one or both directions. See also [N_Port](#).

exported device

A device that has been mapped between fabrics (a host or storage port in one edge fabric can be exported to any other fabric by using LSAN zoning).

F_Port

Fabric port. A port that is able to transmit under fabric protocol and interface over links. It can be used to connect an N_Port to a switch. See also [FL_Port](#), [Fx_Port](#), [N_Port](#).

fabric

A collection of Fibre Channel switches and devices, such as hosts and storage. Also called a *switched fabric*. See also [SAN](#), [topology](#).

fabric name

The unique identifier assigned to a fabric and communicated during login and port discovery.

fabric port count

The number of ports available for connection by nodes in a fabric.

fabric services

Codes that describe the communication to and from any well-known address.

fabric topology

The arrangement of switches that form a fabric.

FAN

Fabric address notification. Retains the AL_PA and fabric address when a loop reinitializes, if the switch supports FAN. See also [AL_PA](#).

FC router

A platform running the HP StorageWorks Fibre Channel Routing Service or FC-to-FC routing (for instance, the MP Router) that enables two or more fabrics to share resources (such as hosts or storage devices) without merging those fabrics. The platform could simultaneously be used as an FC router and as an FCIP tunnel.

FC-0

Lowest layer of Fibre Channel transport. Represents physical media.

FC-2

Layer of Fibre Channel transport that handles framing and protocol, frame format, sequence/exchange management, and ordered set usage.

FC-3

Layer of Fibre Channel transport that contains common services used by multiple N_Ports in a node.

FC-4

Layer of Fibre Channel transport that handles standards and profiles for mapping upper-level protocols such as SCSI and IP onto the Fibre Channel Protocol.

FCC

Federal Communications Commission.

FC-GS

Fibre Channel generic services.

FC-GS-2

Fibre Channel generic services, second generation.

FC-GS-3

Fibre Channel generic services, third generation.

FC_IP

Fibre Channel over IP.

FC-NAT

Fibre Channel network address translation.

FC-PH

The Fibre Channel physical and signaling standard for FC-0, FC-1, and FC-2 layers of the Fibre Channel Protocol. Indicates signaling used for cable plants, media types, and transmission speeds.

FC-PH-2

Fibre Channel Physical Interface, second generation.

FC-PH-3

Fibre Channel Physical Interface, third generation.

FCIP Tunneling Service

The HP StorageWorks Multi-protocol SAN Routing Service that enables SANs to span longer distances than could be supported with native Fibre Channel links. FCIP is a TCP/IP-based tunneling protocol that allows the transparent interconnection of geographically distributed SAN islands through an IP-based network.

FCP

Fibre Channel Protocol. Mapping of protocols onto the Fibre Channel standard protocols. For example, SCSI FCP maps SCSI-3 onto Fibre Channel.

FCRP

Fibre Channel Router Protocol. A protocol that enables LSAN switches to perform routing between different edge fabrics, optionally across a backbone fabric.

FCRS

Fibre Channel Routing Service. The HP StorageWorks Multi-protocol SAN Routing Service that extends hierarchical networking capabilities to Fibre Channel fabrics. Sometimes called *FC-to-FC routing*, FCRS enables devices located on separate fabrics to communicate without merging the fabrics. It also enables the creation of LSANs.

FCS

Fibre Channel switch.

FCS switch

Refers to the HP StorageWorks Secure Fabric OS feature. One or more designated switches that store and manage security parameters and configuration data for all switches in the fabric. They also act as a set of backup switches to the primary FCS switch. See also [backbone fabric](#).

FC-SW-2

The second-generation Fibre Channel Switch Fabric standard defined by ANSI. Specifies tools and algorithms for the interconnection and initialization of Fibre Channel switches to create a multiswitch Fibre Channel fabric.

Fibre Channel

The primary protocol used for building SANs to transmit data between servers, switches, and storage devices. Unlike IP and Ethernet, Fibre Channel was designed to support the needs of storage devices of all types. It is a high-speed, serial, bidirectional, topology-independent, multi-protocol, and highly scalable interconnection between computers, peripherals, and networks.

Fibre Channel transport

A protocol service that supports communication between Fibre Channel service providers. *See also* [FSP](#).

FID

Fabric ID. Unique identifier of a fabric in a meta-SAN.

firmware

The basic operating system provided with the hardware.

FL_Port

Fabric loop port. A port that is able to transmit under fabric protocol and also has arbitrated-loop capabilities. Can be used to connect an NL_Port to a switch. *See also* [F_Port](#), [Fx_Port](#), [NL_Port](#).

flash

Programmable nonvolatile RAM (NVRAM) memory that maintains its contents without power.

frame

The Fibre Channel structure used to transmit data between ports. Consists of a start-of-frame delimiter, header, optional headers, data payload, cyclic redundancy check (CRC), and end-of-frame delimiter. There are two types of frames: link control frames (transmission acknowledgements and so forth) and data frames.

FS

Fibre Channel service. A service that is defined by Fibre Channel standards and exists at a well-known address. For example, the Simple Name Server is a Fibre Channel service. *See also* [FSP](#).

FSP

Fibre Channel Service Protocol. The common protocol for all fabric services, transparent to the fabric type or topology. *See also* [FS](#).

FSPF

Fabric shortest path first. The HP StorageWorks routing protocol for Fibre Channel switches.

FSS

Fabric OS state synchronization. The FSS service is related to high availability (HA). The primary function of FSS is to deliver state update messages from active components to their peer standby components. FSS determines whether fabric elements are synchronized, and thus FSS compliant.

FTP

File Transfer Protocol.

FTS

Fiber Transport Services.

Fx_Port

A fabric port that can operate as either an F_Port or FL_Port. *See also* [F_Port](#), [FL_Port](#).

G_Port

Generic port. A port that can operate as either an E_Port or an F_Port. A port is defined as a G_Port when it is not yet connected or has not yet assumed a specific function in the fabric. *See also* [F_Port](#), [FL_Port](#).

gateway

Hardware that connects incompatible networks by providing translation for both hardware and software. For example, an ATM gateway can be used to connect a Fibre Channel link to an ATM connection.

Gbit/sec

Gigabits per second (1,062,500,000 bits per second).

HA

High availability. A set of features in HP StorageWorks switches that is designed to provide maximum reliability and nondisruptive replacement of key hardware and software modules.

HBA

Host bus adapter. The interface card between a server or workstation bus and the Fibre Channel network.

header

A Fibre Channel frame has a header and a payload. The header contains control and addressing information associated with the frame.

hop count

The number of ISLs a frame must traverse to get from its source to its destination.

host

A computer system that provides end users with services like computation and storage access.

in-band

Transmission of management protocol over Fibre Channel.

initiator

A server or workstation on a Fibre Channel network that initiates communications with storage devices. *See also* [target](#).

interswitch link

See [ISL](#).

IOCTL

I/O control.

IOD

In-order delivery. A parameter that, when set, guarantees that frames are either delivered in order or dropped.

IP

Internet Protocol. The addressing part of TCP.

ISL

Interswitch link. A Fibre Channel link from the E_Port of one switch to the E_Port of another. See also [E_Port](#).

JBOD

Just a bunch of disks. A number of disks connected in a single chassis to one or more controllers. See also [RAID](#).

L_Port

Loop port. A node port (NL_Port) or fabric port (FL_Port) that has arbitrated-loop capabilities. An L_Port can be in either Fabric Mode or Loop Mode.

LAN

Local area network. A network in which transmissions typically take place over fewer than 5 kilometers (3.4 miles).

latency

The time required to transmit a frame. Together, latency and bandwidth define the speed and capacity of a link or system.

Link Services

A protocol for link-related actions.

login server

The unit that responds to login requests.

LSAN

Logical storage area network. An LSAN enables device and storage connectivity that spans two or more fabrics. The path between devices in an LSAN can be local to a fabric or cross one or more FC routers and one or more backbone fabrics.

LSAN zone

The mechanism by which LSANs are administered. An FC router attached to two fabrics listens for the creation of matching LSAN zones on both fabrics. If this occurs, it creates phantom domains and FC-NAT entries as appropriate, and inserts entries for them into the name servers on the fabrics. LSAN zones are compatible with all standard zoning mechanisms.

Mbit/sec

Megabits per second.

meta-SAN

The collection of all devices, switches, edge and backbone fabrics, LSANs, and FC routers that make up a physically connected but logically partitioned storage network. LSANs span between edge fabrics using FC routers. In a data network, this would simply be called *the network*. However, an additional term is required to specify the difference between a single-fabric network (SAN), a multifabric network without cross-fabric connectivity (dual-redundant fabric SAN), and a multifabric network with connectivity (meta-SAN).

metric

A relative value assigned to a route to aid in calculating the shortest path (1000 at 1 Gbit/sec, 500 at 2 Gbit/sec).

MS

Management Server. Allows a SAN management application to retrieve information and administer the fabric and interconnected elements, such as switches, servers, and storage devices. The MS is located at the Fibre Channel well-known address FFFFFAh.

MSD

Management Server daemon. Monitors the MS. Includes the Fabric Configuration Service and the Unzoned Name Server.

MSRS

Multi-protocol SAN Routing Services. An optional, licensed software bundle available on certain HP StorageWorks platforms, such as the MP Router, that includes the Fibre Channel Routing Service and the FCIP Tunneling Service.

MTBF

Mean time between failures. An expression of time, indicating the longevity of a device.

N_Port

Node port. A port on a node that can connect to a Fibre Channel port or to another N_Port in a point-to-point connection.

Name Server (SNS)

A switch service that stores names, addresses, and attributes for up to 15 minutes and provides them as required to other devices in the fabric. SNS is defined by Fibre Channel standards and exists at a well-known address. Also called a *directory service*.

NL_Port

Node loop port. A node port that has arbitrated-loop capabilities. Used to connect an equipment port to the fabric in a loop configuration through an FL_Port. See also [N_Port](#), [Nx_Port](#).

node

A Fibre Channel device that contains an N_Port or NL_Port.

node count

The number of nodes attached to a fabric.

node name

The unique identifier for a node, communicated during login and port discovery.

NR_Port

A normal E_Port used to connect an FC router to a backbone fabric.

NS

Name Server. The service provided by a fabric switch that stores names, addresses, and attributes related to Fibre Channel objects. Can cache information for up to 15 minutes. Also called a *Simple Name Server* or a *directory service*.

Nx_Port

A node port that can operate as either an N_Port or NL_Port.

originator

The Nx_Port that originated an exchange.

out-of-band

Transmission of management protocol outside of the Fibre Channel network, usually over Ethernet.

OX_ID

Originator ID or exchange ID. Refers to the exchange ID assigned by the originator port.

packet

A set of information transmitted across a network. *See also* [frame](#).

parallel

The simultaneous transmission of data bits over multiple lines.

path selection

The selection of a transmission path through the fabric. HP StorageWorks switches use the FSPF protocol. *See also* [FSPF](#).

persistent error log

Error messages of a high enough level (by default, Panic or Critical) are saved to flash memory on the switch instead of to RAM. These messages are saved over reboots and power cycles, constituting the persistent error log. Note that each CP on an HP StorageWorks Core Switch 2/64 has its own unique persistent error log.

phantom address

An AL_PA value that is assigned to a device that is not physically in the loop. Also called *phantom AL_PA*. *See also* [AL_PA](#).

phantom device

A device that is not physically in an arbitrated-loop, but is logically included through the use of a phantom address.

phantom domain

See [xlate domain](#).

PID

Port identifier. *See also* [core PID](#).

port

In an HP StorageWorks switch environment, an SFP or GBIC receptacle on a switch to which an optical cable for another device is attached.

port address

In Fibre Channel technology, the port address is defined in hexadecimal. In the HP StorageWorks Fabric OS, a port address can be defined by a domain and port number combination or by area number. In an ESCON Director, the port address is an address used to specify port connectivity parameters and to assign link addresses for attached channels and control units.

port log

A record of all activity on a switch, kept in volatile memory.

port log dump

A view of what happens on a switch, from the switch's point of view. The `portlogdump` command is used to read the port log.

port name

A user-defined alphanumeric name for a port.

port_name

The unique identifier assigned to a Fibre Channel port and communicated during login and port discovery.

POST

Power-on self-test. A series of tests run by a switch after it is turned on.

protocol

A defined method and set of standards for communication. Determines the type of error-checking, the data-compression method, how sending devices indicate an end of message, and how receiving devices indicate receipt of a message.

R_A_TOV

Resource-allocation timeout value. The maximum time a frame can be delayed in the fabric and still be delivered. *See also* [E_D_TOV](#).

RAID

Redundant array of independent disks. A collection of disk drives that appear as a single volume to the server and are fault tolerant through mirroring or parity checking. *See also* [JBOD](#).

RCS

Reliable Commit Service. Refers to HP-specific ILS command code.

route

With respect to a fabric, the communication path between two switches. Might also apply to the specific path taken by an individual frame, from source to destination. *See also* [FSPF](#).

routing

The assignment of frames to specific switch ports, according to frame destination.

RSCN

Registered state change notification. A switch function that allows notification of fabric changes to be sent from the switch to specified nodes. The fabric controller issues RSCN requests to N_Ports and NL_Ports, but only if they have registered to be notified of state changes in other N_Ports and NL_Ports. This registration is performed through the State Change Registration (SCR) Extended Link Service. An N_Port or NL_Port can issue an RSCN to the fabric controller without having completed SCR with the fabric controller.

SAN

Storage area network. A network of systems and storage devices that communicate using Fibre Channel protocols. *See also* [fabric](#).

SAN architecture

The overall design of a storage network solution, which includes one or more related fabrics, each of which has a topology.

SAN port count

The number of ports available for connection by nodes in the entire SAN.

scalability

One of the properties of a SAN; the size to which a SAN topology can grow port and switch counts with ease.

SCN

State change notification. Used for internal state change notifications, not external changes. This is the switch logging that the port is online or is an Fx_Port, not what is sent from the switch to the Nx_Ports.

SCR

State change registration. Extended Link Service (ELS) requests the fabric controller to add the N_Port or NL_Port to the list of N_Ports and NL_Ports registered to receive the Registered State Change Notification (RSCN) Extended Link Service.

SCSI

Small Computer Systems Interface. A parallel bus architecture and a protocol for transmitting large data blocks to a distance of 15 to 25 meters.

sequence

A group of related frames transmitted in the same direction between two N_Ports.

sequence initiator

The N_Port that begins a new sequence and transmits frames to another N_Port.

sequence recipient

Serializing/deserializing circuitry. A circuit that converts a serial bit stream into parallel characters and vice versa.

serial

The transmission of data bits in sequential order over a single line.

SFP

Small-form-factor pluggable. A transceiver used on 2-GB/sec switches that replaces the GBIC.

soft zone

A zone consisting of zone members that are made visible to each other through client service requests. Typically, soft zones contain zone members that are visible to devices using Name Server exposure of zone members. The fabric does not enforce a soft zone. Note that well-known addresses are implicitly included in every zone.

SSH

Secure shell. Used starting in HP StorageWorks Fabric OS 4.1 to support encrypted telnet sessions to the switch. SSH encrypts all messages, including the client sending the password at login.

switch

A fabric device providing bandwidth and high-speed routing of data via link-level addressing.

switch name

The arbitrary name assigned to a switch.

switch port

A port on a switch. Switch ports can be E_Ports, F_Ports, or FL_Ports.

syslogd

Syslog daemon. Used to forward error messages.

target

A storage device on a Fibre Channel network. *See also* [initiator](#).

TCP/IP

Transmission Control Protocol Internet Protocol. A communications protocol developed under contract from the U.S. Department of Defense to interconnect dissimilar systems.

telnet

A virtual terminal emulation used with TCP/IP. Telnet is sometimes used as a synonym for the HP Fabric OS CLI.

throughput

The rate of data flow achieved within a cable, link, or system. Usually measured in bps (bits per second or b/sec). *See also* [BB fabric](#).

Time Server

A Fibre Channel service that allows for the management of all timers.

topology

In Fibre Channel technology, the configuration of the Fibre Channel network and the resulting communication paths allowed. There are three possible topologies:

- Point to point: A direct link between two communication ports.
- Switched fabric: Multiple N_Ports linked to a switch by F_Ports.
- Arbitrated loop: Multiple NL_Ports connected in a loop.

translate domain

See [xlate domain](#). A mode in which private devices can communicate with public devices across the fabric.

trunking

In Fibre Channel technology, a feature that enables distribution of traffic over the combined bandwidth of up to four ISLs between adjacent switches, while preserving in-order delivery.

trunking group

A set of up to four trunked ISLs.

trunking ports

The ports in a set of trunked ISLs.

TS

See [Time Server](#)

tunneling

A technique for enabling two networks to communicate when the source and destination hosts are both on the same type of network, but are connected by a different type of network.

U_Port

Universal port. A switch port that can operate as a G_Port, E_Port, F_Port, or FL_Port. A port is defined as a U_Port when it is not connected or has not yet assumed a specific function in the fabric.

WAN_TOV

Wide area network timeout value.

well-known address

In Fibre Channel technology, a logical address defined by Fibre Channel standards as assigned to a specific function and stored on the switch.

WWN

World wide name. An identifier that is unique worldwide. Each entity in a fabric has a separate WWN.

xlate domain

Translate domain. A router virtual domain that represents an entire fabric. Device connectivity can be achieved from one fabric to another, over the router and through this virtual domain, without merging the two fabrics. Also called a *phantom domain*.

zone

A set of devices and hosts attached to the same fabric and configured as being in the same zone. Devices and hosts within the same zone have access to others in the zone but are not visible to any outside the zone.

zone configuration

A specified set of zones. Enabling a configuration enables all zones in that configuration. See also [defined zone configuration](#), [enabled zone configuration](#).

zoning

A feature in fabric switches or hubs that allows segmentation of a node by physical port, name, or address.

Index

A

- account passwords, changing 24
- activating
 - licenses 92
 - Ports on Demand 26
- administering iSCSI configurations 80
- aliAdd command 88
- aliCreate command 88
- aliRemove command 88
- aliShow command 82, 88
- allocating frame buffers 22
- altBoot command 35
- arbitrated-loop device support 13
- archive files, maximum number of 95
- assigning domain IDs 27
- audience defined 7
- authorized reseller, HP 9

B

- backbone fabric ID 53
- benefits, interconnectivity 59

C

- cfgCreate command 49, 63
- cfgEnable command 49, 64
- cfgShow command 63, 64, 69, 82, 87, 88
- changing account passwords 24
- CHAP
 - configuring 80
 - displaying information about 81
 - removing secret 80
 - secret 78
- clearing system error log 97
- combining routing and tunneling services 12
- command help, displaying 31
- commands
 - aliAdd 88
 - aliCreate 88
 - aliRemove 88
 - aliShow 82, 88
 - altBoot 35
 - cfgCreate 49, 63
 - cfgEnable 49, 64
 - cfgShow 63, 64, 69, 82, 87, 88
 - configDefault 15
 - configDownload 34, 43
 - configShow 20, 33, 34
 - configUpload 34, 43
 - configure 21, 27, 51, 71, 87
 - dlsReset 29
 - dlsSet 29
 - dlsShow 29
 - errClear 97

- errShow 17, 95, 96
- fabricShow 22, 27, 62, 67, 68
- fanShow 42
- fcipShow 74
- fcPing 58
- fcrConfigure 53
- fcrPhyDevShow 49
- fcrProxyConfig 57
- fcrProxyDevShow 49, 68
- fcrResourceShow 58
- fcrXlateConfig 57
- firmwareDownload 103, 104
- firmwareShow 35
- help 31
- ifcsDisable 81
- ifcsEnable 81, 83
- interopMode 15
- iodReset 28
- iodSet 28
- ipAddrShow 34
- iscsiAuthCfg 80
- iscsiWwnAlloc 82
- licenseRemove 24
- licenseShow 34
- lsanZoneShow 49
- msPlatShow 21
- msplMgmtDeactivate 22
- netLoad 103
- nsAllShow 23, 62, 69
- nsShow 23, 49
- passwd 24
- portCfgDefault 15
- portCfgExPort 51, 54, 59
- portCfgFcip 72
- portCfgGige 72, 78
- portCfgLongDistance 15, 22
- portCfgTopology 14
- portDisable 15, 26
- portEnable 26
- portLogClear 68
- portLogDump 68
- portLogShow 97
- portPerfShow 93
- portShow 41, 53, 61, 74, 79
- portStart 53, 78
- portStop 15, 52, 61, 72, 78
- portType 72, 78
- psShow 43
- reboot 34
- rnPing 74
- secAuthSecret 55
- secModeDisable 22
- secModeShow 22
- supportShow 68

- switchDisable 21, 25, 27, 34, 53
- switchEnable 21, 25, 27, 53
- switchShow 23, 54, 60, 68
- switchStatusShow 41
- syslogDlpAdd 99
- syslogDlpRemove 99
- tempShow 43
- topologyShow 30, 93
- trunking 93
- trunkReset 93
- trunkSet 91, 93
- trunkShow 93
- tsClockServer 24, 71
- urouteConfig 30
- urouteShow 30
- zoneAdd 49, 88
- zoneCreate 49, 88
- zoneRemove 88
- zoneShow 82, 88
- zoning 89
- compatibility
 - feature 15
 - zone server 86
- configDefault command 15
- configDownload command 34, 43
- configShow command 20, 33, 34
- configUpload command 34, 43
- configuration
 - administering iSCSI 80
 - backing up 33
 - completing 68
 - displaying settings 33
 - printing information 34
 - restoring 34
 - summary of iSCSI steps 78
 - zoning 86
- configure command 21, 27, 51, 71, 87
- configuring
 - CHAP 80
 - fabrics for interconnectivity 61
 - FCIP interswitch link 72
 - FCIP ISL 72
 - interfabric link 52
 - iSCSI gateway zones 79
 - iSCSI portal 78
 - long-distance connection 22
 - McDATA fabric for interconnection 64
 - MP Router 61
 - one-way authentication 80
 - remote MP router 73
 - two-way authentication 80
 - zones 87
- confirming fabric connectivity 22
- connecting ISLs 21
- connecting to McDATA SANs 59
- connection, verifying 74
- connectivity
 - features 59
 - limitations 59

- McDATA 11
 - preparing switch for 63
 - verifying 23
- controlling routing 28
- conventions, document 8
- creating LSANs 45

D

- daemon overseer service 15
- databases, Platform Services 21
- delivery order, frame 28
- devices, proxy 47
- DH-CHAP secret, configuring 56
- disabling
 - DLS 29
 - failover 82
 - FCIP interswitch link 76
 - iFCS 81
 - ports 25
 - switches 25
 - syslogd 99
- displaying
 - CHAP information 81
 - command help 31
 - configuration settings 33
 - domain IDs 27
 - fabric-wide device count 23
 - iFCS information 81
 - installed version 36
 - iSCSI portal information 79
 - Name Server information 23
 - zoning information 81
- DLS
 - disabling 29
 - enabling 29
- dlsReset command 29
- dlsSet command 29
- dlsShow command 29
- document conventions 8
- domain IDs
 - assigning 27
 - displaying current list of 27
 - setting 27
- domains
 - front phantom 51
 - phantom 50
 - translate phantom 51
 - xlate 60
- dynamic load sharing 28

E

- E_port 14, 45
- edge fabric 45
- edge fabrics, zoning 48
- EFCM zone 65
- ELS 58
- enabling
 - DLS 29
 - failover 82

- FCIP interswitch link 76
- iFCS 81
- ports 25
- switches 25
- syslogd 99
- trunking 92
- errClear command 97
- error log 95
 - buffer number 97
- errShow command 17, 95, 96
- event log 95
- EX_port 15, 45, 59
 - connecting to edge fabric 51
- exchange-based trunking 13, 91
- extended link service request 58
- extended-edge PID mode 15

F

- F_port 14, 86, 101
- fabric connectivity, confirming 22
- fabric ID 45
 - backbone 53
- fabric parameters, matching 51
- Fabric Watch 15
- fabricShow command 22, 27, 62, 67, 68
- fabric-wide device count, displaying 23
- fan status, viewing 42
- fanShow command 42
- FC-FC routing service 11
- FCIP
 - configuring interswitch link 72
 - configuring ISL 72
 - link 59
 - tunneling service 12
- fcipShow command 74
- fcPing command 58
- fcrConfigure command 53
- fcrPhyDevShow command 49
- fcrProxyConfig command 57
- fcrProxyDevShow command 49, 68
- fcrResourceShow command 58
- fcrXlateConfig command 57
- feature compatibility 15
- features
 - connectivity 59
 - Fibre Channel switch 13
- Fibre Channel
 - over IP 71
 - proxy initiator 77
 - routing 45
 - switch features 13
- Fibre Channel NAT 50
- firmware, maintaining 35
- firmwareDownload command 103, 104
- firmwareShow command 35
- frame buffers, allocating 22
- frame delivery order, specifying 28
- front phantom domain 51
- FSPF protocol 50

G

- Gateway Service, iSCSI 12
- Gigabit Ethernet support 13

H

- hard zoning 86, 101
- hardware checks, performing 41
- hardware status, viewing 41
- help command 31
- help, obtaining 9
- high availability 17
- HP
 - authorized reseller 9
 - storage web site 9
 - subscriber's choice web site 9
 - technical support 9

I

- iFCS information, displaying 81
- ifcsDisable command 81
- ifcsEnable command 81, 83
- in-order frame delivery, forcing 28
- installed version, displaying 36
- installing
 - package 36
 - XPath OS 37
- interconnection, configuring McDATA fabric for 64
- interconnectivity
 - benefits 59
 - configuring fabrics for 61
- interfabric link, configuring 52
- Internet Small Computer Systems Interface, see iSCSI
- interopMode command 15
- iodReset command 28
- iodSet command 28
- ipAddrShow command 34
- iSCSI
 - administering configurations 80
 - configuring gateway zones 79
 - configuring portal 78
 - displaying portal information 79
 - Gateway Service 12, 77
 - mapping names 88
 - portals 77
 - proxy target 77
 - summary of configuration 78
- iscsiAuthCfg command 80
- iscsiWwnAlloc command 82
- ISLs, connecting 21

L

- licenseRemove command 24
- licenses
 - activating 92
- licenseShow command 34
- limitations, connectivity 59
- log entries, maximum 97
- logs

- error 95
- event 95
- port 95, 97
- long-distance connection, configuring 22
- loop mode topology 14
- LSANs
 - creating 45
 - naming convention 48
 - naming scheme 67
 - zoning 48, 67
- lsanZoneShow command 49
- LUN sharing 59

M

- maintaining
 - firmware 35
 - router configuration 33
- managing
 - port log 97
 - trunking 92
- mapping iSCSI names 88
- matching fabric parameters 51
- maximum log entries 97
- McDATA connectivity 11
- McDATA SANs, connecting to 59
- message severity levels 95
- meta-SAN 45
- modes
 - supported 59
 - Switch PID Address 21
- monitoring resources 58
- MP Router
 - configuring 61
- MP router
 - configuring remote 73
- msPlatShow command 21
- msplMgmtDeactivate command 22
- multiple user accounts 13
- multiprotocol routing services 11

N

- Name Server information, displaying 23
- Name Server support 13
- netLoad command 103
- nsAllShow command 23, 62, 69
- nsShow command 23, 49
- NTP server, synchronizing time with 23

O

- one-way authentication, configuring 80
- open E_port 15
- out-of-order frame delivery, restoring 28

P

- package, installing 36
- passwd command 24
- passwords, changing 24
- performing hardware checks 41

- phantom domains 50
- PID mode
 - extended-edge 15
 - requirements 15
 - verifying 20
- Platform Services database 21
- point-to-point topology 14
- port log 95, 97
 - field descriptions 98
 - management 97
- port status, viewing 41
- portals, iSCSI 77
- portCfgDefault command 15
- portCfgExPort command 51, 54, 59
- portCfgFcip command 72
- portCfgGige command 72, 78
- portCfgLongDistance command 15, 22
- portCfgTopology command 14
- portDisable command 15, 26
- portEnable command 26
- portLogClear command 68
- portLogDump command 68
- portLogShow command 97
- portPerfShow command 93
- ports
 - disabling 25
 - enabling 25
- Ports on Demand
 - activating 26
- portShow command 41, 53, 61, 74, 79
- portStart command 53, 78
- portStop command 15, 52, 61, 72, 78
- portType command 72, 78
- power supply status, modifying threshold 43
- primary iFCS router 80
- printing configuration information 34
- protocols, FSPF 50
- proxy
 - devices 47
 - Fibre Channel initiator 77
 - iSCSI target 77
 - setting ID 57
 - topology 47
- psShow command 43

R

- rack stability, warning 8
- reboot command 34
- recovery kernel, using 103
- related documentation, finding 7
- removing CHAP secret 80
- resources, monitoring 58
- restoring configuration 34
- rnPing command 74
- router configuration, maintaining 33
- router information, viewing 19, 29
- routing
 - between HP and McDATA fabrics 52
 - controlling 28

- FC-FC service 11
- Fibre Channel 45
- services 11
- unicast 29

routing and tunneling services, combining 12

S

- SAN scalability 51
- scalability 59
 - SAN 51
- SCC list 54
- secAuthSecret command 55
- secModeDisable command 22
- secModeShow command 22
- Secure Fabric OS 54
- service ready state 17
- services
 - combining 12
 - daemon overseer 15
 - FCIP tunneling 12
 - multiprotocol routing 11
- setting
 - domain ID 27
 - proxy ID 57
- severity levels, message 95
- shared secret 54
- soft zoning 86
- software installation support 103
- specifying frame delivery order 28
- states, service ready 17
- subscriber's choice, HP 9
- support
 - arbitrated-loop device 13
 - Gigabit Ethernet 13
- supported modes 59
- supportShow command 68
- Switch PID Address mode 21
- switchDisable command 21, 25, 27, 34, 53
- switchEnable command 21, 25, 27, 53
- switches
 - disabling 25
 - enabling 25
- switchShow command 23, 54, 60, 68
- switchStatusShow command 41
- symbols in text 8
- synchronizing time 71
- syslog daemon, using 98
- syslogd
 - disabling 99
 - enabling 99
 - on Linux systems 98
 - on UNIX systems 98
- syslogdIpAdd command 99
- syslogdIpRemove command 99
- system error log 95
- system error log, clearing 97
- system error log, viewing 96

T

- tempShow command 43
- text symbols 8
- threshold management 14
- time, synchronizing 71
- topology
 - loop mode 14
 - point-to-point 14
 - proxy 47
- topologyShow command 30, 93
- translate phantom domain 51
- trunking
 - commands 93
 - enabling 92
 - exchange-based 13, 91
 - groups 91
 - managing 92
- trunkReset command 93
- trunkSet command 91, 93
- trunkShow command 93
- tsClockServer command 24, 71
- two-way authentication, configuring 80

U

- unicast routing 29
 - viewing 30
- urouteConfig command 30
- urouteShow command 30
- user accounts, multiple 13
- user interfaces 14

V

- verifying
 - connection 74
 - connectivity 23
 - PID mode 20
- viewing
 - fan status 42
 - hardware status 41
 - port status 41
 - router information 19, 29
 - system error log 96
 - unicast routing information 30
- virtual E_ports 71

W

- warning
 - rack stability 8
- web sites
 - HP storage 9
 - HP subscriber's choice 9
- wwn mapping table, working with 82

X

- xlate domain 60
- XPath OS
 - and Secure Fabric OS 54
 - features 11

installing 37
syslogd CLI commands 98

Z

zone server 14
zoneAdd command 49, 88
zoneCreate command 49, 88
zoneRemove command 88
zones, EFCM 65
zoneShow command 82, 88
zoning
 commands 89
 configuration 86
 configuring 87
 displaying information about 81
 edge fabrics 48
 enforcement 86
 hard 86, 101
 LSANs 48, 67
 objects 86
 server compatibility 86
 soft 86
 terminology 86