reference guide

hp StorageWorks SAN design

Twenty-First Edition (June 3, 2005)

Part Number: AA-RMPNX-TE

This document is a guide to designing and building HP StorageWorks storage area networks (SANs). It describes how Hewlett-Packard storage systems, storage management tools, and Fibre Channel products can be used in heterogeneous SANs. Refer to the following URL for updates to this document.

http://h18006.www1.hp.com/products/storageworks/san/documentation.html

SAN design reference guide
Twenty-First Edition (June 3, 2005)
Part Number: AA-RMPNX-TE

contents

Figures

Tables

about this guide

About this guide topics include:

Changes from Previous Version

Significant changes from the May 2005 version of the SAN design reference guide include:

- Chapters 4, 9, and 10: added MP Router iSCSI support information.
- Chapters 8 and 9: updated EVA4000/6000/8000 support information.

Related Documentation

HP provides additional SAN information in these documents:

Topic	Documents
B-Series Product Line Switches and MP Router	For the latest information on B-Series switches see the storage infrastructure web site: http://h18006.www1.hp.com/storage/saninfrastructure.html.
C-Series Product Line Switches	For the latest information on C-Series switches see the storage infrastructure web site: http://h18006.www1.hp.com/storage/saninfrastructure.html.
M-series Product Line Switches	For the latest information on M-Series switches and firmware versions, see the storage infrastructure web site: http://h18006.www1.hp.com/storage/saninfrastructure.html.
Fabric Interoperability	■ *Fabric Interoperability: Merging Fabrics Based on C-Series and B-Series Fibre Channel Switches Application Notes* ■ *Fabric Interoperability: Merging Fabrics Based on C-Series and M-Series Fibre Channel Switches Application Notes* ■ *Fabric Interoperability: Merging Fabrics Based on M-Series and B-Series Fibre Channel Switches Application Notes* See these documents at the storage infrastructure web site: http://h18006.www1.hp.com/storage/saninfrastructure.html.
Continuous Access and SAN Extension	*Continuous Access and Data Replication Manager SAN Extensions Reference Guide* See this document at the storage infrastructure web site: http://h18006.www1.hp.com/storage/saninfrastructure.html.
Enterprise Backup Solution	EBS documents are available through the EBS web site: http://www.hp.com/go/ebs The *EBS Compatibility Matrix* is available through the tape compatibility web site: http://h18000.www1.hp.com/products/storageworks/tapecompatibility.html

Conventions

Conventions consist of the following:

- Document Conventions
- Text Symbols
- Getting Help

Document Conventions

The document conventions included in Table 1 apply in most cases.

Table 1: Document Conventions

Element	Convention
Cross-reference links	Blue text: Figure 1
Key and field names, menu items, buttons, and dialog box titles	**Bold**
File names, application names, and text emphasis	*Italics*
User input, command and directory names, and system responses (output and messages)	`Monospace font` `COMMAND NAMES` are uppercase monospace font unless they are case sensitive
Variables	*`<monospace, italic font>`*
Web site addresses	Blue, underlined sans serif font text: http://www.hp.com

Text Symbols

The following symbols may be found in the text of this guide. They have the following meanings.

WARNING: Text set off in this manner indicates that failure to follow directions in the warning could result in bodily harm or death.

Caution: Text set off in this manner indicates that failure to follow directions could result in damage to equipment or data.

Note: Text set off in this manner presents commentary, sidelights, or interesting points of information.

Getting Help

For additional information about HP SANs, or SAN based products, contact an HP authorized service provider or access our website: http://www.hp.com.

HP Technical Support

In North America, call technical support at 1-800-652-6672, available 24 hours a day, 7 days a week.

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

Outside North America, call technical support at the nearest location. Telephone numbers for worldwide technical support are listed on the HP website under support: http://thenew.hp.com/country/us/eng/support.html.

Be sure to have the following information available 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

HP Storage Web site

The HP web site has the latest information on this product, as well as the latest drivers. Access storage at: http://thenew.hp.com/country/us/eng/prodserv/storage.html. From this website, select the appropriate product or solution.

HP Authorized Reseller

For the name of your nearest HP authorized reseller:

- In the United States, call 1-800-345-1518
- In Canada, call 1-800-263-5868
- Elsewhere, see the HP website for locations and telephone numbers: http://www.hp.com.
-

Volume 1

Architecture

SAN architecture is presented in these chapters:

SAN design overview

1

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

This chapter describes the following topics:

SAN solutions

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

- **Open systems**

  SANs support various operating systems and servers to meet your operational requirements. A robust storage infrastructure accommodates new business models, unexpected growth, and corporate reorganizations.

- **Fast backup and restore**

  SANs remove backup and recovery traffic from the LAN, reducing congestion, improving backup windows, and efficiently utilizing storage resources. You can use centrally managed, high-performance tape libraries to reduce backup overhead.

- **Business continuance**

  SANs can eliminate single points of failure, incorporate failover software, and support mirroring at geographically dispersed data centers for disaster recovery. You can quickly restore productivity after a power failure or component downtime.

- **High availability**

  Redundant fabric designs, storage replication, dynamic failover protection, traffic rerouting, and server clustering enable SANs to provide enterprise-class availability to open systems servers.

- **Server and storage consolidation**

  Multiple servers and backup systems can share storage for efficient processing and increased availability.

- **Cost savings**

  SAN total cost of ownership (TCO) is typically less than direct-attached storage (DAS). The business realizes a higher return on investment (ROI) because sharing storage among servers utilizes capacity more efficiently, and expenses for backup hardware are reduced. Increased system availability can help prevent costly downtime and lost data.

- **Centralized management**

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

- **Security**

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

- **Online scalability**

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

- **Modularity**

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

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

HP SAN implementations

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

- **Flexible design and deployment**

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

- **Incremental scaling**

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

- **Interoperability**

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

- **Geographically dispersed installations**

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

For information about:

- SAN infrastructure solutions and cost efficiency, see the SAN Information Center web site: http://h18006.www1.hp.com/storage/saninfrastructure.html.
- SAN deployment in small-to-medium sized businesses, see the HP StorageWorks Solution Center: http://www.hp.com/sbso/serverstorage/san.html.
- The Small Business SAN Kit, see the MSA1000 web site: http://h18006.www1.hp.com/products/storageworks/msa1000smb/index.html.

SAN components

A SAN consists of the following hardware and software components:

- **Switches**

  A Fibre Channel switch creates the fabric of the SAN. By interconnecting switches, you can create scalable SANs with thousands of port connections.

- **Routers, bridges, and gateways**

  Router functionality provides high levels of scalability, dynamic device sharing, and Fibre Channel network fault isolation. Routers extend the SAN over long distances and enable integration of multiprotocol technologies.

- **Storage devices**

  A SAN can integrate multiple storage system types, such as disk arrays and tape libraries, to allocate storage efficiently.

- **Servers and HBAs**

  Host bus adapters (HBAs) connect the server to the SAN. HBA drivers provide an intelligent interface to the switches and minimize CPU overhead.

- **Cabling and cable connectors**

  Fiber optic cables provide the physical connections between SAN components.

- **SAN management applications**

  HP applications manage and monitor components and ensure optimal SAN operation.

Fibre Channel technology

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

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

- RAID arrays
- Tape devices and backup libraries

Fibre Channel technology supports simultaneous use of these transport protocols:

- IP
- SCSI
- iSCSI

For the latest information on Fibre Channel and related technologies visit this website:

www.incits.org

Storage area networks

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

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

SAN infrastructure

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

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

Fabrics

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

A collection of switches can be connected as a single fabric or partitioned into logical, separate fabrics (LSANs for B-Series or VSANs for C-Series) to form an interconnected network of independent fabrics.

SAN scaling

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

- **Fibre Channel architecture**

  Fibre Channel allows a maximum of 256 ports per switch and 239 switches in a single fabric. HP specifies support based on rules for the maximum number of switches and maximum number of ports in a single fabric or multifabric SAN. Using many switches to obtain a high number of ports is unacceptable if the fabric exceeds the total switch count limit. Likewise, using large-capacity switches can create a network that exceeds the maximum number of ports. For the HP-supported switch and port count fabric maximums, see “B-Series switches and fabric rules” on page 81, “C-Series switches and fabric rules” on page 97, and “M-Series switches and fabric rules” on page 105.

- **Supported configurations**

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

- **Fabric services**

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

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

Fibre Channel switches

A switch is identified by its function in a SAN:

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

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

Switch rules

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

SAN design approaches

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

- **HP standard design**

  HP standard designs specify the arrangement of Fibre Channel switches in a SAN fabric, and are optimized for specific data access requirements and typical workloads. Implementing a standard design is the simplest approach to SAN design. HP recommends this approach for users who are designing a SAN for the first time.

- **Modified HP standard design**

  Select a standard SAN design that satisfies most of your requirements, and then modify it to meet your data access and connectivity requirements. HP recommends this approach for users with an intermediate level of SAN experience.

- **Custom design using the HP SAN design rules**

  Use a custom SAN design for specific storage and data access requirements. The SAN design rules in this guide specify guidelines for configuring custom topologies. HP recommends this approach for users with an intermediate or advanced level of SAN experience.

For information about:

- Standard SAN designs, see "SAN fabric topologies" on page 41.
- Customizing a SAN design, see
  - "B-Series switches and fabric rules" on page 81.
  - "C-Series switches and fabric rules" on page 97.
  - "M-Series switches and fabric rules" on page 105.
- Heterogeneous SAN design, see
  - "Heterogeneous server rules" on page 127.
  - "SAN storage system rules" on page 171.
- Recommended SAN solutions and conventions, see “Best practices” on page 349.

Design considerations

To design or modify a SAN, evaluate the following design considerations:

- **Geographic layout**

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

  For information about supported distances, see “SAN fabric connectivity and switch interoperability rules” on page 113.

- **Data availability**

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

  For information about availability levels, see “Data availability” on page 59.

- **Connectivity**

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

  For information about the connections available in a SAN topology, see "SAN fabric topologies" on page 41.

- **Storage capacity**

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

  For storage systems information, see "SAN storage system rules" on page 171.

- **Heterogeneous platforms and operating systems**

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

  For information about configuring systems in a heterogeneous environment, see "Heterogeneous server rules" on page 127.

- **Scalability and migration**

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

  For information on scaling and migrating, see “Best practices” on page 349.

- **Backup and restore**

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

  For information about centralized backup, see “Enterprise Backup Solution” on page 205.

- **Disaster tolerance**

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

  For information about disaster tolerance and failover protection, see:

  — "SAN management" on page 273

  — "SAN extension" on page 209

  - Continuous Access and Data Replication Manager SAN extensions reference guide

- **Switch and hop counts**

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

  For information about switches and hop counts, see:

  - "B-Series switches and fabric rules" on page 81
  - "C-Series switches and fabric rules" on page 97
  - "M-Series switches and fabric rules" on page 105

- **Oversubscription**

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

  For information about oversubscription, see "Recommended ISL ratios" on page 51.

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

  Provide an adequate level of performance based on application workloads. For frequent data reference and quick response times, use local, high-capacity paths to connect servers and storage systems. Deploy servers and storage in your SAN based on your data access requirements. See "SAN fabric topologies" on page 41.

- **Manageability**

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

  For information about SAN management, see "SAN management".

- **Fabric zoning**

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

  For information about zoning, see, "B-Series switches and fabric rules" on page 81, "C-Series switches and fabric rules" on page 97, and "M-Series switches and fabric rules" on page 105.

- **Selective Storage Presentation**

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

  For information about SSP, see "Selective Storage Presentation" on page 290.

- **SAN security**

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

  For information about security, see "SAN security" on page 307.

- **Fibre Channel routing functionality**

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

  For routing functionality information, see "Fibre Channel routing", on page 65.

- **Virtual SANs (C-Series switches)**

  To create a SAN with multiple logical SANs with separate fabric services, implement virtual SANs (VSANs). Add the inter-VSAN routing feature to enable device sharing across VSANs. For information about VSANs, see "Fibre Channel routing", on page 65.

SAN fabric topologies

2

This chapter discusses HP standard storage area network (SAN) fabric topologies. It describes the following topics:

Overview

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

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

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

- "B-Series switches and fabric rules" on page 81
- "C-Series switches and fabric rules" on page 97
- "M-Series switches and fabric rules" on page 105
- "Heterogeneous server rules" on page 127
- "SAN storage system rules" on page 171.

Fabric topologies

A *SAN fabric topology* defines the arrangement of Fibre Channel switches in a fabric. HP supports the following SAN fabric topologies to meet your data access and connectivity requirements:

- Single-switch fabric, page 43
- Cascaded fabric, page 44
- Meshed fabric, page 46
- Ring fabric, page 48
- Core-edge fabric, page 50

Routed SAN fabrics

HP standard fabric topologies support Fibre Channel routing. Fibre Channel routing enables connectivity between devices in multiple fabrics or multiple virtual SANs (VSANs). HP supports the following routed fabric technologies:

- B-Series Meta SAN, page 57
- C-Series VSANs with IVR, page 58

Benefits

With HP standard SAN fabric topologies, you can:

- Create a SAN fabric for each department or application in your organization.
- Perform centralized management and backups.
- Update a SAN fabric to accommodate changing capacity or data access needs. You can also convert to another SAN fabric topology as needed.
- Connect devices over long distances using extended Fibre Channel or IP connections. (See "SAN fabric connectivity rules" on page 114 and "SAN extension" on page 209.)
- Connect multiple SAN fabrics using routing technology. (See "B-Series Meta SAN" on page 57.)
- Deploy multiple VSANs. (See "C-Series VSANs with IVR" on page 58.)
- Incorporate a range of SAN availability levels. (See "Data availability" on page 59.)

Single-switch fabric

This section describes the following topics:

Overview

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

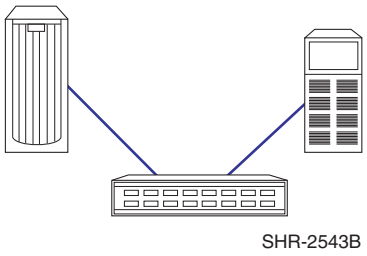

Figure 1: Single-switch fabric

Switch models

For a small, single-switch SAN fabric, use an HP SAN, Fabric, or Edge switch (4, 8, or 16 ports). For a larger single-switch SAN fabric, use a SAN, Fabric, or Edge switch (32 to 40 ports), or a Core or Director switch (64 to 240 ports), which have higher port counts. For a high-availability SAN, use two switches configured in a dual-fabric SAN.

Benefits

The benefits of a single-switch fabric are:

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

Cascaded fabric

This section describes the following topics:

Overview

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

Cascading enables you to:

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

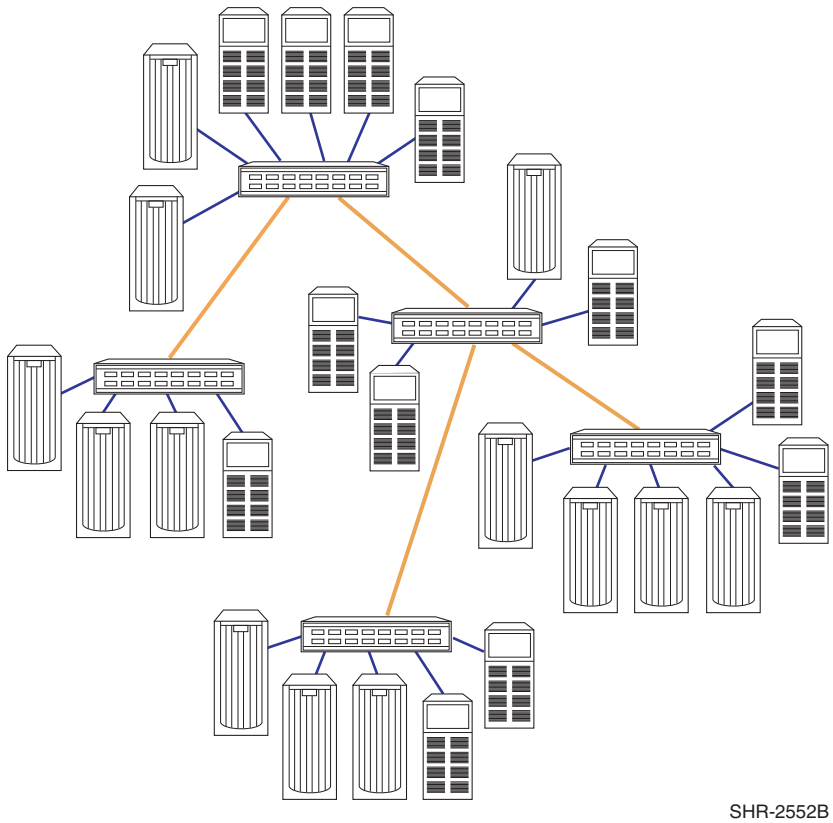

Figure 2: Cascaded fabric

Switch models

All HP Fibre Channel switches are supported for use in a cascaded fabric topology. Cascaded fabric topologies typically use SAN, Fabric, or Edge switches, which support smaller incremental growth.

Note: Over time, a cascaded fabric topology can result in increased hops between switches. B-Series and C-Series fabrics must not exceed seven hops; M-Series fabrics must not exceed three hops. For additional switch hop information, see "B-Series switches and fabric rules" on page 81, "C-Series switches and fabric rules" on page 97, and "M-Series switches and fabric rules" on page 105.

Benefits

The benefits of a cascaded fabric are:

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

Meshed fabric

This section describes the following topics:

Overview

A meshed fabric is a group of interconnected switches that have two or more ISLs between different switches for fabric resiliency (Figure 3). If one ISL fails, the switch automatically reroutes data through an alternate path in the fabric. If the alternate path includes other switches, the data must pass through those switches to reach its destination.

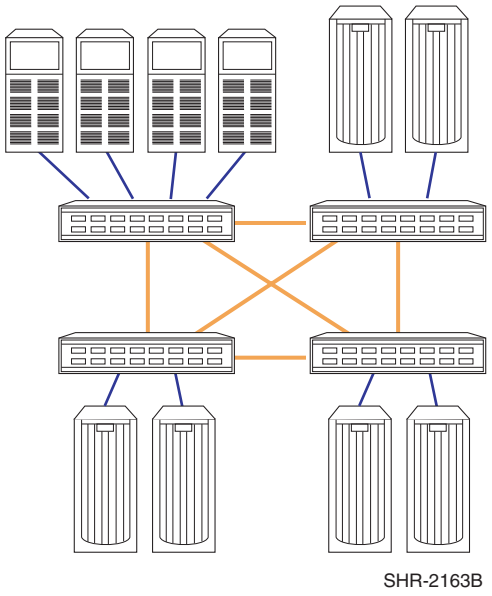

Figure 3: Meshed fabric

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

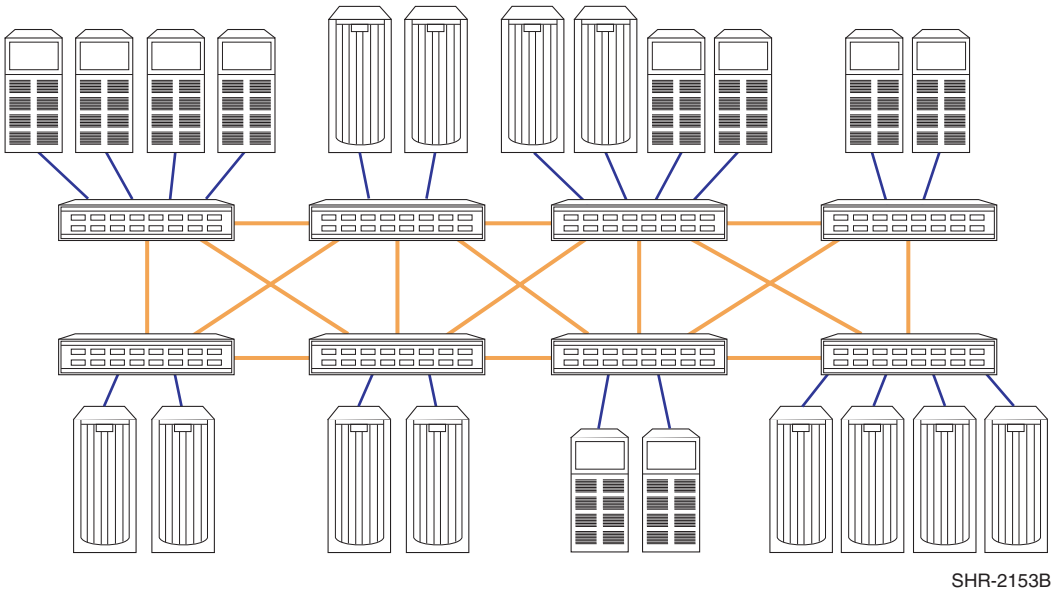

Figure 4: ISL connections in a meshed fabric

Switch models

All HP Fibre Channel switches are supported for use in a meshed fabric topology. Meshed fabric topologies typically use SAN, Fabric, or Edge switches, which support smaller incremental growth. To meet higher port-count requirements, use Core or Director switches.

Benefits

The benefits of a meshed fabric are:

- Ability to meet multiple data access needs
- Multiple paths for internal fabric resiliency
- Ease of scalability
- Shared backup and management support
- Support for a mix of local and distributed data access (see “Topology data access” on page 54)
- Less impact on performance due to intraswitch traffic

Ring fabric

This section describes the following topics:

Overview

A ring fabric is a continuous ring of interconnected switches (Figure 5). Each switch in the ring connects to the adjacent switch; the last switch connects to the first switch. The ring fabric provides a similar level of fabric resiliency as the meshed fabric and ensures full fabric connectivity with a minimum of two paths for each switch.

The ring fabric enables you to:

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

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

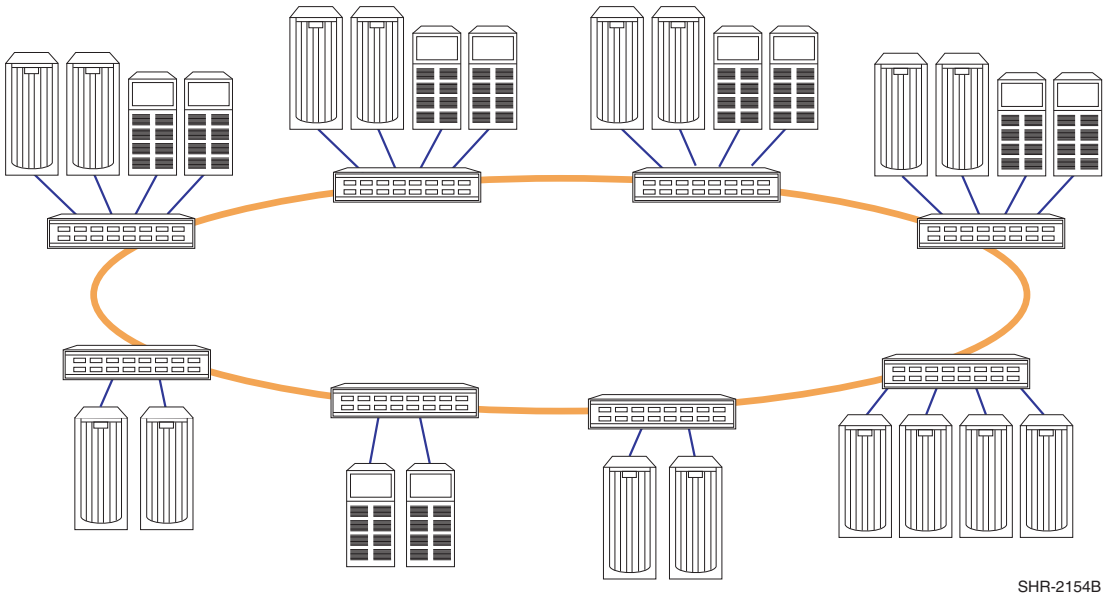

Figure 5: Ring fabric

If the ring fabric has fewer than 12 switches, you can add switches (called satellite switches) outside the ring to create more user ports. For example, you can connect 23 satellite switches to an 11-switch ring to create a 34-switch fabric that adheres to the 7-hop maximum of B-Series and C-Series switches (Figure 6). M-Series switches have a 3-hop maximum, reducing the total number of switches that can be configured in a ring to 7; satellite switches are not supported.

Note: For further information on switch-specific fabric maximums, see "B-Series switches and fabric rules" on page 81, "C-Series switches and fabric rules" on page 97, and "M-Series switches and fabric rules" on page 105.

Note: Adding satellite switches slightly reduces fabric availability.

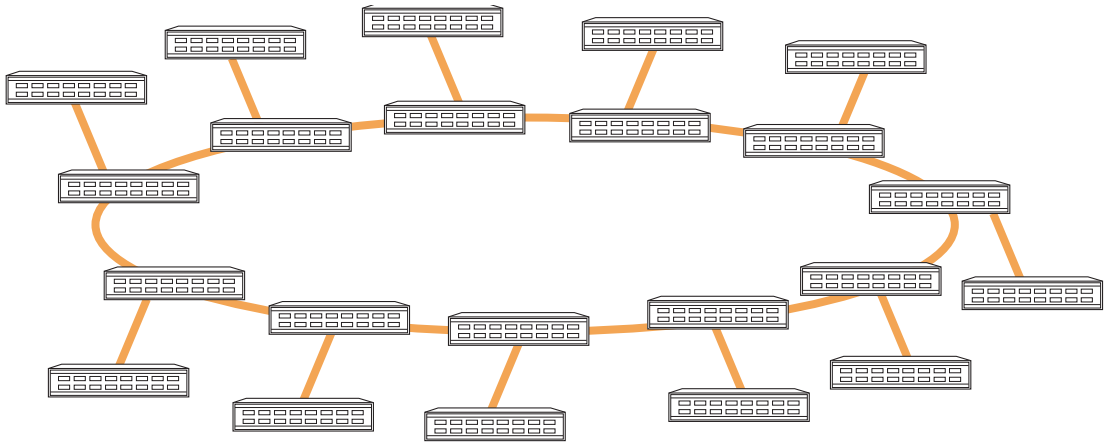

SHR-2544B

Figure 6: Ring fabric with satellite switches

Switch models

All HP Fibre Channel switches are supported for use in a ring fabric topology. Ring fabric topologies typically use SAN, Fabric, or Edge switches, which support smaller incremental growth. To meet higher port-count requirements, use Core or Director switches.

Benefits

The benefits of a ring fabric are:

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

Core-edge fabric

HP recommends using a core-edge fabric wherever possible.

This section describes the following topics:

Overview

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

The core-edge fabric is optimal for:

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

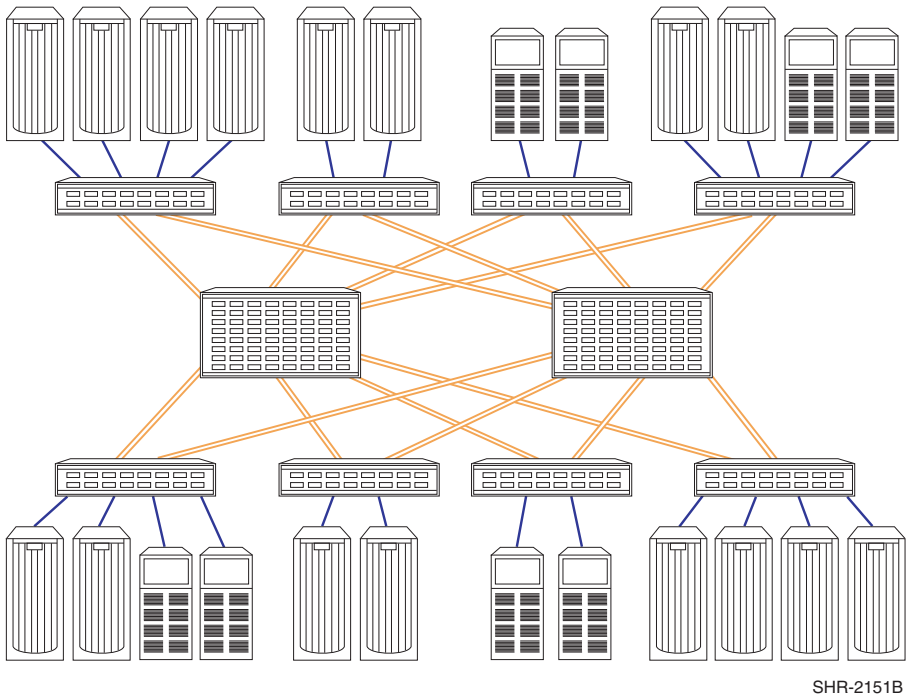

Figure 7: Core-edge fabric

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

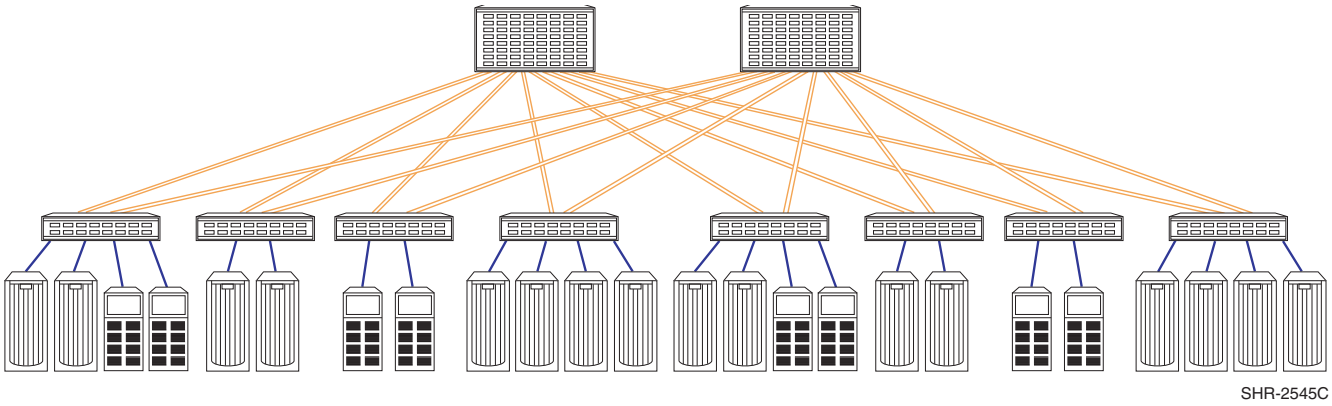

Figure 8: Core-edge fabric (hierarchical)

Core-edge fabric types

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

This section describes the following topics:

Fat and skinny trees

There are two core-edge fabric topology types: fat tree and skinny tree. Table 2 describes fat and skinny trees.

Table 2: Core-edge fabric topology types

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

Recommended ISL ratios

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

The ISL ratio is also affected by the speed of the server, storage, and ISL ports. The highest Fibre Channel speed (4 Gb/s) is supported between switches. Consider these factors when determining the ideal ISL ratio.

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

- Location of server and storage fabric connection
- Server and storage hardware performance
- Data access type (see "Topology data access" on page 54)
- Server application performance requirements

Table 3 describes the recommended core-edge fabric ISL ratios.

Table 3: Recommended core-edge fabric ISL ratios

I/O intensity	Recommended ratios
Most I/O data intensive application needs (> 70 MB/s)	1:1 3:1
Less I/O data intensive application needs (< 70 MB/s)	7:1 15:1

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

Numeric representation

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

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

Figure 9 shows a 4 x 12 core-edge fabric with 4 core switches and 12 edge switches, each connected to the core with 4 ISLs. This provides a total of 304 user ports, including 160 in the core and 144 on the edge switches.

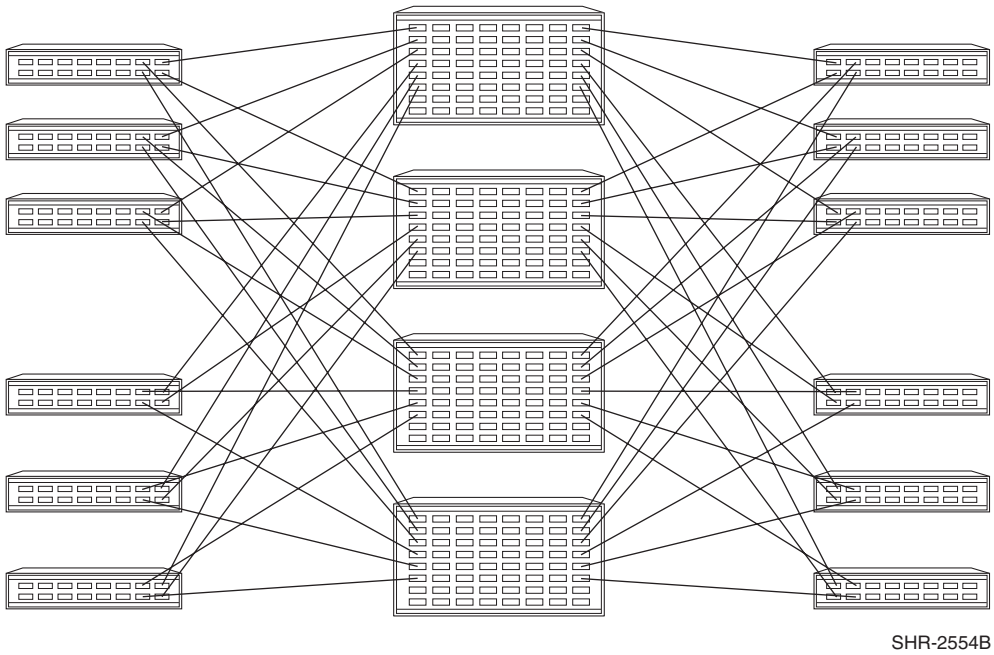

Figure 9: 4 x 12 core-edge fabric

Switch models

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

Benefits

The benefits of a core-edge fabric are:

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

Topology data access

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

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

Table 4 lists the data access performance ratings for each SAN fabric topology.

Table 4: Data access performance by SAN fabric topology

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

Topology maximums

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

The tables in this section describe the switch and port maximums for each topology type by switch family.

Consider the following:

- User ports are for server and storage connections.
- It is assumed that you have the minimum number of ISLs. If you require more ISLs, this reduces the number of user ports available for server and storage connections. See the following chapters for configuration limits:
  - "B-Series switches and fabric rules" on page 81
  - "C-Series switches and fabric rules" on page 97
  - "M-Series switches and fabric rules" on page 105
- If you connect a Storage Management Appliance to the fabric, this further reduces the number of ports available for server and storage connections. (See "SAN management" on page 273 for more information.)

B-Series switches

Table 5 lists the B-Series switch and port maximums for specific fabric topologies.

Table 5: B-Series switch and port topology maximums

SAN topology	Number of switches	Total number of ports	Number of user ports
Single-switch fabric	1	128	128
Cascaded fabric	34	1,280	1,234
Meshed fabric	34	1,280	1,118
Ring fabric	15	1,280	1,250
Ring fabric with satellite switches	34	1,280	1,212
Core-edge fabric	34	1,280	1,214 single core 1,152 dual core

C-Series switches

Table 6 lists the C-Series switch and port maximums for specific fabric topologies.

Table 6: C-Series switch and port topology maximums

SAN topology	Number of switches	Total number of ports	Number of user ports
Single-switch fabric	1	224	224
Cascaded fabric	40	4,000 (maximum of 12 Director switches)	3,500 (cascaded with 12 Director switches and 10 Fabric switches)
Core-edge fabric	40	4,000 (maximum of 12 Director switches)	3,500

M-Series switches

Table 7 lists the M-Series switch and port maximums for specific fabric topologies.

Table 7: M-Series switch and port topology maximums

SAN topology	Number of switches	Total number of ports	Number of user ports
Single-switch fabric	1	140	140
Cascaded fabric	24 (maximum of 8 Director switches)	1,632	1,024 (cascaded with 8 Director switches and 16 Edge switches)
Meshed fabric	N/A (topology exceeds 3-hop-count limit)	N/A	N/A
Ring fabric	7	980	966
Ring fabric with satellite switches	N/A	N/A	N/A
Core-edge fabric	24	1,632	1,024

Routed fabric topologies

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

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

- B-Series Meta SAN—Implemented using the B-Series Multi-protocol Router (MP Router), which provides selective Fibre Channel routing connectivity between multiple B-Series fabrics.
- C-Series VSANs with IVR—Implemented using the C-Series switch inter-VSAN routing (IVR) functionality. IVR functionality provides selective Fibre Channel routing connectivity between devices in different VSANs. Ports on one or more switches can be assigned to different VSANs.

B-Series Meta SAN

This section describes the following topics:

- Overview
- Switch models and fabric topologies
- Benefits

Overview

A Meta SAN contains multiple B-Series fabrics connected using the B-Series MP Router. The MP Router implements the Fibre Channel routing service, which allows selective access between devices in different fabrics without having to merge fabrics. This provides a high level of isolation between fabrics. This isolation can be viewed as individual Fibre Channel subnetworks within the Meta SAN.

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

For more information about Meta SANs and Fibre Channel routing, see "Fibre Channel routing" on page 65.

Switch models and fabric topologies

HP supports Meta SANs with all B-Series switches in either HP standard or customized topologies. You must follow all B-Series and Meta SAN fabric rules. For more information, see "B-Series switches and fabric rules" on page 81.

Benefits

A Meta SAN:

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

- Enables higher levels of storage consolidation
- Provides centralized backup for multiple fabrics
- Allows higher levels of fabric management and management consolidation

C-Series VSANs with IVR

This section describes the following topics:

Overview

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

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

Switch models and fabric topologies

- HP supports VSANs with all C-Series switches in either HP standard or customized topologies. You must follow all C-Series and VSAN fabric rules. See "C-Series switches and fabric rules" on page 97 for more information.

Benefits

A VSAN:

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

Data availability

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

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

Note: See "High-availability MP Router configurations" on page 75 for information on high-availability configurations when using the HP B-Series MP Router or C-Series VSANs.

This section describes the following topics:

Factors

Several factors affect SAN data availability:

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

Levels

The data availability levels are:

Level 1: single connectivity fabric

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

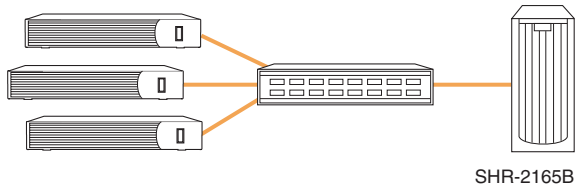

Figure 10: Level 1: single connectivity fabric

Level 2: single resilient fabric

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

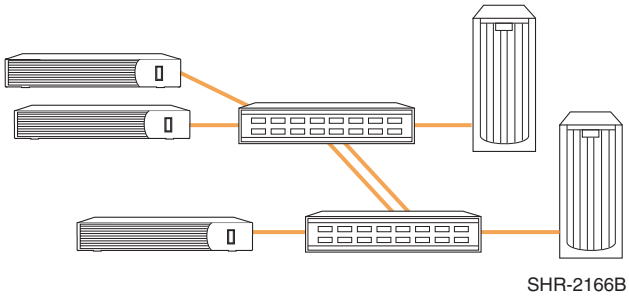

Figure 11: Level 2: single resilient fabric

Level 3: single resilient fabric with multiple device paths

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

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

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

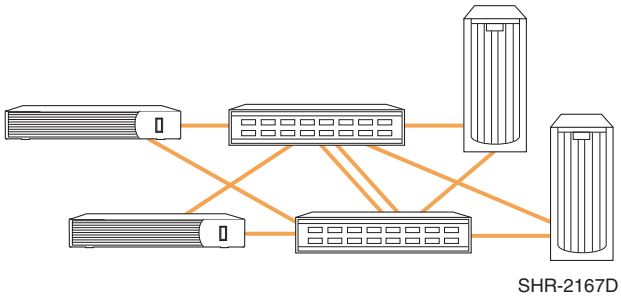

Figure 12: Level 3: single resilient fabric with multiple device paths

Level 4: multiple fabrics and device paths (NSPOF)

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

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

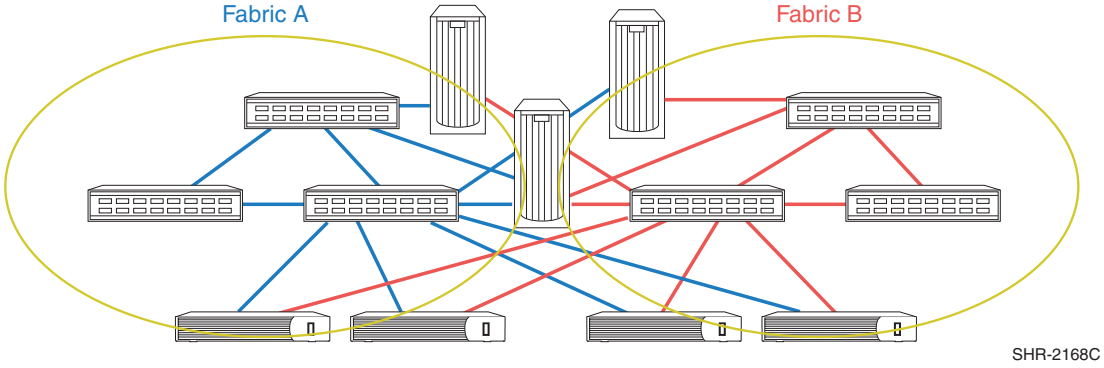

Figure 13: Level 4: multiple fabrics and device paths (NSPOF)

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

Table 8 indicates data availability and supported topologies for each level.

Table 8: Fabric design data availability

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

Considerations

When choosing a data availability level, you must consider:

- Cost
- Access to critical data

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

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

Table 9: Calculating data availability level costs

Data availability level	Hardware cost	Number of available ports
Level 1: single connectivity fabric	x*	# ports = n – number of ISL ports†
Level 2: single resilient fabric	x + additional ISLs	# ports = n – number of ISL ports
Level 3: single resilient fabric with multiple device paths	x + additional ISLs + additional HBAs	# ports = n – number of ISL ports - additional HBA ports‡
Level 4: multiple fabrics and device paths (NSPOF)	x + additional ISLs + additional HBAs + additional switches	# ports = 2n – number of ISL ports - additional HBA ports

*. x is the cost of a single connectivity fabric.
†. n is the total number of ports for servers and storage systems.
‡. May require use of zoning to define a minimum of two data paths in a single fabric—operating system dependent.

Topology migration

To increase SAN connectivity and capacity:

- Increase the number of switches.
- Use switches with more ports.
- Implement multiple fabrics.
- Migrate to another fabric topology.

This section describes the following topics:

Nondisruptive migration

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

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

Migrating a cascaded fabric SAN

Migration paths for a cascaded fabric SAN are:

Cascaded to meshed

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

Cascaded to ring

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

Cascaded to core-edge

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

Migrating a meshed fabric SAN

Migration paths for a meshed fabric SAN are:

- Meshed to ring
- Meshed to core-edge

Meshed to ring

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

Meshed to core-edge

Use the method described in "Cascaded to core-edge" on page 63.

Migrating a ring fabric SAN

Migration paths for a ring fabric SAN are:

- Ring to meshed
- Ring to core-edge

Ring to meshed

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

Ring to core-edge

This migration is less disruptive if you have two ISLs between all switches in the ring fabric SAN. Use the method described in "Cascaded to core-edge" on page 63.

Fibre Channel routing

3

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

Fibre Channel routing overview

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

Note: This chapter uses the terms "fabric" and "VSAN" interchangeably. HP does not support using the B-Series MP Router and C-Series VSAN inter-VSAN routing (IVR) functionality in the same SAN configuration.

Routing enables independent fabrics or VSANs with IVR to dynamically share devices without the need to reconfigure or recable physical connections.

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

Fibre Channel routing features include:

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

This section describes the following topics:

Fabric and VSAN independence

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

Fabric services

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

The fabric services manage:

- Device names and addresses
- Timestamps
- Switch utilities

Routing connects devices in multiple fabrics or VSANs without extending fabric services from one routed fabric to another. Devices in a routed network can communicate across logical SANs (LSANs) or VSANs despite differing fabric services configurations.

World wide name

A recognized naming authority assigns each Fibre Channel device a unique identifier, the world wide name (WWN). Use the device WWNs to:

- assign devices to zones
- define devices to export from one fabric or VSAN to another

Import and export

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

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

Routing table

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

SAN scaling and routing

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

- Increase Fibre Channel switch capability within a fabric. (See "Switch scaling" on page 68.)
- Provide connections between independent fabrics using a Fibre Channel router or VSANs with IVR. (See "Scaling by routing" on page 69.)

Switch scaling

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

- Switch scaling limits
- Fabric services limits

Switch scaling limits

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

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

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

- "B-Series switches and fabric rules" on page 81
- "C-Series switches and fabric rules" on page 97
- "M-Series switches and fabric rules" on page 105
- "SAN fabric connectivity and switch interoperability rules" on page 113

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

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

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

Fabric services limits

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

Fabric services

Fabric services include:

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

Sample fabric service

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

Coordinating fabric services

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

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

Scaling by routing

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

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

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

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

Fibre Channel routing implementations

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

- Connecting several fabrics using a router (Figure 14 on page 70)
- Dividing a single fabric into several smaller VSANs (Figure 15 on page 70)

This section describes the following topics:

Fibre Channel routing techniques

There are two Fibre Channel routing techniques:

- A B-Series Multi-Protocol Router (MP Router) connects independent fabrics (SAN islands), as shown in Figure 14.
- A C-Series switch with IVR connects multiple VSANs, as shown in Figure 15.

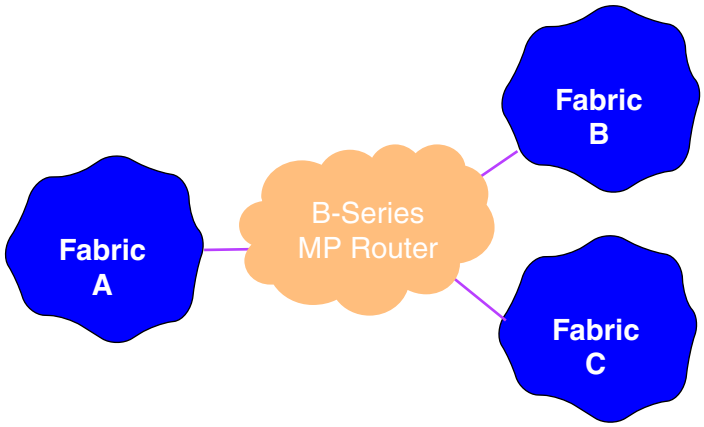

Figure 14: Basic MP Router configuration

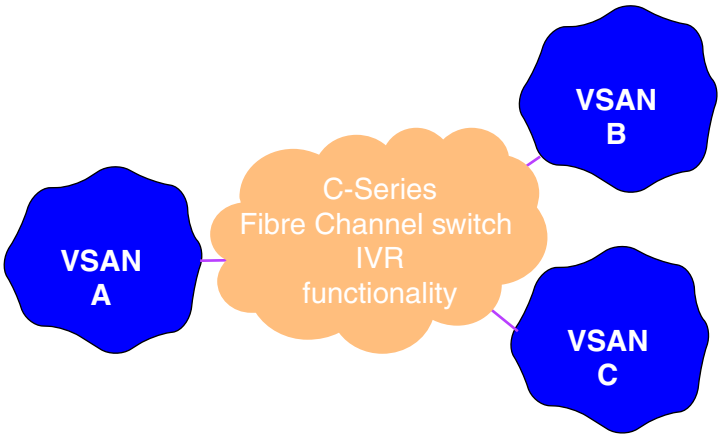

Figure 15: Basic IVR configuration

B-Series fabric groups

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

Figure 14 on page 70 shows Fabric A, Fabric B, and Fabric C, each containing one or more switches. Any B-Series switch can be used in these fabrics. In each fabric, the switches must run the same version of switch firmware and must have the same variable settings (for example, R_A_TOV). Each fabric has a unique set of fabric services. See "B-Series switches and fabric rules" on page 81 for fabric restrictions.

Fabrics connected by an MP router must comply with configuration rules for a routed fabric. (See "MP Router fabric rules" on page 90.) The fabrics may have identical domain names and zoning definitions.

The MP Router also provides FCIP capabilities, allowing implementation of Fibre Channel routing and FCIP SAN extension. (See "Integration of Fibre Channel routing and FCIP" on page 77.)

C-Series fabric division

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

Figure 15 on page 70 shows VSAN A, VSAN B, and VSAN C, each a set of switch ports on one or more C-Series switches. A VSAN can extend across multiple switches. Each VSAN has a unique set of fabric services with independent fabric management. VSANs can share devices by using the license-enabled IVR function. IVR is distributed across all switches in the SAN, and there is no separate router hardware. Because the switches are a connected set, they must run the same version of switch firmware.

B-Series and C-Series routing differences

An MP Router or VSAN IVR function can connect existing fabrics or VSANs. Using an MP Router, existing fabrics are physically connected to the router, and the router creates the Meta SAN. Using C-Series switches, existing fabrics are physically connected, and the routing function in the switches is configured using IVR.

Figure 16 and Figure 17 show the differences between B-Series and C-Series routing.

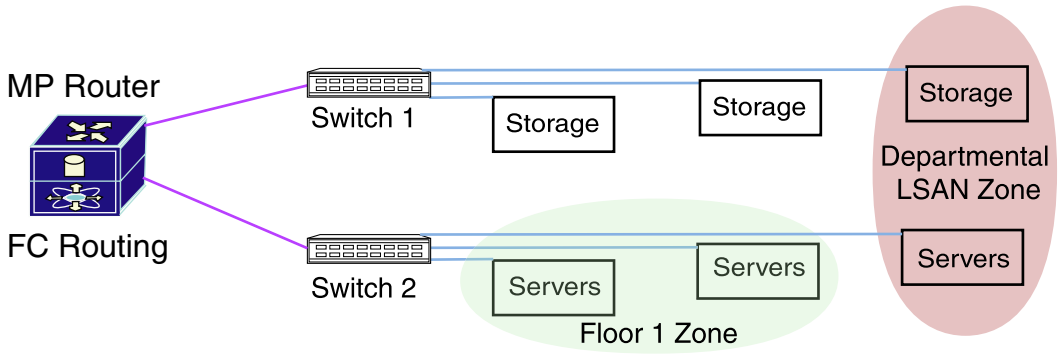

Figure 16: B-Series routing

Independent fabrics are connected at the MP Router. An LSAN can include devices connected to the same switch (for example, Floor 1 LSAN) or to multiple switches (for example, Departmental LSAN.)

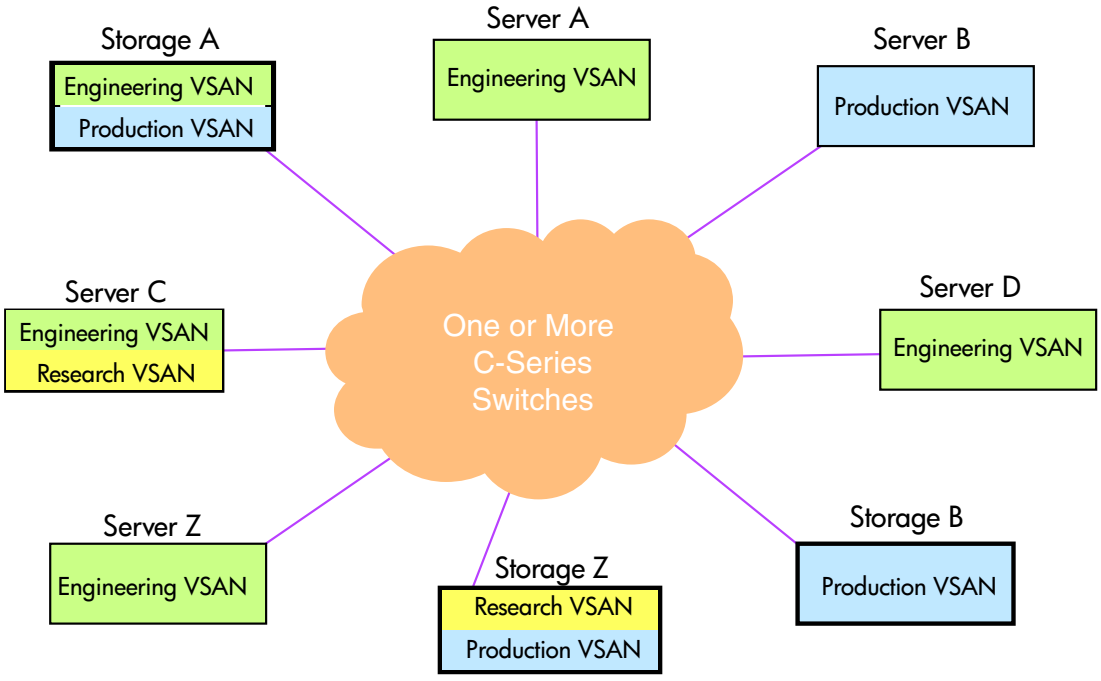

Figure 17: C-Series routing

VSANs can include devices that reside on a single switch or multiple switches in the SAN. Devices in different VSANs can communicate by using IVR. Multiple switches connect in any supported fabric configuration.

Fabric redundancy and routing

An MP router can connect one group of fabrics, or a single fabric can connect multiple VSANs with IVR. For a high-availability, fully redundant implementation, you can have two routers and two groups of fabrics, or two multi-VSAN fabrics with IVR.

This section describes the high-availability dual-redundant routed SAN.

High-availability dual-redundant routed SAN

Using two routers, you can configure a second set of fabrics as a dual-redundant Meta SAN (Figure 18). See "High-availability MP Router configurations" on page 75.

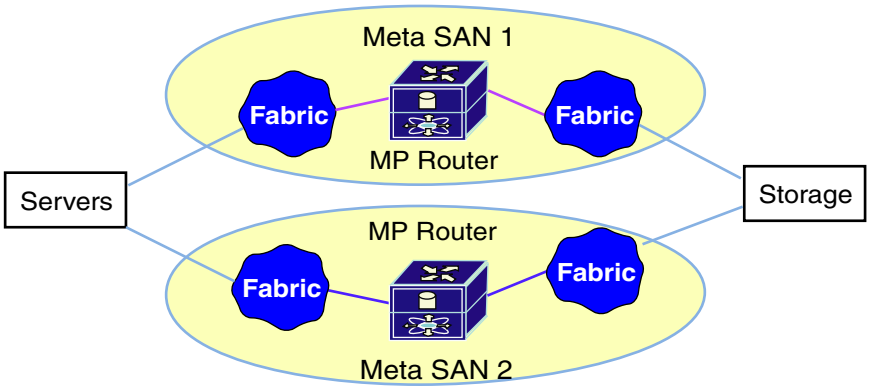

Figure 18: Dual-redundant Meta SAN

You can configure two IVR-connected VSAN groups for a dual-redundant VSAN (Figure 19).

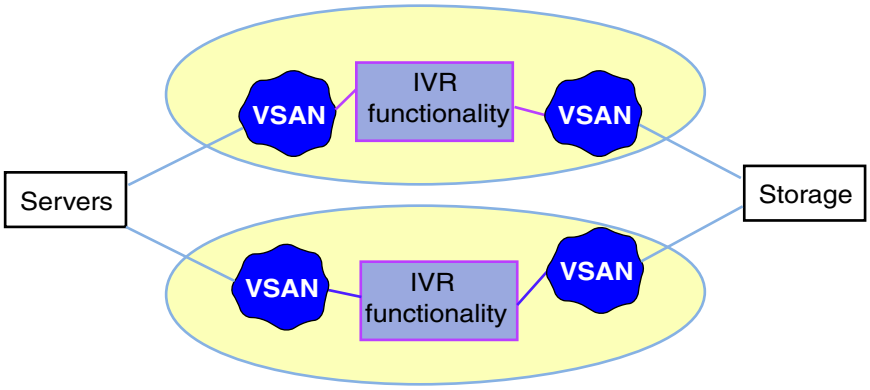

Figure 19: Dual-redundant VSAN

Supported routing configurations

Routing requires additional configuration rules for fabrics. For details about routing configuration rules, see “MP Router fabric rules” on page 90 and “C-Series switches and fabric rules” on page 97.

This section describes the following topics:

The typical configuration is a router connected to two or more fabrics, as shown in Figure 20. An MP Router is required for fabrics that include B-Series switches.

Routing and core-edge fabrics

For core-edge fabrics, connect the core switches as shown for the B-Series MP Router Meta SAN in Figure 20, or as shown for the C-Series VSAN with IVR in Figure 21.

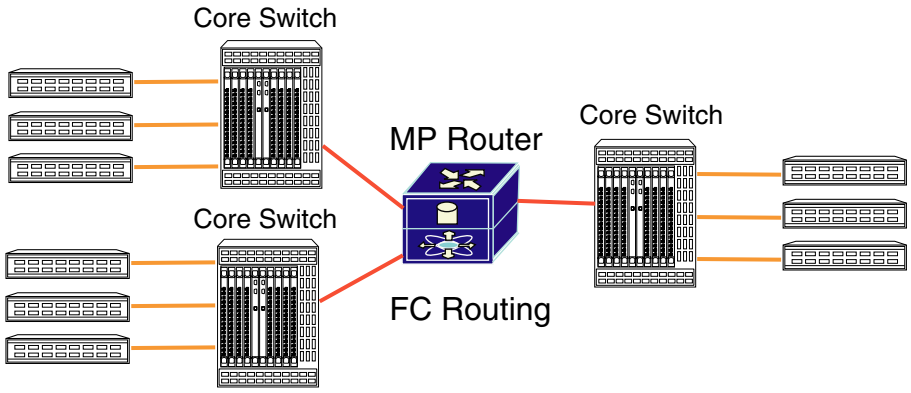

Figure 20: MP Router connecting at core switches

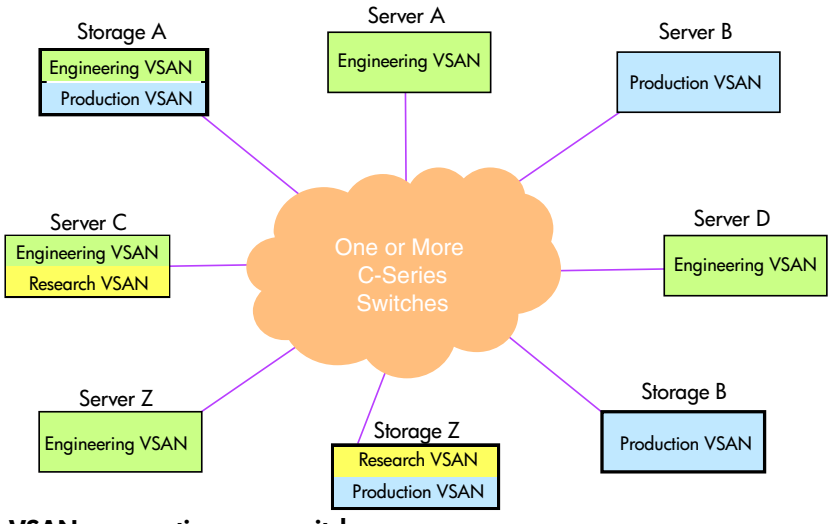

Figure 21: VSANs connecting core switches

Routing through an IP network

When connecting fabrics through IP, the MP Router can serve as an FCIP gateway for Fibre Channel routing. Routers that communicate with the FCIP protocol must be installed in pairs. (Figure 22).

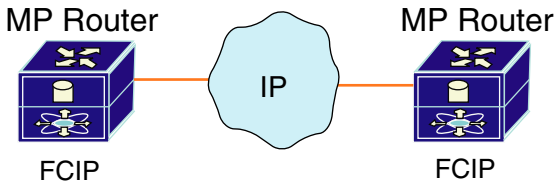

Figure 22: Routers connecting fabrics through IP

VSANs can be connected through IP using the FCIP feature of the C-Series Fibre Channel switches (see C-Series switches" on page 98).

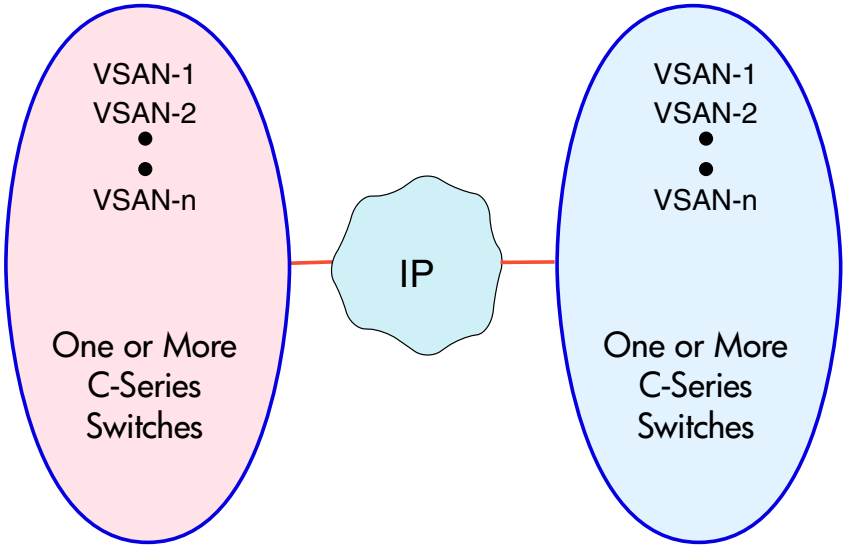

Figure 23: VSAN IVR over FCIP

High-availability MP Router configurations

In high-availability configurations, use router pairs to provide redundant paths between fabrics. Figure 24 shows valid configurations.

- The first configuration shows servers and storage connected using a pair of redundant fabrics in a level 4 NSPOF configuration. (See "Data availability" on page 59 for information about high-availability levels.)
- The second configuration shows routers cross-wired to provide full fabric connectivity in case a router fails.
- The third configuration shows independent fabrics for the servers and storage, connected using a single router. You can increase availability by adding ISLs between the MP Router and the fabrics.

For simplicity, the configurations show a small number of SAN fabrics connected to each router. The same principles apply to configurations with a higher number of fabrics connected to a router, and a higher number of routers in the Meta SANs. See "B-Series switches and fabric rules" on page 81 for scalability rules, such as the maximum number of fabrics and MP Routers.

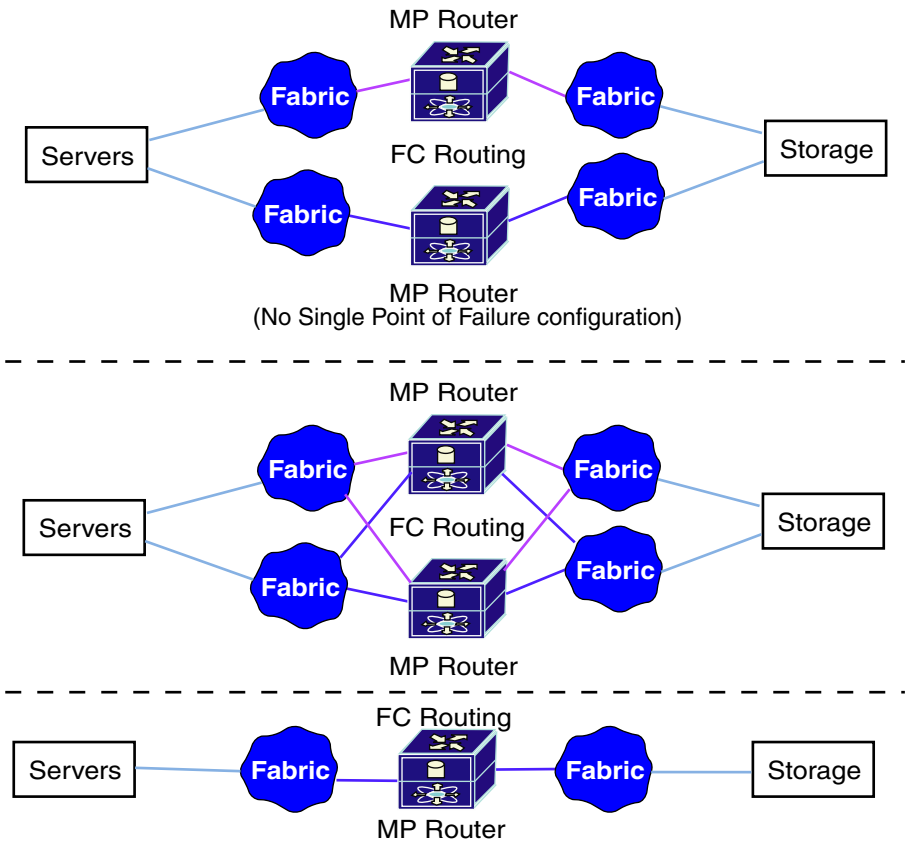

Figure 24: High-availability MP router configurations

MP Router use cases

This section describes the following topics:

- SAN island consolidation and scaling, page 76
- Integration of Fibre Channel routing and FCIP, page 77
- Tape backup consolidation, page 77

SAN island consolidation and scaling

The MP Router consolidates multiple independent fabrics or SAN islands into a Meta SAN. This modular SAN design offers:

- Simplified scalability that allows you to scale a SAN without having to merge fabrics.
- Selective sharing of devices in different fabrics so that only devices required for specific functions are seen across fabrics.
- Limited sharing or specific times for data migrations and storage consolidation.
- Ability to access equipment without changing its physical location. Connecting multiple fabrics to the MP Router enables sharing of devices located anywhere in the Meta SAN.

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

Multiple SANs can be centralized and consolidated into one Meta SAN, or partitioned into different administrative domains as required. HP recommends the use of Fabric Manager to simplify management procedures when implementing an MP Router-based Meta SAN.

Figure 25 shows a typical configuration for SAN island consolidation.

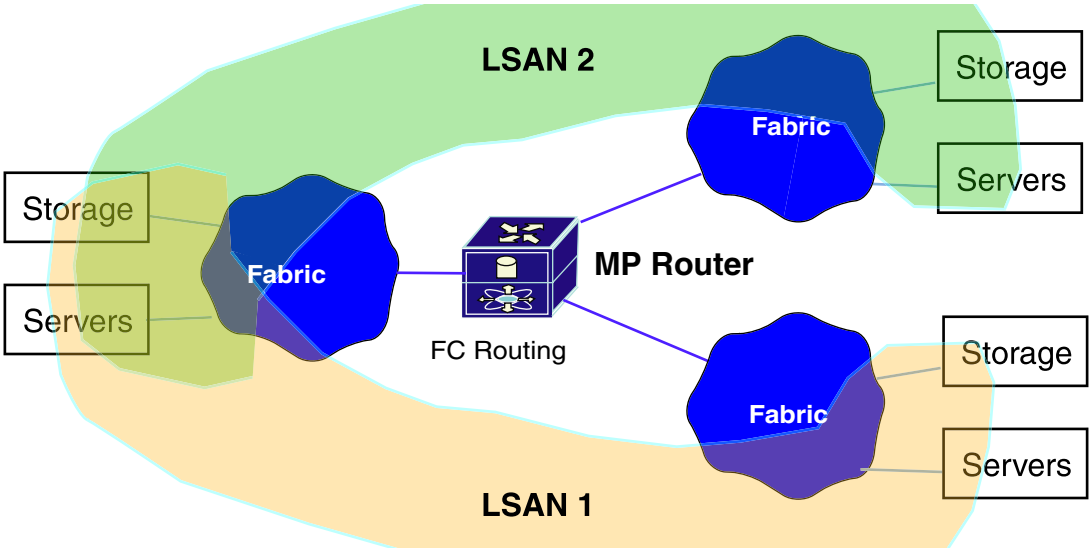

Figure 25: Consolidating SAN islands with the MP Router

Integration of Fibre Channel routing and FCIP

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

In typical Continuous Access configurations, local and remote fabrics merge when connected through FCIP. The IP connection is like an ISL in a single fabric. By using the MP Router Fibre Channel routing feature along with FCIP, the local and remote fabrics connect without merging. You can create an LSAN that contains local and remote storage arrays and servers.

Figure 26 shows a typical Continuous Access NSPOF configuration where the MP Router provides Fibre Channel routing and FCIP.

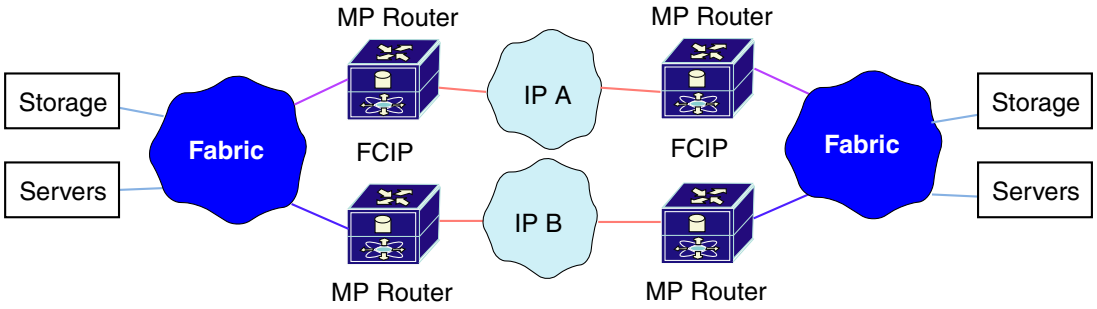

Figure 26: MP Router providing FCIP and FC routing for NSPOF configuration

Tape backup consolidation

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

Figure 27 shows a configuration where an MP Router consolidates tape backup in a SAN. See the *HP StorageWorks Enterprise Backup Solutions design guide* and the EBS compatibility matrix for information about supported HP tape products and backup applications. These documents are available at the following web sites:

http://h10025.www1.hp.com/ewfrf/wc/manualCategory?lc=en&cc=us&product=406722

lhttp://h18004.www1.hp.com/products/storageworks/ebslegacymatrices.html

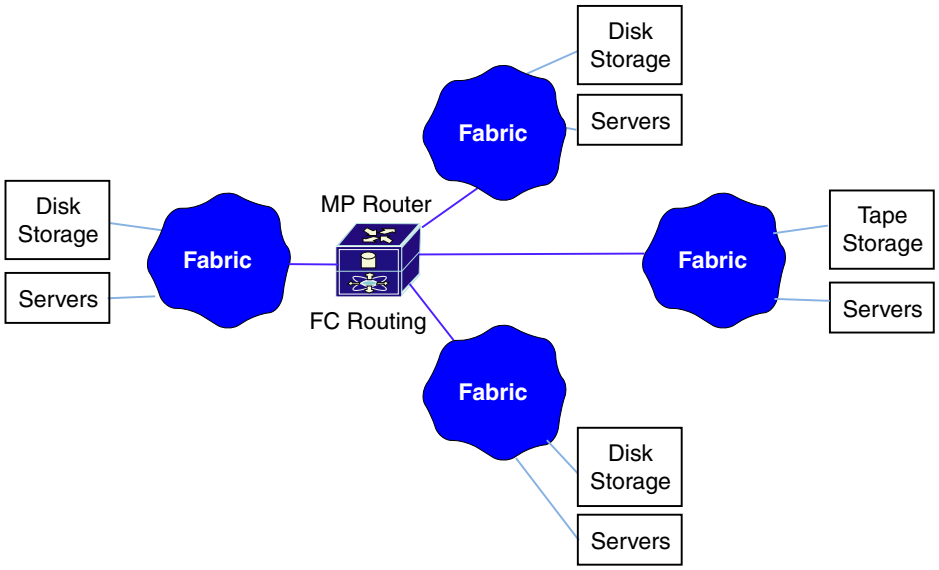

Figure 27: Tape backup consolidation

Note: Independent fabrics connected through an MP Router must not have ISL connections between the fabrics (Figure 28). A direct ISL path between fabrics bypasses the MP Router, resulting in a full fabric merge.

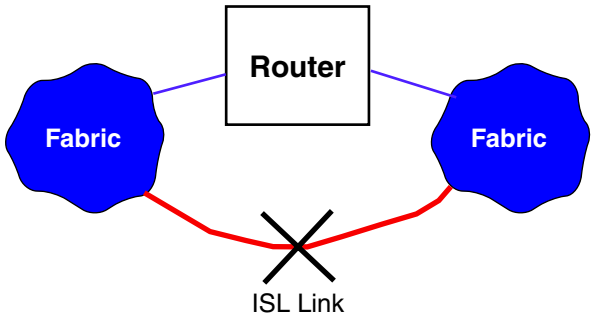

Figure 28: Unsupported configuration

Volume 2

Fabric infrastructure rules

Fabric infrastructure rules are presented in these chapters:

B-Series switches and fabric rules

4

This chapter describes the B-Series switches, Multi-protocol (MP) Router, and the fabric rules for building B-Series fabrics. It describes the following topics:

Note: See "SAN fabric connectivity and switch interoperability rules" on page 113 for information about using switches from the different series in the same SAN or the same fabric.

B-Series switches and MP Router

B-Series switches:

- Can be core or edge switches. When configured in a core-edge fabric topology, a *core switch* typically connects to other switches in the SAN; an *edge switch* typically connects to servers and storage.
- Have 8, 16, 32, 64, or 128 ports.
- Include entry-level and high-availability switches.
- Support plug-and-play compatibility.
- Support Fibre Channel routing (see “MP Router fabric rules” on page 90) and Fibre Channel over IP (FCIP) through the MP Router (see "SAN extension" on page 209) and iSCSI over IP through the MP Router (see “iSCSI storage” on page 235).
- Support iSCSI (see “iSCSI storage” on page 235) and FCIP through the SR2122-2 storage router (see "SAN extension" on page 209).

The B-Series switches offer:

- High availability
- Scalability
- Cost efficiency

This section describes the following topics:

Model numbering

The B-Series switches use the numbering scheme `x/y`:

- `x`—The highest speed at which the switch ports can operate, measured in gigabits per second (Gb/s)
- `y`—The number of switch ports

For example, the HP StorageWorks SAN Switch 2/16 is a 2 Gb/s switch with 16 ports.

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

Model naming

The B-Series switches are called HP StorageWorks SAN, Core, or Director switches. Core and Director switches are core (enterprise-class) switches. SAN switches are edge (entry-level or mid-range) switches. The suffix “-EL,” “V,” or “N” identifies entry-level switches.

The B-Series embedded switch for HP p-Class BladeSystem servers is called the Brocade 4Gb SAN Switch. The B-Series router is called the HP StorageWorks Multi-protocol (MP) Router.

Switch models

Table 10 describes the B-Series switches supported by HP (1 Gb/s, 2 Gb/s, and 4 Gb/s B-Series Fibre Channel switches). Table 10 includes HP and Compaq legacy switches.

HP supports all B-Series switches in a fabric if you:

- Use the firmware versions listed.
- Follow the fabric rules. (See "Fabric rules" on page 87.)

Table 10: B-Series switches and MP Router

HP StorageWorks switch name	Firmware version	Fabric Manager version	Number of ports
Brocade 4Gb SAN Switch for HP p-Class BladeSystem	5.00	4.4.0	8 internal, 4 external
HP StorageWorks SAN Switch 2/8V TAA, 2/8V TAA power pack	4.4.0c	4.1.1	8
HP StorageWorks SAN Switch 2/16V, 2/16V TAA, 2/16N FF, 2/16N power pack, 2/16N FF TAA, 2/16N TAA power pack			16
HP StorageWorks SAN Switch 2/32, 2/32 power pack			32
HP StorageWorks Core Switch 2/64, 2/64 power pack			64 (2 switches / chassis, total 128 ports / chassis)
HP StorageWorks SAN Director 2/128, 2/128 power pack			128
HP StorageWorks SAN Switch 4/32 base/full		4.4.0	16/32
HP StorageWorks SAN Switch 4/32 power pack			32
HP StorageWorks MSA SAN Switch 2/8	3.2.0	3.0.2c	8
HP StorageWorks SAN Switch 2/8 EL, 2/8 power pack			8
HP StorageWorks SAN Switch 2/16, 2/16 EL, 2/16 power pack			16
HP StorageWorks Multi-Protocol Router	7.3.0b	4.4.0	8/16 See Table 18 for scalability rules

Legacy HP switch name	Legacy Compaq StorageWorks switch name			Number of ports
HP Brocade 2400 (HP reseller)	Compaq StorageWorks SAN Switch 8	2.6.2c	3.0.2c	8
N/A	Compaq StorageWorks SAN Switch 8-EL			8
HP Brocade 2800 (HP reseller)	Compaq StorageWorks SAN Switch 16			16
N/A	Compaq StorageWorks SAN Switch 16-EL			16
HP Surestore FC Switch 6164 (64 ISL Ports)	Compaq StorageWorks SAN Switch Integrated/32 (64 ISL Ports)			32 (counts as 6 switches and 2 hops when applying configuration rules)
HP Surestore FC Switch 6164 (32 ISL Ports)	Compaq StorageWorks SAN Switch Integrated/64 (32 ISL Ports)			64 (counts as 6 switches and 2 hops when applying configuration rules)

Table 10: B-Series switches and MP Router (Continued)

HP StorageWorks switch name		Firmware version	Fabric Manager version	Number of ports
HP Surestore FC 1Gb/2Gb Entry Switch 8B	N/A	3.2.0	3.0.2c	8
N/A	Compaq StorageWorks SAN Switch 2/8-EL			8
N/A	Compaq StorageWorks SAN Switch 2/16-EL			16
HP Surestore FC 1Gb/2Gb Switch 8B	N/A			8
HP Surestore FC 1Gb/2Gb Switch 16B	Compaq StorageWorks SAN Switch 2/16			16

For the latest information on supported B-Series switches and firmware versions, see the HP storage web site: http://h18006.www1.hp.com/storage/saninfrastructure.html.

Features

Features of the B-Series switches are:

- **Advanced Performance Monitor**—Analyzes resource utilization throughout the fabric.
- **Advanced WebTools**—Centralizes and simplifies switch management through a browser-based application.
- **Advanced Zoning—**Provides secure access control over fabric resources. Uses the switch hardware to enforce port and WWN zoning.
- **Extended Fabrics**—Enables Fibre Channel SAN connectivity up to 100 km to improve disaster recovery operations and ensure business continuity.
- **Fabric Manager**—Centralizes fabric management through a host-based application.
- **Fabric Watch**—Proactively monitors the health and performance of switches and the fabric.
- **ISL Trunking**—Combines multiple links between switches to form a single, logical interswitch link (ISL) with a total bandwidth of 32 Gb/s. Enables dynamic load balancing of data across ISLs.
- **Remote Switch**—Creates one logical SAN that spans remote fabrics at unlimited distances. All SAN components appear as local devices.
- **Secure Fabric OS**—Provides flexible security and policy-based administration to protect data from unauthorized access and corruption.

Table 11 provides a comparison of the high-availability features for B-Series switches.

Table 11: B-Series switch high-availability feature comparison

Model	Redundant/ hot-swappable power	Redundant/ hot-swappable cooling	Redundant control processor	Nondisruptive code activation	Nondisruptive port expansion	Redundant active components
Brocade 4 Gb SAN Switch for HP p-Class BladeSystem	N/A	N/A	No	Yes	No	No
SAN Switch 8, 16, 2/8, 2/8-EL	No / No	Yes / No	No	No	No	No
SAN Switch 2/8V, 2/16V, 2/16N, 2/16N FF	Yes /No	Yes /No	No	Yes	No	No
SAN Switch 2/16	Yes / Yes	Yes / Yes	No	No	No	No
SAN Switch 2/32, 4/32	Yes / Yes	Yes / Yes	No	Yes	No	No
Core Switch 2/64	Yes / Yes	Yes / Yes	Yes	Yes	Yes	Yes
SAN Director 2/128	Yes / Yes	Yes / Yes	Yes	Yes	Yes	Yes
MP Router	Yes /Yes	Yes /Yes	No	No	Yes	No

Usage

Table 12 describes the use of B-Series switches as core switches.

Table 12: Using B-Series switches as core switches

Model	1-96 user ports	97-224 user ports	225-500 user ports	501-728 user ports	729-1280 user ports
Core Switch 2/64 Core Switch 2/128	Excellent	Excellent	Excellent	Excellent	Excellent
SAN Switch 2/32, 4/32	Excellent	Very good	Good	Not recommended	Not recommended
SAN Switch 2/8, 2/8-EL, 2/8V, 2/16, 2/16V, 2/16N, 2/16N FF, Brocade 4 Gb SAN Switch for HP p-Class BladeSystem	Good	Good	Good	Not recommended	Not recommended
SAN Switch 8, 16	Good	Good	Not recommended	Not recommended	Not supported

Note: HP does not recommend using SAN Switch 8 or SAN Switch 16 as a core switch if it connects to other 1 Gb/s SAN switches.

Table 13 describes the use of B-Series switches as edge switches.

Table 13: Using B-Series switches as edge switches

Model	1-96 user ports	97-224 user ports	225-500 user ports	501-728 user ports	729-1280 user ports
Core Switch 2/64, Core Switch 2/128	Excellent	Excellent	Excellent	Excellent	Excellent
SAN Switch 2/32, 4/32	Excellent	Excellent	Excellent	Excellent	Excellent
SAN Switch 2/8, 2/8-EL, 2/8V, 2/16, 2/16V, 2/16N, 2/16N FF, Brocade 4 Gb SAN Switch for HP p-Class BladeSystem	Excellent	Very good	Very good	Good	Good
SAN Switch 8, 16	Very good	Very good	Good	Good	Not supported

Fabric rules

This section describes the fabric rules for the B-Series switches, the MP Router, and other factors you should consider when building B-Series fabrics.

When using B-Series switches in a fabric, consider the following:

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

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

This section describes the following topics:

Operating systems and storage models

The fabric rules for B-Series switches and the MP Router apply to SANs that include the following operating systems and storage models:

Operating systems:	Storage models:
■ HP-UX	■ XP12000
■ OpenVMS	■ XP128/1024
■ Tru64 UNIX	■ XP48XP48/512/256
■ IBM AIX	■ VA7100/7110/7400/7410
■ Linux	■ EVA3000/4000/5000/6000/8000
■ Microsoft Windows	■ EMA/ESA12000
■ Novell NetWare	■ EMA16000
■ Sun Solaris	■ MA/RA8000
■ VMware ESX	■ MA6000
	■ MSA1000/1500
	■ RA4000/4100

See "Heterogeneous server rules" on page 127, and “SAN storage system rules” on page 171 to determine operating system support for each storage model.

Fabric rules for B-Series switches

The following fabric rules apply to all B-Series SANs. They also apply, in general, to Continuous Access XP, Continuous Access EVA, and Data Replication Manager (DRM) configurations. However, additional rules apply to Continuous Access and DRM implementations. See the *HP StorageWorks Continuous Access EVA planning guide* for detailed information.

Table 14 lists the rules for creating a SAN with B-Series switches.

Table 14: B-Series fabric rules

Rule number	Description
1	A maximum of 55 switches, 2,560 total ports, and 2,300 user ports in a single fabric with all switches using 4.4x or later firmware.
2	A maximum of 34 switches and 1,280 total ports in a single fabric with switches using 2.6x, 3.x, and 4x firmware.
3	A fabric that contains 2 Gb/s switches only, or 2 Gb/s and 4 Gb/s switches only, supports a maximum of 1,200 user ports. With security enabled, the fabric supports a maximum of 728 user ports.
4	A fabric that contains 1 Gb/s switches using 2.6.1x (or later) firmware supports a maximum of 728 user ports. With security enabled, the fabric supports a maximum of 500 user ports.
5	Brocade 4Gb SAN Switch for HP p-Class BladeSystem and HP StorageWorks SAN Switch 4/32—4 Gb/s port speed supported only for ISL connectivity between these switch models.
6	HP StorageWorks SAN Director 2/128—Maximum of 20 chassis per fabric. Each chassis adds one switch to the fabric switch count if configured as one domain.
7	HP StorageWorks Core Switch 2/64—Maximum of 20 chassis per fabric. Each chassis contains two logical switches, adding two switches to the fabric switch count.
8	In a fabric that contains the Brocade 4Gb SAN Switch for HP p-Class BladeSystem, HP StorageWorks SAN Switch 4/32 SAN Switch 4/32, 2/32, Core Switch 2/64, or SAN Director 2/128, the Core switch addressing mode is required on all other switches in that fabric. (See "Core switch addressing mode" on page 94.)
9	Within a fabric, assign a unique domain number (domain ID) and a unique world wide name (WWN) to each switch. All switch configuration parameters for like switches must be the same. Do not configure any switches with a domain ID of 8, which is reserved for HP-UX.
10	All switches in a single fabric or multi-fabric SAN must use the same firmware version for all switches in the same firmware family (for example, 2.x, 3.x, 4.x, 5.x). Two successive switch firmware versions can be used temporarily in one fabric or multiple fabrics when updating switch firmware.
11	StorageWorks SAN Switch Integrated 32 or 64 and HP Surestore FC Switch 6164 support a maximum of four chassis per fabric. Each chassis adds six switches to the fabric switch count. The maximum fabric configuration is four chassis with ten additional SAN switches.
12	Maximum of seven hops (eight switches) between any two communicating devices. The SAN Switch Integrated 32 or 64 and the HP Surestore FC Switch 6164 add a maximum of two hops between devices, depending on the device-to-switch connections and device-to-device access.

Table 14: B-Series fabric rules (Continued)

Rule number	Description
13	StorageWorks SAN Switch 2/8-EL and 2/16-EL support a maximum of four switches in a fabric by default. A license upgrade is available to support fabrics with more than four switches.
14	StorageWorks SAN Switch 2/8-V and 2/16-V support a maximum of two switches in a fabric by default. A license upgrade is available to support fabrics with more than two switches.
15	HP Surestore FC 1Gb/2Gb Entry Switch 8B and Compaq StorageWorks SAN Switch 8-EL support one E-port connection. (See "SAN fabric connectivity rules" on page 114 for more information.)
16	Compaq Fibre Channel Storage Switch 8 and Storage Switch 16 support a maximum of four switches in fabrics with these switches only, or in fabrics with these switches and any 1 Gb/s SAN switches. VC Encoded Address Mode must be set for 1 Gb/s SAN switches (see the SAN switch documentation). These switches are not supported in fabrics with 2 Gb/s or 4 Gb/s switches.
17	Compaq FC-AL Switch 8 supports cascaded attachment in a fabric using an FL-port on a Compaq SAN Switch 8, SAN Switch 16, SAN Switch 8-EL, or SAN Switch 16-EL. In this configuration, HP supports only servers directly attached to the FC-AL switch with access to RA4000/4100 storage systems. HP does not support attachment to 2 Gb/s or 4 Gb/s switches.

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

Switch database size

Table 15 describes the database size rules for B-Series switches in a fabric.

Table 15: Database size rules for B-Series switches

Rule number	Description
1	Secure Fabric OS—With security enabled in a fabric, the maximum security database size is as follows: ■ If the fabric contains 1 Gb/s and 2 Gb/s switches, the maximum size is 32 KB, with only 16 KB active. The maximum number of Device Connection Control (DCC) policies is 620. ■ If the fabric contains 2 Gb/s switches only, the maximum size is 128 KB, with 64 KB active.
2	Advanced Zoning—The maximum zoning database size is as follows: ■ If the fabric contains 1 Gb/s switches with 2.6.1x (or later) firmware, the database must not exceed 96 KB. ■ If the fabric contains 2 Gb/s or 4 Gb/s switches with 3.3.x/4.4.x (or later) firmware, the database must not exceed 256 KB. **Note:** Use the cfgSize command to determine the size of the zoning database.

ISL maximums

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

Note: Some switches require licensing for additional ISL ports.

MP Router fabric rules

This section describes the fabric rules for the MP Router and other factors you should consider when building B-Series fabrics that contain MP Routers. The fabric rules for the MP Router apply to SANs that include the same operating system and storage models as the B-Series switches (see "Operating systems and storage models" on page 87).

Table 16 describes the rules for creating fabrics with MP Routers.

Table 16: MP Router fabric rules

Rule number	Description
1	All configurations must use the default settings for R_A_TOV (10000 seconds) and E_D_TOV (2000 seconds).
2	Devices connected directly to the MP Router or the Backbone Fabric are for iSCSI access only and cannot be routed via an LSAN zone to an Edge Fabric (Figure 29).
3	Devices connected to Edge Fabrics are for local and routed access only and cannot be routed via an LSAN to an iSCSI Port on the MP Router (Figure 30).
4	SAN boot through the MP Router is not supported.
5	For redundancy, HP recommends using a minimum of 2 EX-Ports (MP Router ISLs) for each fabric connected to the MP Router. The number of EX-Ports required is based on the performance requirements for all devices shared by fabrics through the MP Router. ■ For initial settings, calculate the number of EX-Ports using a 5:1 ratio of standard ISLs or switch device ports to MP Router EX-Port ISLs. Monitor port performance to determine if this ratio is acceptable based on usage (use the portperfshow command to monitor MP Router port performance). ■ For very high bandwidth requirements and applications, use a ratio of 3:1 or 1:1.
6	Scalability rules (see "Scalability rules" on page 92.)

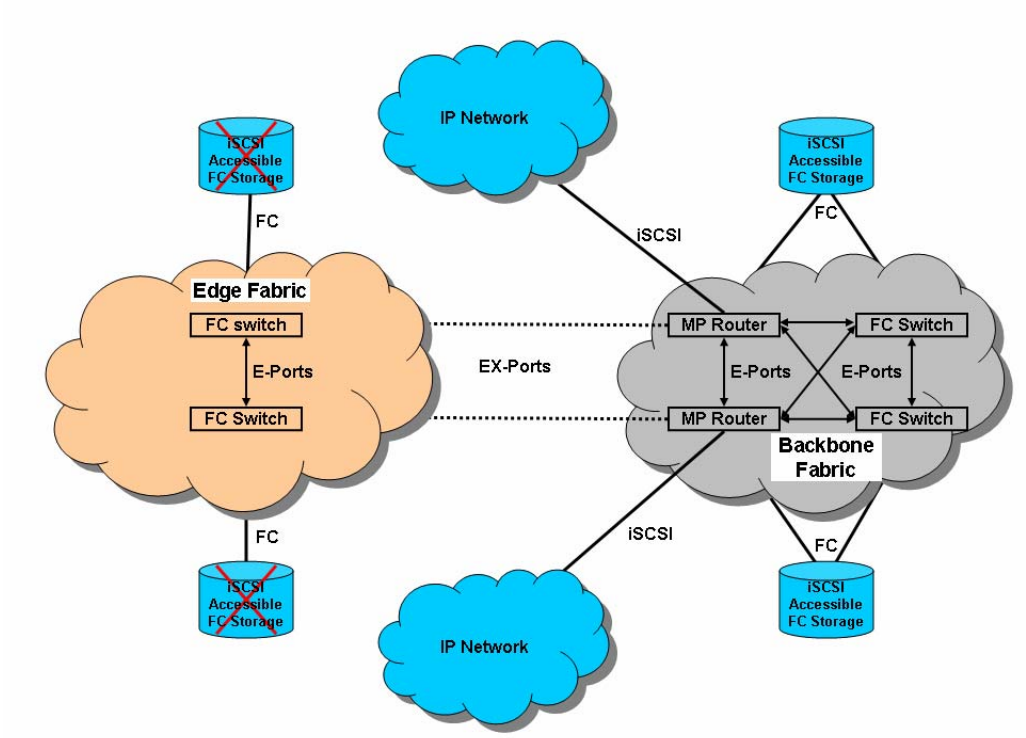

Figure 29: iSCSI storage device connectivity for the MP Router

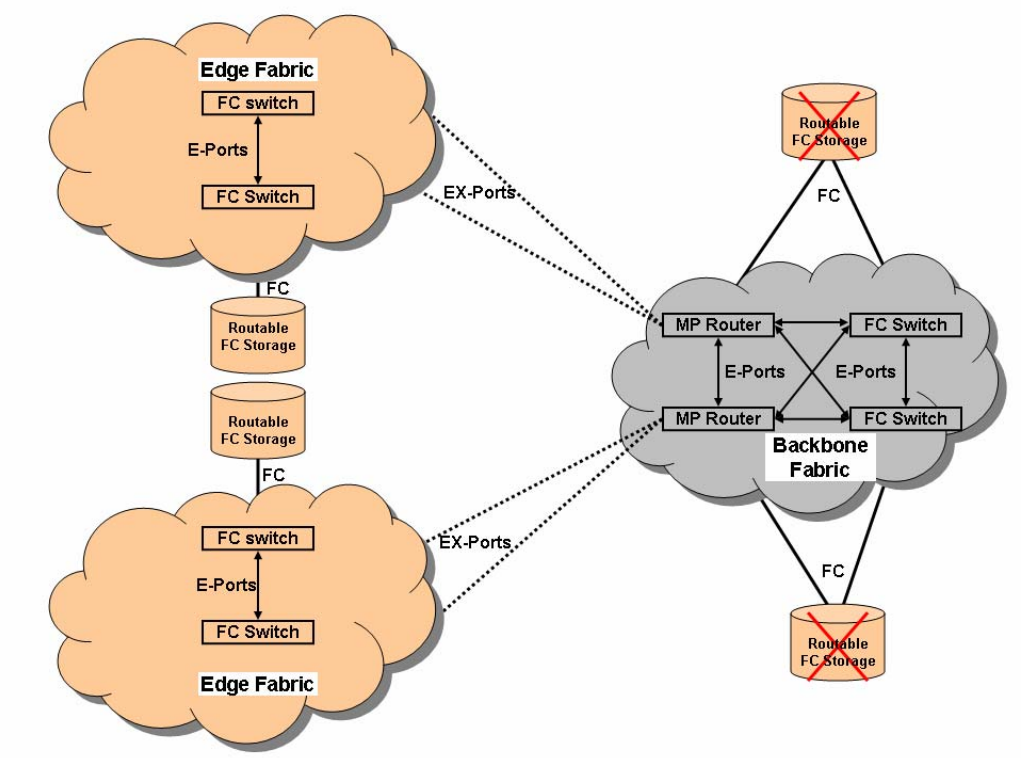

Figure 30: FC routing storage device connectivity for the MP Router

The remainder of this section describes additional considerations and rules for using the MP Router in a fabric:

- XPath OS compatibility, page 92
- Scalability rules, page 92

XPath OS compatibility

Table 17 lists the B-Series switches and firmware versions that are supported in a SAN with the MP Router XPath OS.

Table 17: B-Series switches and firmware supported with XPath OS v7.3.0b

	1 Gb	2 Gb 2/8, 2/8-EL, 2/16, 2/16-EL, Surestore 8B, 16B	2 Gb, 2/8V, 2/16V, 2/16N, 2/32, 2/64, 2/128	4 Gb 4/32	Fabric Manager[1]
HP recommended versions	2.6.2c	3.2.0	4.4.0c	4.4.0c	4.4.0
HP minimum versions	2.6.2b	3.1.2b	4.4.0b	4.4.0b	4.2.x 4.4.0 (with 4/32)

1. HP recommends that you use Fabric Manager in all configurations using the MP Router.

Scalability rules

Table 18 lists the scalability rules for Meta SANs using the MP Router. These scalability rules are for using the MP Router in fabrics built using any of the B-Series switches listed in Table 10.

The following terms describe MP Router scalability:

- **Front phantom domains**—Individual EX-port connections from the MP Router to the edge fabric(s). There can be 24 EX-ports to each fabric.
- **Translate phantom domains**—Assigned to each fabric connected to the MP Router. There is one unique translate domain for each fabric, regardless of how many EX-ports are connected to the fabric. There can be up to 33 translate domains for the router.
- **Backbone fabric**—The MP Router backbone fabric consists of one MP Router at a minimum and includes all MP Routers and B-Series switches connected to the MP-Router via E-ports.
- **Edge fabric**—A fabric that is attached to an EX-Port on an MP Router.

Table 18: MP Router scalability rules

Edge fabric scalability		
Domains (see Figure 32)	Maximum number of front domains (number of EX-ports) per edge fabric	24
	Maximum number of translate domains (number of remote fabrics) per edge fabric	33
	Maximum number of real domains in an edge fabric	34
	Maximum number of domains per edge fabric (real domains + front domains + translate domains) Note: The total number of domains in the first three rows must not exceed 80.	80

Table 18: MP Router scalability rules (Continued)

Devices (user ports)	Maximum number of local and remote devices per edge fabric (For edge fabrics with more than 600 devices, v4.2.0c or later is required for all switches in the fabric when using the MP Router.)	1280
	Maximum number of imported devices per edge fabric	1000
LSAN zone scalability		
Zoning	Maximum number of entries per LSAN zone	200
Routed fabrics (Meta SAN) scalability		
Edge fabrics	Maximum number of edge fabrics connected to a routed fabric	34
Meta SAN	Maximum number of total ports per routed fabric	12800
Routers	Maximum number of MP Routers per backbone fabric	7
LSAN	Maximum number of LSAN device entries per routed fabric	9000
	Maximum number of LSAN zones per routed fabric	1000
Hop count scalability		
Hop count	Maximum number of hops between switches (including routers) in a Meta SAN (Figure 31)	12

MP Router hop count

The MP Router is counted just as a Fibre Channel switch is counted with respect to fabric hop count. Devices communicating across fabrics through the MP Router must adhere to both the B-Series seven-hop limit and the MP Router twelve-hop limit (Figure 31).

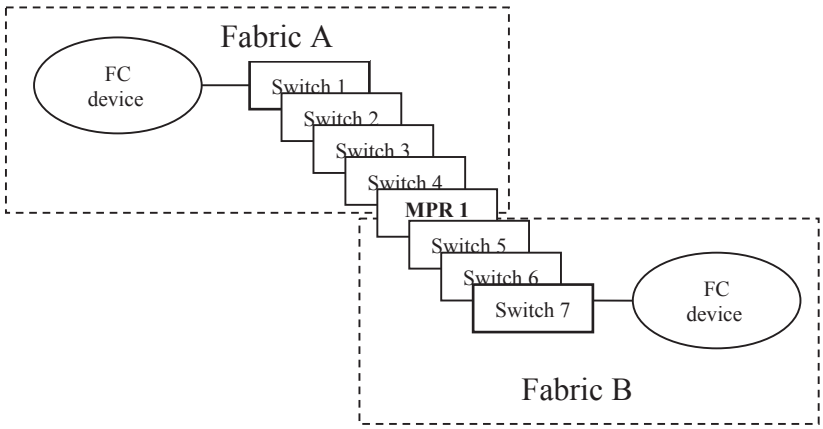

Figure 31: MP Router with seven hops

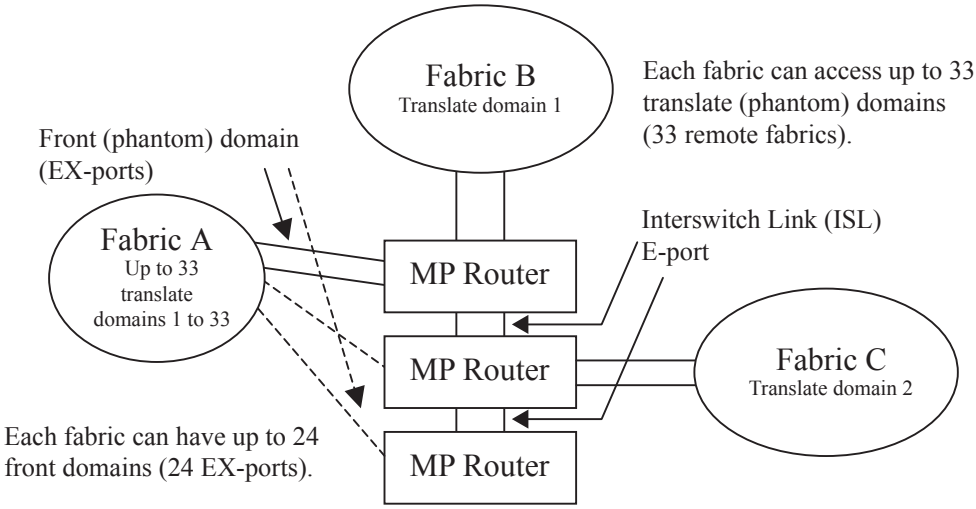

Figure 32: Front and translate domains

MP Router backbone fabric

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

Core switch addressing mode

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

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

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

For more information about Core switch addressing mode, see the HP storage web site: http://h18006.www1.hp.com/storage/saninfrastructure.html.

From this web site, select a B-Series switch, and then select **Technical documents**.

Zoning limits and enforcement

Table 19 describes zone enforcement for B-Series switches.

Table 19: Zone enforcement for B-Series switches and MP Router LSANs

Switches	Configuration	Enforcement	Comments
Compaq StorageWorks SAN Switch-8, SAN Switch 8-EL;	Define zones using domain number, port number	Access authorization at frame level in hardware	Hard zoning
SAN Switch-16 SAN Switch 16-EL SAN Switch Integrated/32 SAN Switch Integrated/64 (-6164)	Define zones using WWNs only	Discovery authentication Name Servers (NS) directory-based	Soft zoning
	Define zones using combination of domain/port numbers and WWNs	Discovery-based authentication	Soft zoning
HP StorageWorks SAN Switch 2/8-EL, 2/8V, 2/16, 2/16-EL, 2/16V, 2/16N, 2/16N FF	Define zones using domain number, port number	Access authorization at frame level in hardware	Hard zoning
HP reseller FC-8B FC-16B	Define zones using WWNs only	Access authorization at frame level in hardware	Hard zoning
HP StorageWorks SAN Director 2/128 Core Switch 2/64 Brocade 4Gb SAN Switch for HP p-Class BladeSystem SAN Switch 2/32, 4/32	Define zones using combination of domain/port numbers and WWNs	Name service plus login authentication	Soft zoning, NS authentication, and login protection
Switches that support QuickLoop	Define zones using ALPAs, domain/port numbers, or combination thereof	Implemented in hardware tables, access prevented by hardware between unauthorized devices	Hard zoning

C-Series switches and fabric rules

5

This chapter describes the C-Series switches and the fabric rules for building C-Series fabrics. It describes the following topics:

Note: See "SAN fabric connectivity and switch interoperability rules" on page 113 for information about using switches from the different series in the same SAN or fabric.

C-Series switches

C-Series switches:

- Can be core or edge switches. When configured in a core-edge topology, a *core switch* typically connects to other switches in the SAN; an *edge switch* typically connects to servers and storage.
- MDS 9506 and MDS 9509 Director switches can accommodate 16-port and 32-port Fibre Channel modules, and IPS-4, IPS-8, and 14/2 services modules.
  - The 16-port Fibre Channel module is recommended for enterprise-level host connections, storage connections, and ISL connections.
  - The 32-port Fibre Channel module uses 3.2:1 internal oversubscription and is recommended for low- to mid-range host connections and tape device connections.
  - The IP services modules provide MDS iSCSI and FCIP functionality:
    - IPS-4 and IPS-8 provide 4 and 8 Gigabit Ethernet (GE) IP ports, respectively.
    - 14/2 provides 14 Fibre Channel ports and 2 GE IP ports, respectively.
- MDS 9216/9216A switches have 2 slots:
  - One slot is a fixed configuration with a 16-port Fibre Channel module.
  - The second slot can accommodate a 16-port or 32-port Fibre Channel module, or a 14/2, IPS-4, or IPS-8 module for iSCSI and FCIP support.
- The MDS 9216i switch has 2 slots:
  - One slot is a fixed configuration with a 14/2 module.
  - The second slot can accommodate a 16-port or 32-port Fibre Channel module, or a 14/2, IPS-4, or IPS-8 for iSCSI and FCIP support.
- MDS 9120/9140 switches have fixed configurations with 20 and 40 Fibre Channel ports, respectively:
  - MDS 9120 has 4 full-rate Fibre Channel ports and 16 oversubscribed ports.
  - MDS 9140 has 8 full-rate Fibre Channel ports and 32 oversubscribed ports.
- All C-Series Fibre Channel switches are also supported with iSCSI and FCIP through the SR2122-2 storage router. (See "SAN extension" on page 209 for more information.)

The C-Series switches offer:

- High availability
- Scalability
- Cost efficiency

This section describes the following topics:

Model naming

The C-Series switches are named MDS 9*xnn*. The 95*nn* switches (Multilayer Directors) are director or core switches. The *nn* value indicates the number of slots available for supervisors and port modules. The 92*nn* switches are mid-range switches; the 91*nn* switches are entry-level switches (Multilayer Fabric switches). The *nn* value indicates the number of fixed ports.

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

Switch models

Table 20 describes the C-Series switches. HP supports all C-Series switches in a fabric if you:

- Use the firmware versions listed.
- Do not exceed the maximum number of Fibre Channel ports listed.
- Follow the fabric rules. (See "Fabric rules" on page 102.)

Table 20: C-Series switches

Switch	SAN-OS	Maximum number of Fibre Channel ports
MDS 9506 Director	2.0(1b), 2.1(1a)	128
MDS 9509 Multilayer Director		224
MDS 9216/9216A Multilayer Fabric		48
MDS 9216i Multilayer Fabric		46
MDS 9120 Fabric		20
MDS 9140 Fabric		40

For the latest information on supported C-Series switches and firmware versions, see the HP storage web site: http://h18006.www1.hp.com/storage/saninfrastructure.html.

Features

Features of the C-Series switches are:

- Nonblocking architecture using virtual output queuing (VOQ)
- Virtual storage area network (VSAN) deployment over the physical infrastructure (VSANs are separate instances of all fabric services, including address space)
- Advanced diagnostics and troubleshooting (FC Ping, FC Traceroute, SPAN, RSPAN, and Call Home)
- Comprehensive security (SSH, SFTP, RADIUS, SNMPv3, and RBAC)
- Comprehensive fabric management (CLI, SNMP, and Java-based GUI)
- Traffic management (FCC and QoS)
- High-availability, fault-tolerant software
- PortChannel (ISL aggregation for highly resilient SAN architectures)
- Integrated multiprotocol capability (MDS 95*nn* and 92*nn*) for SAN extension—FCIP and iSCSI

Additional features of the Director switches are:

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

Table 21 provides a comparison of the high-availability features for C-Series switches.

Table 21: C-Series switch high-availability feature comparison

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

Usage

Table 22 describes the use of C-Series switches as core switches.

Table 22: Using C-Series switches as core switches

Model	1–48 total ports	49–224 total ports	225-512 total ports
MDS 9140 Fabric	Excellent (40 port maximum)	Not recommended	Not recommended
MDS 9120 Fabric	Excellent (20 port maximum)	Not recommended	Not recommended
MDS 9216/9216A/9216i Multilayer Fabric	Excellent up to 48 ports (46 ports for 9216i)	Not recommended	Not recommended
MDS 9506 Director	Good	Excellent (128 port maximum)	Excellent
MDS 9509 Multilayer Director	Good	Excellent	Excellent

Table 23 describes the use of C-Series switches as edge switches.

Table 23: Using C-Series switches as edge switches

Model	1–48 total ports	49–224 total ports	225-512 total ports
MDS 9140 Fabric	Excellent	Excellent	Excellent
MDS 9120 Fabric	Excellent	Excellent	Excellent
MDS 9216/9216A/9216i Multilayer Fabric	Excellent	Excellent	Very good
MDS 9506 Director	Excellent (select modules to optimize performance or user port count)		
MDS 9509 Multilayer Director	Excellent (select modules to optimize performance or user port count)		

Fabric rules

This section describes the fabric rules for C-Series switches and other factors you should consider when building C-Series fabrics. It describes the following topics:

Operating systems and storage models

The fabric rules for C-Series switches apply to SANs that include the following operating systems and storage models:

Operating systems:

- HP-UX
- OpenVMS
- Tru64 UNIX
- IBM AIX
- Linux
- Microsoft Windows
- Novell NetWare
- Sun Solaris
- VMware ESX

Storage models:

- XP12000
- XP128/1024
- XP48XP48/512/256
- VA7100/7110/7400/7410
- EVA3000/4000/5000/6000/8000
- EMA/ESA12000
- EMA16000
- MA/RA8000
- MA6000
- MSA1000/1500

See "Heterogeneous server rules" on page 127, and "SAN storage system rules" on page 171 to determine operating system support for each storage model.

Fabric rules for C-Series switches

The following fabric rules apply to all C-Series SANs. They also apply, in general, to Continuous Access XP, Continuous Access EVA, and Data Replication Manager (DRM) configurations. However, additional rules apply to Continuous Access and DRM implementations. See the *HP StorageWorks Continuous Access EVA planning guide* for detailed information.

Table 24 lists the rules for creating a SAN with C-Series switches.

Table 24: C-Series fabric rules

Rule number	Description
1	Up to 40 MDS switches with up to 4,000 total ports and 3,500 user ports in a fabric
2	MDS 9506 switch supports up to 128 ports over four modular chassis (four 32-port modules). (See "C-Series switches" on page 98 for information about the 32-port module.)
3	MDS 9509 switch supports up to 224 ports over seven modular chassis (seven 32-port modules). (See "C-Series switches" on page 98 for information about the 32-port module.)
4	Maximum of seven switch hops (eight switches) between any two communicating devices
5	Maximum of 80 VSANs per fabric

Table 24: C-Series fabric rules (Continued)

Rule number	Description
6	Maximum of 2,000 IVR shared devices per routed fabric
7	Maximum of 2,000 IVR zones per routed fabric

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

ISL maximums

You can use all full-rate ports on all C-Series switches for ISLs, with a maximum of one half of the total ISL port count configured to the same destination switch. Table 25 lists the ISL maximums for switches with higher port counts.

Table 25: ISL maximums

Switch	Total number of available user ports	Number of ports allowed as ISLs
MDS 9000 32 Port Fibre Channel Module	32	8 Three user ports are disabled for each port configured as ISL.
MDS 9000 16 Port Fibre Channel Module	16	16
MDS 9000 14/2	14	14
MDS 9120	20	4 For each ISL port configured in addition to the initial four, three user ports must be disabled.
MDS 9140	40	8 For each ISL port configured in addition to the initial eight, three user ports must be disabled.

Zoning limits and enforcement

Table 26 lists the zoning limits for C-Series switches.

Table 26: Zoning limits for C-Series switches

Rule number	Description
1	Maximum number of zones for a fabric with VSANs is 2,048.
2	Maximum number of zone members for a fabric with VSANs is 20,000.

Table 27 describes zone enforcement for C-Series switches.

Table 27: Zone enforcement for C-Series switches

Switch	Configuration	Enforcement	Comments
MDS 9506 MDS 9509 MDS 9216/9216A/9216i MDS 9120 MDS 9140	Domain ID and port number WWNs only Domain ID and port number with WWNs	Access authorization at frame level in hardware	Hard zoning

C-Series VSAN high availability

Figure 33 shows a typical high-availability configuration with server and storage connections to different fabrics. It provides two paths for data access between servers and storage. The addition of an ISL between Fabric A and Fabric B enables you to manage both fabrics through a single management station. This configuration is classified as a level 3 high-availability configuration. See "Data availability" on page 59 for information about SAN data availability levels.

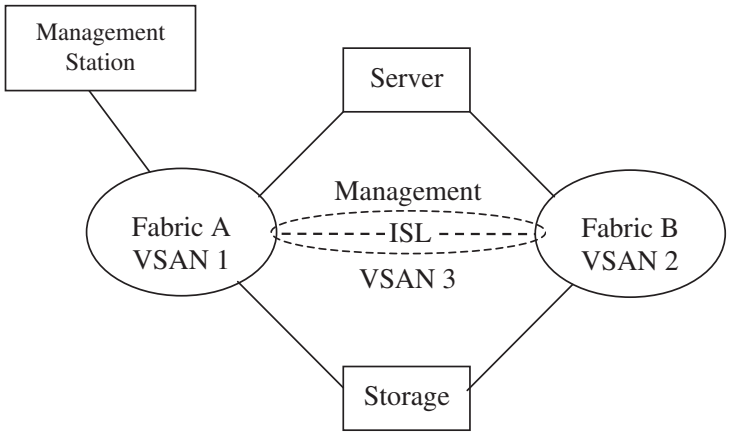

Figure 33: C-Series high-availability VSAN management configuration

M-Series switches and fabric rules

6

This chapter describes the M-Series switches and the fabric rules for building M-Series fabrics. It describes the following topics:

Note: See "SAN fabric connectivity and switch interoperability rules" on page 113 for information about using switches from the different series in the same SAN or the same fabric.

M-Series switches

M-Series switches:

- Can be core or edge switches. When configured in a core-edge fabric topology, a *core switch* typically connects to other switches in the SAN; an *edge switch* typically connects to servers and storage.
- Have 4, 8, 12, 24, 32, 64, or 140 ports.
- Include entry-level and high-availability switches.
- Use the same level of internal firmware.
- Support plug-and-play compatibility.
- Support iSCSI and Fibre Channel over IP (FCIP) through the SR2122-2 storage router. (See "SAN extension" on page 209 for more information.)

The M-Series switches offer:

- High availability
- Scalability
- Cost efficiency

This section describes the following topics:

Model numbering

The M-Series switches use the numbering scheme `x/y`:

- `x`—The highest speed at which the switch ports can operate, measured in gigabits per second (Gb/s)
- `y`—The number of switch ports

For example, the HP StorageWorks Edge Switch 2/32 is a 2 Gb/s switch with 32 ports.

The 2 Gb/s switch ports autonegotiate the signaling speed. When you connect a 2 Gb/s port directly to a 1 Gb/s port, both ports run at 1 Gb/s in each direction. If the ports are not directly connected, the fabric switch that connects the ports determines the signaling speed.

Model naming

The M-Series switches are called HP StorageWorks Edge or Director switches. Director switches are core (enterprise-class) switches. The Edge Switch 2/32 is a mid-range edge switch. The Edge Switch 2/12 and 2/24 are entry-level edge switches.

Switch models

Table 28 describes the M-Series switches supported by HP (1 Gb/s and 2 Gb/s M-Series Fibre Channel switch models). These switches include legacy products released by HP and Compaq.

HP supports all M-Series switches in a fabric if you:

- Use the firmware versions listed.
- Do not exceed the maximum number of Fibre Channel ports listed.
- Follow the fabric rules. (See "Fabric rules" on page 110.)

Table 28: M-Series switches

HP StorageWorks switch name		Firmware version	Number of ports
HP StorageWorks Edge Switch 2/12		07.00.00-84	4 to 12
HP StorageWorks Edge Switch 2/16			16
HP StorageWorks Edge Switch 2/24			8 to 24
HP StorageWorks Edge Switch 2/32			16 to 32
HP StorageWorks Director Switch 2/64			32 to 64
HP StorageWorks Director Switch 2/140			64 to 140
Legacy HP switch name	**Legacy Compaq switch name**		**Number of ports**
N/A	McDATA ES-3016 (Compaq reseller)	07.00.00-84	16
N/A	McDATA ES-3032 (Compaq reseller)		32
McDATA ED-5000 (McDATA reseller)		04.01.01-2	32
HP Director FC-64	Compaq StorageWorks SAN Director 64	07.00.00-84	64

For the latest information on supported M-Series switches and firmware versions, see the HP storage web site: http://h18006.www1.hp.com/storage/saninfrastructure.html.

Features

Features of the M-Series switches are:

- Hot Code Activation Technology (HotCAT) for nondisruptive code activation
- Full nonblocking performance across all ports
- Consistent latency across all ports
- Nondisruptive port expansion
- Redundant, hot-swappable power and cooling systems
- Redundant power cords for separate connections
- Hot-swappable short- and long-wave optical transceivers
- Web server for managing devices and small fabrics
- Fibre Channel public loop connectivity support (Edge Switch 2/12 and 2/24)

Additional features of the Director switches are:

- Four ports on each card to minimize service disruption
- Nonblocking port density to minimize floor-space usage
- Low power consumption and heat generation
- Redundant, hot-swappable switching and processor logic cards
- Nondisruptive failover of redundant components
- Automatic health checks of redundant field-replaceable units (FRUs)

Table 29 compares the high-availability features of M-Series switches.

Table 29: M-Series high-availability feature comparison

Model	Size (1U=1.75")	Redundant control processor/switching	Redundant active components
Edge Switch 2/12	1 U	No	No
Edge Switch 2/16	1 U	No	No
Edge Switch 2/24	1 U	No	No
Edge Switch 2/32	1.5 U	No	No
Director Switch 2/64	9 U	Yes	Yes
Director Switch 2/140	12 U	Yes	Yes

Usage

Table 30 describes the use of M-Series switches as director switches.

Table 30: Using M-Series switches as director switches

Model	12-500 total ports	501-1000 total ports	1000-1632 total ports
Director Switch 2/140	Excellent	Excellent	Excellent
Director Switch 2/64	Excellent	Excellent	Excellent
Edge Switch 2/32	Very good	Good	Good
Edge Switch 2/24	Very good	Good	Good
Edge Switch 2/12	Good	Not recommended	Not recommended

Table 31 describes the use of M-Series switches as edge switches.

Table 31: Using M-Series switches as edge switches

Model	12-500 total ports	501-1000 total ports	1000-1632 total ports
Director Switch 2/140	Good	Very good	Excellent
Director Switch 2/64	Good	Very good	Excellent
Edge Switch 2/32	Excellent	Very good	Good
Edge Switch 2/24	Excellent	Very good	Good
Edge Switch 2/12	Very good	Very good	Good

Fabric rules

The following SAN fabric rules apply to all M-Series SANs. They also apply, in general, to Continuous Access XP, Continuous Access EVA, and Data Replication Manager (DRM) configurations. However, additional rules apply to Continuous Access and DRM implementations. See the *HP StorageWorks Continuous Access EVA planning guide* for detailed information.

When using M-Series switches in a fabric, consider the following:

- Use the HP default settings for all current HP switch model.
- Legacy HP default switch settings and legacy Compaq default switch settings are the same for equivalent switches.
- All switch and fabric rules apply to fabrics implemented with the switch firmware versions listed in Table 28 and HAFM version 07.x or 08.x. The ED-5000 switch requires version 04.01.00-16 firmware and a minimum of HAFM version 04.02.00.

This section describes the following topics:

Operating systems and storage models

The fabric rules for M-Series switches apply to SANs that include the following operating systems and storage models:

Operating systems:

- HP-UX
- OpenVMS
- Tru64 UNIX
- IBM AIX
- Linux
- Microsoft Windows
- Novell NetWare
- Sun Solaris
- VMware ESX

Storage models:

- XP12000
- XP128/1024
- XP48/256/512
- VA7100/7110/7400/7410
- EVA3000/4000/5000/6000/80000
- EMA/ESA12000
- EMA16000
- MA/RA8000
- MA6000
- MSA1000
- RA4000/4100

See "Heterogeneous server rules" on page 127, and "SAN storage system rules" on page 171 to determine operating system support for each storage model.

Fabric rules for M-Series switches

Table 32 lists the rules for creating a SAN with M-Series switches.

Table 32: M-Series fabric rules

Rule number	Description
1	A maximum of 24 switches, 1,632 total ports, and 1,024 total user ports are supported in a single fabric. Both 1 Gb/s and 2 Gb/s switches are supported in the fabric if you do not exceed the following limits: ■ Maximum of eight Director Switch 2/64 or Director Switch 2/140 switches ■ For a fabric that contains ED-5000 switches, a maximum of 16 switches and 512 total ports Eight fully populated Director Switch 2/140 switches exceed the1,024 user-port maximum. Zoning restricts the configuration to 1,024 unique zone members (user ports). However, you can use the remaining ports as ISL connections in the fabric.
2	A maximum of three switch hops (four switches) between any two devices.
3	A ring fabric topology is supported with a maximum of seven M-Series switches.
4	Assign a unique domain number (domain ID) and a unique world wide name (WWN) to each switch. The switch configuration parameters must be the same for each switch. Do not configure any switches with a domain ID of 8, which is reserved for HP-UX.
5	HP requires that all switches in a single-fabric SAN or multifabric SAN use the same firmware version for all like switches. You can use two successive switch firmware versions temporarily in one fabric or multiple fabrics when updating switch firmware.

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

ISL maximums

When designing a fabric using 4-port, 8-port, 12-port, 16-port, and 32-port switches, all ports can be used as ISLs (with a maximum of one half of the total ISL port count configured to the same destination switch). Table 33 lists the maximums for switches with higher port counts.

Note: Some switches require licensing for additional ISL ports.

Table 33: ISL maximums

Switch	Total number of available user ports	Number of E-Ports allowed
McDATA ED-5000	32	4
HP Director FC-64 Compaq StorageWorks SAN Director 64	64	32
HP StorageWorks Director 2/64	64	48 (75% of installed ports)
HP StorageWorks Director 2/140	140	70 (50% of installed ports)

Zoning limits and enforcement

Table 34 lists zoning limits for M-Series switches.

Table 34: M-series zoning limits

Zoning parameter	Maximum value
Number of zone members in a zone	2,048
Number of zones in a zone set	1,024
Number of unique zone members in a zone set	2,048
Total number of zone members in a zone set (where a zone member can be in multiple zones)	4,096
Characters per zoning name	32
Number of unique zone members in HAFM Zoning Library	2,048
Number of zones in HAFM Zoning Library	1,024
Number of zone sets in HAFM Zoning Library	64
Number of end ports	1,024
Number of devices supported (including loop devices)	1,024

Table 35 describes zoning enforcement for M-Series switches.

Table 35: Zoning enforcement for M-Series switches

Switches	Configuration	Enforcement	Comments
HP Director FC-64 Compaq StorageWorks SAN Director 64 HP StorageWorks Edge Switch 2/12 Edge Switch 2/16 Edge Switch 2/24 Edge Switch 2/32 Director Switch 2/64 Director Switch 2/140	Define zones using domain number, port number Define zones using WWNs only Define zones using combination of domain/port numbers and WWNs	Discovery authentication Name Servers (NS) directory based Access authorization at frame level in hardware	Soft zoning or hard zoning (5.01.00-24 and later)

SAN fabric connectivity and switch interoperability rules

7

This chapter describes SAN fabric connectivity and interoperability rules. It describes the following topics:

SAN fabric connectivity rules

This section describes the following topics:

Switch port interfaces

The switch port interfaces are:

- **E_Port**—Provides switch-to-switch connectivity for interswitch links (ISLs).
- **EX_Port**—Connects a Fibre Channel router to an edge fabric.
- **F_Port**—Provides fabric-attached device connectivity for initiators (HBAs) and targets (storage ports).
- **FL_Port**—Provides fabric-aware public loop connectivity with 24-bit Fibre Channel addressing capability.
- **FCAL_Port**—Provides private loop connectivity for 8-bit Fibre Channel addressable devices (requires B-Series QuickLoop feature).

  QuickLoop enables private FC-AL initiators and targets to communicate through the switch. A target that is not configured with QuickLoop cannot communicate with a QuickLoop initiator. QuickLoop is supported only for specific B-Series switches and legacy storage systems. For more information, contact an HP storage representative.

Fiber optic cables

All 4 Gb/s and 2 Gb/s components use industry-standard LC connectors for fiber optic cable connections; all 1 Gb/s components use industry-standard SC connectors. Cables and adapters are available with SC connectors on one end and LC connectors on the other end.

Some fiber optic cable configurations require SC or LC connector sleeves (duplex couplers) to couple the cable connector ends (for example, if you use wall jacks or connect to existing installed cables). HP supports the use of duplex couplers, provided that you do not exceed the overall cable loss budget for that cable segment.

For cable part number information, see the Fibre Channel switch QuickSpecs. QuickSpecs are available at the HP storage web site:
http://h18006.www1.hp.com/storage/saninfrastructure.html

Select a switch from the Fibre Channel Switches section, and then select **Specifications**.

Table 36 lists the rules for fiber optic cable connections.

Table 36: Rules for fiber optic cable connections

Rule number	Description
1	The minimum bend radius is 25 mm for 50, 62.5, and 9 micron fiber optic cable. HP recommends 50 micron fiber optic cable for new installations that require multimode fiber connections. The 62.5 micron fiber optic cable is acceptable for existing installations. See Table 41 for supported maximum cable distances.
2	HP does not support 50, 62.5, and 9 micron cable in the same cable segment.
3	The minimum cable segment length between Fibre Channel devices (transmitter and receiver) is: ■ 0.5 meters for 50 and 62.5 micron cable ■ 2.0 meters for 9 micron cable The minimum length does not apply to patch cords through a passive patch panel; it applies only to the total distance between the transmitter and receiver of each device connected through the patch panel.
4	For fiber optic cable lengths greater than 50 meters, contact a third-party provider. 50 micron cable must be duplex, tight-buffered multimode 50/125 µm (Belcore GR-409 compliant). The connectors must be SC or LC duplex low metal (Belcore and IEC compliant). 9 micron cable must be duplex, tight-buffered, single-mode 9/125 µm (Belcore GR-409 compliant). The connectors must be SC or LC duplex low metal (NTT-SC Belcore 326, IEC-874-19 SC compliant).

Fiber optic cable loss budgets

Cable loss budgets are determined by the Fibre Channel Physical Interface Specification (see the standards at www.incits.org). The maximum distances specified are based on the use of nominal bandwidth fiber optic cable. This specifies modal bandwidth of 500 MHz-km for 50 micron fiber optic cable, and 200 MHz-km for 62.5 micron fiber optic cable.

Table 37 through Table 39 list the maximum loss budgets for different interconnect speeds at specific distances. Table 40 through Table 46 list the maximum supported distances based on switch-to-switch (ISL) or device-to-switch connectivity.

Note: The following tables do not specify media losses due to variances between different fiber optic cable manufacturers. In all cases, the specification that must be followed is the total channel insertion loss, which includes media losses.

HP supports the use of optical fiber patch panels. The total channel insertion loss between the transmitter and receiver for the cable segment routed through the patch panel must not exceed the maximum listed for the connector and cable type.

Table 37 lists the 4 Gb/s fiber optic cable loss budgets.

Table 37: 4 Gb/s fiber optic cable loss budgets

Cable	Maximum distance per cable segment	Total channel insertion loss[1]	Loss per mated connector pair
50/125 micron	150 meters	2.06 dB	0.75 dB
62.5/125 micron	70 meters	1.78 dB	0.75 dB

1. Channel insertion loss is the combined passive loss from connectors, splices, and media between the transmitter and receiver.

Table 38 lists the 2 Gb/s fiber optic cable loss budgets.

Table 38: 2 Gb/s fiber optic cable loss budgets

Cable	Maximum distance per cable segment	Total channel insertion loss[1]	Loss per mated connector pair
50/125 micron	300 meters	2.62 dB	0.75 dB
62.5/125 micron	150 meters	2.1 dB	0.75 dB
9/125 micron	10 km (6.2 miles)	7.8 dB	0.75 dB
9/125 micron	35 km (21.7 miles)	21.5 dB	0.75 dB

1. Channel insertion loss is the combined passive loss from connectors, splices, and media between the transmitter and receiver.

Table 39 lists the 1 Gb/s fiber optic cable loss budgets.

Table 39: 1 Gb/s fiber optic cable loss budgets

Cable	Maximum distance per cable segment	Total channel insertion loss[1]	Loss per mated connector pair
50/125 micron	500 meters	3.85 dB	0.75 dB
62.5/125 micron	200 meters	3.0 dB	0.75 dB
9/125 micron	10 km (6.2 miles)	7.8 dB	0.75 dB
9/125 micron	35 km (21.7 miles)	21.5 dB	0.75 dB
9/125 micron	100 km (62 miles)	19 dB	0.75 dB

1. Channel insertion loss is the combined passive loss from connectors, splices, and media between the transmitter and receiver.

Storage product interface and transport distance rules

Table 40 through Table 46 describe the maximum distances supported for each cable segment type, switch-to-switch or device-to-switch for each interface and transport type. Unless otherwise specified, the distances specified apply to both switch-to-switch connectivity (ISL) and device-to-switch connectivity.

See the DRM and Continuous Access EVA documentation for interconnect and distance rules for those products. For more information, see the HP storage web site: http://h18006.www1.hp.com/products/sanworks/drm/index.html

Table 40 describes the distance rules for 4 Gb/s Fibre Channel connections.

Table 40: 4 Gb/s Fibre Channel distance rules

Interface/transport	Supported distances	Supported storage products
50 micron multimode fiber optic cable and short-wave SFPs	150 meters ISL at 4 Gb/s 300 meters at 2 Gb/s 500 meters at 1 Gb/s	Heterogeneous SAN host-to-disk storage systems DRM and Continuous Access EVA and XP Enterprise Backup Solutions (EBS)
62.5 micron[1] multimode fiber optic cable and short-wave SFPs	70 meters ISL at 4 Gb/s	
9 micron single-mode fiber optic cable and long-wave SFPs	10 km ISL at 2 Gb/s with 2 Gb/s SFP	
9 micron single-mode fiber optic cable and extended-reach SFPs	35 km ISL at 2 Gb/s with 2 Gb/s SFP	
Fibre Channel using WDM[2]	100 km at 4 Gb/s 200 km at 2 Gb/s 400 km at 1 Gb/s	

1. Information for 62.5 micron fiber optic cable is provided to facilitate use of installed cable. HP recommends 50 micron fiber optic cable for new installations that require multimode fiber.
2. WDM distance is the maximum distance for the WDM link, distances specified when using B-Series v4.4x or later firmware.

Table 41 describes the distance rules for 2 Gb/s Fibre Channel connections.

Table 41: 2 Gb/s Fibre Channel distance rules

Interface/transport	Supported distances	Supported storage products
50 micron multimode fiber optic cable and short-wave SFPs	300 meters at 2 Gb/s 500 meters at 1 Gb/s	Heterogeneous SAN host-to-disk storage systems DRM and Continuous Access EVA and XP Enterprise Backup Solutions (EBS)
62.5 micron[1] multimode fiber optic cable and short-wave SFPs	150 meters at 2 Gb/s	
9 micron single-mode fiber optic cable and long-wave SFPs	10 km ISL at 2 Gb/s	
9 micron single-mode fiber optic cable and extended-reach SFPs	35 km ISL at 2 Gb/s	
Total distance[2]	200 km (B-Series, C-Series) 105 km (M-Series)	
Fibre Channel using WDM[3]	100 km at 2 Gb/s	

1. Information for 62.5 micron fiber optic cable is provided to facilitate use of installed cable. HP recommends 50 micron fiber optic cable for new installations that require multimode fiber.
2. Total distance is the sum of all cable segments.
3. WDM distance is the maximum distance for the WDM link.

Table 42 describes the distance rules for 1 Gb/s Fibre Channel connections.

Table 42: 1 Gb/s Fibre Channel distance rules

Interface/transport	Supported distances	Supported storage products
50 micron multimode fiber optic cable and short-wave GLMs[1]	500 meters at 1 Gb/s	Heterogeneous SAN host-to-disk storage systems DRM and Continuous Access EVA and XP Enterprise Backup Solutions (EBS)
62.5 micron multimode fiber optic cable and short-wave GBICs[2]	200 meters at 1 Gb/s	
9 micron single-mode fiber optic cable and long-wave GBICs	10 km ISL at 1 Gb/s 35 km ISL (C-Series, M-Series)	
9 micron single-mode fiber optic cable and very long distance GBICs	100 km ISL (B-Series) at 1 Gb/s	
Total distance[3]	200 km (B-Series, C-Series) 105 km (M-Series)	
Fibre Channel using WDM[4, 5]	200 km at 1 Gb/s	

1. Gigabit Link Modules (GLMs) are used in HSG80-based (MA/RA8000, EMA/ESA12000, EMA16000) and HSG60-based (MA6000) storage systems.
2. Information for 62.5 micron fiber optic cable is provided to facilitate use of installed cable. HP recommends 50 micron fiber optic cable for new installations that require multimode fiber.
3. Total distance is the sum of all cable segments.
4. WDM distance is the maximum distance for the WDM link
5. Up to 200 km over a WDM link at 1 Gb/s Fibre Channel with reduced performance levels. The performance levels depend on the number of buffers available in the switch and the application data transfer size.

Table 43 describes the distance rules for ATM extension Fibre Channel connections.

Table 43: ATM extension Fibre Channel distance rules

Interface/transport	Heterogeneous SAN host-to-disk storage systems	DRM	Continuous Access XP[1 2]	Enterprise Backup Solutions (EBS)
One T1/E1 WAN per fabric	Not supported	FC-to-ATM converter, no IP network delay limit Not supported as one ISL in a single-fabric implementation	Intersite backbone with FC-to-IP converter, 100 ms of IP network delay Not supported as one ISL in a single-fabric implementation	Not supported
One T1/E1 WAN per fabric (inverse multiplexing)				
T3/E3 WAN	Not supported	FC-to-ATM converter, no IP network delay limit	Intersite backbone with FC-to-IP converter, 100 ms of IP network delay	
Fractional and/or shared T3/E3 and Optical Carrier (OC) 3 WAN				

1. For supported IP bandwidth levels, see the product documentation.
2. For details on Continuous Access support limits, see the Continuous Access and Data Replication Manager SAN extensions reference guide available at http://www.hp.com/go/SANDesignGuide.

Table 44 describes the distance rules for FCIP (Fibre Channel over IP) extension Fibre Channel connections.

Table 44: FCIP extension Fibre Channel distance rules

Interface/transport	Heterogeneous SAN host-to-disk storage systems	DRM and Continuous Access EVA and XP[1]	Enterprise Backup Solutions (EBS)
HP B-Series MP Router (FCIP)	EVA, XP, VA, EMA/ESA12000, EMA16000, MA/RA8000 100 ms of IP network delay MSA1000/1500 200 km maximum MP Router to switch	100 ms of IP network delay 200 km maximum MP Router to switch **Note**—200 km requires connectivity to a Fibre Channel switch using 4.4 or later FW.	Not supported
HP C-Series MDS 9216i, IPS-4, IPS-8, 14/2 (FCIP)	EVA, XP, VA, EMA/ESA12000, EMA16000, MA/RA8000 100 ms of IP network delay 200 km total Fibre Channel (C-Series) MSA1000/1500 200 km total Fibre Channel (C-Series)	100 ms of IP network delay 200 km total Fibre Channel (C-Series)	
HP SR2122-2 IP Storage Router (FCIP)	EVA, XP, VA, EMA/ESA12000, EMA16000, MA/RA8000, MSA1000/1500 100 ms of IP network delay 200 km total Fibre Channel (B-Series, C-Series) 105 km total Fibre Channel (M-Series) MSA1000/1500 200 km total Fibre Channel (C-Series) 105 km total Fibre Channel (M-Series)	100 ms of IP network delay 200 km total Fibre Channel (B-Series, C-Series) 105 km total Fibre Channel (M-Series)	100 ms of IP network delay 200 km total Fibre Channel (B-Series, C-Series) 105 km total Fibre Channel (M-Series)
Third-party SAN extension devices	See the *HP StorageWorks Continuous Access and Data Replication Manager SAN extensions reference guide*	See the *HP StorageWorks Continuous Access and Data Replication Manager SAN extensions reference guide*	Not supported

1. For details on Continuous Access support limits, see the Continuous Access and Data Replication Manager SAN extensions reference guide available at http://www.hp.com/go/SANDesignGuide.

Table 45 describes the distance rules for iSCSI (Internet SCSI) bridging Fibre Channel connections.

Table 45: iSCSI bridging Fibre Channel distance rules

Interface/transport	Heterogeneous SAN host-to-disk storage systems	DRM and Continuous Access EVA and XP[1]	Enterprise Backup Solutions (EBS)
HP C-Series MDS 9216i, IPS-4, IPS-8, 14/2 (iSCSI)	EVA, XP, VA, EMA/ESA12000, EMA16000, MA/RA8000, MSA1000/1500 200 km total Fibre Channel (C-Series)	Not supported for iSCSI bridging	
HP SR2122-2 IP Storage Router (iSCSI)	EVA, XP, VA, EMA/ESA12000, EMA16000, MA/RA8000, MSA1000/1500 200 km total Fibre Channel (B-Series, C-Series) 105 km total Fibre Channel (M-Series)	Not supported for iSCSI bridging	

1. For details on Continuous Access support limits, see the Continuous Access and Data Replication Manager SAN extensions reference guide available at http://www.hp.com/go/SANDesignGuide.

Table 46 describes the distance rules for Fibre Channel routing connections.

Table 46: Fibre Channel routing distance rules

Interface/transport	Heterogeneous SAN host-to-disk storage systems	DRM and Continuous Access EVA and XP[1]	Enterprise Backup Solutions (EBS)
HP B-Series MP Router (Fibre Channel routing, B-Series only)	EVA, XP, VA, EMA/ESA12000, EMA16000, MA/RA8000 100ms of IP network delay (when combined with FCIP) MSA1000/1500 200 km maximum MP Router to switch	100 ms of IP network delay (when combined with FCIP) 200 km maximum MP Router to switch **Note**—200 km requires connectivity to a Fibre Channel switch using 4.4 or later FW.	200 km total Fibre Channel 200 km maximum MP Router to switch
HP C-Series IVR (Inter-VSAN routing)	See Table 41 and Table 42	See Table 44.	See Table 41 and Table 42.

1. For details on Continuous Access support limits, see the Continuous Access and Data Replication Manager SAN extensions reference guide available at http://www.hp.com/go/SANDesignGuide.

SAN fabric switch interoperability rules

HP supports the following heterogeneous switch SAN configurations:

- Dual interoperable, heterogeneous SAN fabrics, page 121
- Interoperable, heterogeneous switch fabrics, page 121

Dual interoperable, heterogeneous SAN fabrics

A dual interoperable, heterogeneous SAN consists of one fabric that contains all B-Series switches and another fabric that contains all M-Series switches.

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

- HP recommends that you use the same fabric topology and configuration in both fabrics to maintain balanced SAN performance.
- Design both fabrics using the lowest common denominator. For example, B-Series switches support seven switch hops and M-Series switches support three switch hops. Therefore, design both fabrics with a maximum of three switch hops.
- High-availability (HA) configurations require support for common HBA, driver, multipathing software, and storage array firmware versions because servers and storage connect to both fabrics.

Note: HP does not support MSA1000 with the MSA SAN Switch 2/8 or Continuous Access/DRM storage systems in dual interoperable, heterogeneous SAN configurations.

Interoperable, heterogeneous switch fabrics

An interoperable, heterogeneous switch fabric can contain different series of switches. Table 47 lists the switch combinations.

Table 47: Heterogeneous switches in the same fabric

Heterogeneous switch combinations	Reference
C-Series with B-Series	*Fabric interoperability: Merging fabrics based on C-Series and B-Series Fibre Channel switches application notes*
C-Series and M-Series	*Fabric interoperability: Merging fabrics based on C-Series and M-Series Fibre Channel switches application notes*
M-Series and B-Series	*Fabric interoperability: Merging fabrics based on M-Series and B-Series Fibre Channel switches application notes*

These documents are available at the storage web site: http://www.hp.com/go/SANDesignGuide.

Note: HP does not support MSA1000/MSA1500 or Continuous Access/DRM storage systems in interoperable heterogeneous switch fabric configurations.

Third-party switch support

HP offers support for certain switches from other vendors if you purchase third-party support through the HP SAN Environmental Services Group (SAN-ES).

SAN performance considerations

The following SAN components affect SAN application performance:

- Host CPU(s)
- Fibre Channel HBAs
- SAN topology and the number of fabrics
- I/O transfer sizes and usage patterns
- RAID controllers
- Disk configuration

This section describes the following topics:

Infrastructure factors

A single-switch fabric provides the highest level of performance. In a fabric with multiple switches, the following factors can affect performance:

- **Latency**

  Switch latency is less than 5% (at 1 Gb/s) of the time for a data transfer; therefore, the number of switches and hops between devices is not a major performance factor. However, as devices send frames through more switches and hops, other data traffic in the fabric routed through the same ISL or path can cause oversubscription.

- **Oversubscription**

  Oversubscription degrades Fibre Channel performance. When devices must contend for the same ISL or path, each device receives an equal share or 1/*n*th of the available bandwidth on the path (where *n* is the number of contending devices). Oversubscription occurs when one or more devices sends more data than the total bandwidth available on the ISL or path.

- **Fabric interconnect speeds**

  Fibre Channel supports 4 Gb/s, 2 Gb/s, and 1 Gb/s speeds. The highest performance is attained by configuring a fabric with all components at the same, highest available speed. Additional factors such as distance, number of switch and device port buffers, and device response times can also affect performance.

- **Mixed Fibre Channel speeds**

  For fabrics consisting of 4 Gb/s, 2 Gb/s, and 1 Gb/s switches and devices, the fabric segment connections negotiate the speed at which specific devices communicate.

  The presence of 1 Gb/s devices in a fabric does not force other independent 2 Gb/s devices or paths or 4 Gb/s paths to a lower speed. Fibre Channel requires that all ports be able to negotiate to one of three supported speeds. Switch ports or user ports in a fabric communicate at the highest mutually supported speed.

Performance guidelines

Although the topology and size of the fabric affect performance, adhering to the rules and recommendations outlined in this guide minimizes these factors. The topology designs have been defined to accommodate specific data access types. Recommendations on the number of ISLs based on device-to-device access ratios ensure that adequate bandwidth is available across the fabric, minimizing oversubscription.

To maximize fabric performance, HP recommends the following guidelines:

- Implement dual-fabric SANs.
- In a cascaded or core-edge fabric, position switches with the highest port speeds near the center of the fabric.
- Use the highest speed available for all infrastructure components and devices.
- Ensure that communicating devices have the same speed connectivity path through the fabric.
- Connect devices that communicate frequently to the same Fibre Channel switch.
- When possible, ensure that there is an equal number of high-bandwidth application servers and storage systems (for one-to-one access).
- Ensure that Fibre Channel Congestion Control (FCC) is enabled on all C-Series switches.

  FCC allows C-Series switches to intelligently regulate traffic across ISLs and ensure that each initiator-target pair of devices has the required bandwidth for data transfer. C-Series switches can also prioritize frames using the Quality of Service (QoS) feature.

Volume 3

Host and storage system rules

Host and storage system rules are presented in these chapters:

Heterogeneous server rules

8

This chapter describes rules related to specific servers and operating system platforms. It describes the following topics:

For additional information, see the platform and individual product-specific documentation or contact an HP storage representative. See “About this guide” on page 23 for a list of related documentation.

General platform/operating system and storage system rules

1. Each platform listed is supported in all Fabric topology configurations unless otherwise noted in this guide or the applicable platform documentation.
2. Any mix of heterogeneous servers, clustered and standalone, is allowed in a SAN provided that you follow all individual platform rules, fabric rules, applicable server application rules, and the maximums listed in this guide and in the platform specific documentation.
3. All HP and multi-vendor hardware platforms and operating systems that are supported in a homogeneous SAN are supported in a heterogeneous SAN. See Table 51, and Table 53 to determine if zoning is required for specific combinations of supported heterogeneous platforms.
4. Servers can attach to multiple fabrics. The number of separate fabrics per server is based on the specific server model capabilities and the maximum number of Fibre Channel host bus adapters supported.
5. See "High-availability configuration considerations" on page 199 for cabling scheme options for platforms that support high availability multipathing.
6. Any mix of storage systems is allowed in a SAN, provided that you follow all applicable fabric rules, platform/operating system rules, storage system rules, and mixed storage common SAN rules.
   - Fabric Rules - See:
   - "B-Series switches and fabric rules" on page 81
   - "C-Series switches and fabric rules" on page 97
   - "M-Series switches and fabric rules" on page 105
   - "SAN fabric connectivity and switch interoperability rules" on page 113
   - "SAN extension" on page 209
   - Platform/Operating System Rules - Based on the storage system type(s) and switch product line(s) utilized, see:
   - "Specific platform/operating system rules – EVA3000/5000 (VCS v3), EVA4000/6000/8000 (XCS v5), EMA/ESA12000, EMA16000, MA/RA8000, MA6000 (ACS 8.7, 8.8) storage systems, B-Series and M-Series switches" on page 147.
   - "Specific platform/operating system rules – EVA3000/5000 (VCS v3), EVA4000/6000/8000 (XCS v5), EMA/ESA12000, EMA16000, MA/RA8000, MA6000 (ACS 8.7, 8.8) storage systems, C-Series switches" on page 157.
   - "Specific platform/operating system rules – HP XP and VA storage systems" on page 141.
   - "Specific platform/operating system rules – MSA1500, MSA1000, RA4100, RA4000" on page 184.
   - Storage System Rules - Based on the storage system type(s) utilized, see:
   - "HP XP and VA configuration rules" on page 172
   - "EVA3000/4000/5000/6000/8000 configuration rules" on page 173
   - "EMA/ESA12000, EMA16000, MA/RA8000, MA6000 configuration rules" on page 177
   - "MSA1000 configuration rules" on page 188
   - "RA4100 and RA4000 configuration rules" on page 191

- "Mixed storage type SAN rules - B-Series, C-Series, M-Series switches" on page 135 (For SAN fabrics containing a mix of different storage system types).

7. See "Platform Interoperability for Single Shared EMA/ESA12000, EMA16000, MA/RA8000, MA6000 Storage Systems – ACS 8.7, 8.8" on page 166**,** for information related to mixing heterogeneous platforms on a single shared EMA/ESA12000, EMA16000, MA/RA8000, or MA6000 storage system. In certain situations multiple storage systems may be required to accommodate the requirements of different platforms or operating systems.
8. There are different limits relative to the number of switches supported in a Fabric based on the Fibre Channel switch product line in use. See B-Series switches and fabric rules on page 81, C-Series switches and fabric rules on page 97, and M-Series switches and fabric rules on page 105, for more information.

Blade Server support

The BL2xp and BL40p Blade Server products are supported with all B-Series, C-Series, and M-Series product line switches. The BL3xp Blade Server models are supported with all B-Series and C-Series product line switches, and M-Series Edge Switch 2/24 and Edge Switch 2/12 models. See the HP ProLiant BL p-Class SAN product support page at:

http://h18006.www1.hp.com/products/storageworks/bladesystemmatrix/index.html

MSA1000 Small Business SAN

The HP StorageWorks Modular Smart Array 1000 Small Business SAN Kit is a 2 Gb Fibre Channel storage system for the entry-level storage area network (SAN). It provides a low-cost, scalable, high performance storage consolidation system with investment protection. It is designed to reduce the complexity and risk of SAN deployments.

The MSA1000 Small Business SAN is a standalone SAN configuration that is not supported for connectivity to B-Series, C-Series, or M-Series SAN fabrics. See the HP StorageWorks Modular Smart Array 1000 Small Business SAN Kit web page technical documents section for specific configuration support information at:

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

For support information on MSA1000 Fibre Channel storage arrays that are not part of the Small Business SAN, see "Specific platform/operating system rules – MSA1500, MSA1000, RA4100, RA4000" on page 184, for B-Series, C-Series, or M-Series fabric support.

NonStop server support

The NonStop S-series servers (processor versions S78, S780, S760, S7800, S76000, S86000, S88000, and later) are supported with XP 128, XP 1024, and XP 12000 arrays in a direct-connect configuration as well as SAN configurations comprising of selected Fibre Channel switches.

The following sections briefly describe the rules and guidelines surrounding the design of these configurations with the XP arrays.

Direct connect configuration

- Requires a minimum of one XP Subsystem for storage connectivity.
- Requires a minimum of one I/O Adapter Module from the server side.

- Requires a minimum of two FCSA adapters in an IOAM enclosure as shown in the figures below. Note that each FCSA adapter has two FC ports.
- Host based mirroring is strongly recommended**-** i.e. Each LDEV (P) is mirrored to a separate LDEV (M) on separate XP ports (p, b, m, mb paths used); but a non-mirrored volume is allowed- i.e., Each LDEV (P) is not mirrored to a separate LDEV (M) on separate XP ports (only p, b paths used).
- Each LDEV requires two LUNs.
- For availability, the Primary (P) LDEVs and Mirror (M) LDEVs are to be configured on separate array ACP pairs.
- For availability, the p and b paths should be in separate XP array clusters. The m and mb paths should be in separate clusters.
- Recommended that the p and mb paths be in the same XP array cluster and b and m paths be in the same XP array cluster for a volume.
- No Boot Support.
- Requires mode 0C to be set on XP1024/XP128 ports.
- The FCSA HBA is supported with 1Gb/2Gb CHIPs for XP1024/128

Figure 34 on page 132 illustrates a minimum direct connect configuration.

SAN configuration

- Requires a minimum of one XP Subsystem for storage connectivity
- Requires a minimum of one I/O Adapter Module from the server side
- Requires a minimum of two FCSA adapters in an IOAM enclosure as shown in the figures below. Note that each FCSA adapter has two FC ports.
- Requires dual redundant fabrics (level 4, non-interconnected NSPOF high-availability SAN configuration. See “Data availability” on page 59 for information about data availability levels.) Each fabric consists of either a single switch or two cascaded switches as shown in Figure 35 on page 132, Figure 36 on page 133, and Figure 37 on page 134.
- No more than three switches are supported in a single fabric.
- Select B-Series and C-Series switches are supported. Contact HP support for the list of supported switches and firmware versions.
- Fabrics must consist of a single series of switches only. That is, fabrics with mixed switch series are not supported. For example, having B-Series and C-Series switches in the same fabric is not supported for use in NonStop server configurations.
- Requires configuring separate zones, each consisting of the set of HBAs on a single NonStop system, and the XP array ports that those ports need to communicate with, that needs to be seen by the HBA. Configure WWN based zoning only.
- Only NonStop homogeneous connections are allowed to the same switch or to the SAN. In other words, no other OS's can share the same switch/SAN, even if they are in different zones.
- HP strongly recommends host-based mirroring. For example, each LDEV (P) is mirrored to a separate LDEV (M) on separate XP ports (p, b, m, mb paths used); but a non-mirrored volume is allowed. For example, each LDEV (P) is not mirrored to a separate LDEV (M) on separate XP ports (only p, and b paths are used).
- Each LDEV requires two LUNs.

- For availability, Primary (P) LDEVs and Mirror (M) LDEVs are to be configured on separate array ACP pairs.
- For availability, the p and b paths should be in separate XP array clusters. The m and mb paths should be in separate clusters.
- HP recommends the p and mb paths be in the same XP array cluster, and b and m paths be in the same XP array cluster for a volume.
- No Boot Support.
- Requires mode 0C to be set on XP1024/XP128 ports.
- The FCSA HBA and Fibre Channel switches are supported with 1Gb/2Gb CHIPs for XP1024/128.

See Figure 35 on page 132, Figure 36 on page 133, and Figure 37 on page 134.

Table 48 describes the minimum are recommended details for each of the supported configurations.

Table 48: NonStop supported configurations

	Direct connect configuration min/rec	SAN configuration min/rec	Maximum availability configuration min/rec
Number of FC SAN Fabrics	0/0	2/2	2/4
Number of XP Subsystems	1/1	1/1	1/2
Number of I/O Adapter Modules	1/1	1/1	2/2
Number of FC ServerNet Adapters	2/4	2/4	4/4

Note: The following figures show high-availability configurations. Figure 36 and Figure 37 show configurations with physical IOAM redundancy as well.

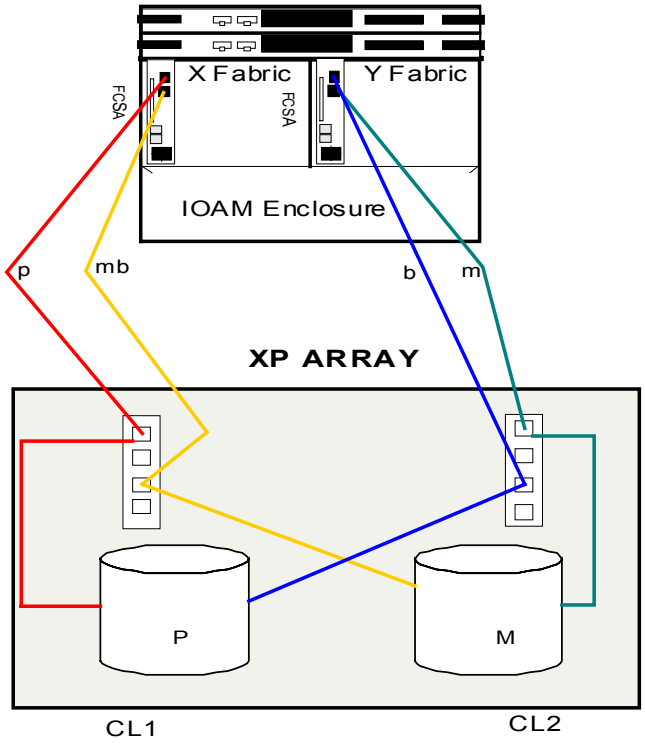

Figure 34: Minimum Direct Connect Configuration for XP

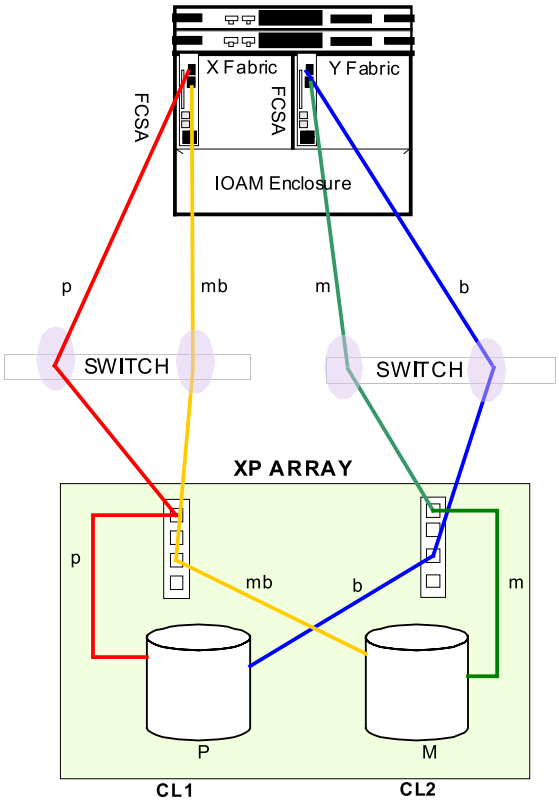

Figure 35: Minimum SAN Configuration for XP

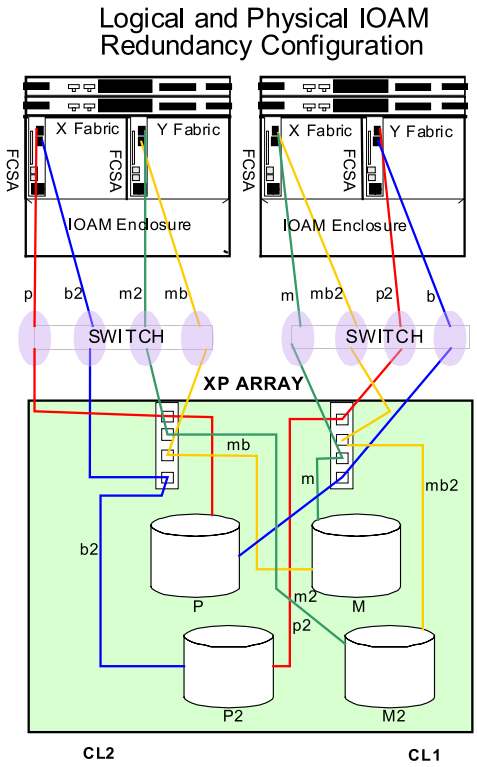

Figure 36: Minimum SAN configuration with logical and physical redundancy for XP

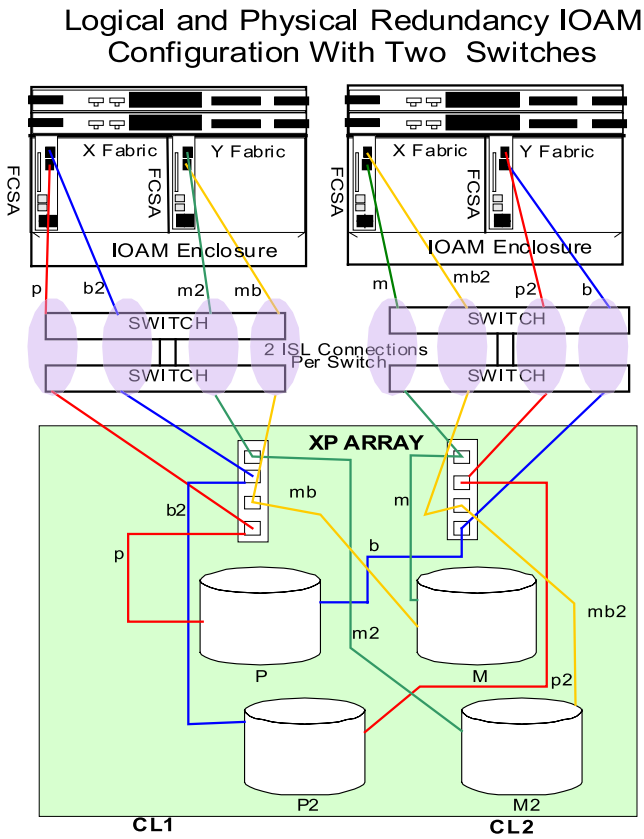

Figure 37: SAN configuration (two cascaded switches) with logical and physical redundancy for XP

Mixed storage type SAN rules - B-Series, C-Series, M-Series switches

HP supports SANs consisting of a mixture of storage system types. This section defines the rules for mixing different storage system types within a SAN using switch models exclusively from one product line of Fibre Channel switch products. For information about interoperability rules and support for SANs or fabrics with mixed switch product line types, see “Interoperable, heterogeneous switch fabrics” on page 121.

Common SAN access

In general, support in the same SAN for mixed storage system families is provided by implementing separate zones for servers and the storage system families being accessed as shown in:

- Figure 38 on page 139, for the B-Series switches
- Figure 39 on page 139, for the M-Series switches
- Figure 40 on page 140 and Figure 41 on page 140, for the C-Series.

As depicted in Figure 41 on page 140, C-Series switches provide a capability to build secure virtual fabrics using the VSAN feature. Each VSAN is a separate virtual fabric that can be dedicated to a different type of storage system and zoning can be implemented on a per-VSAN basis for additional security.

Inter VSAN Routing (IVR), an optional feature in C-Series switches, allows selected hosts and storage in one VSAN to communicate with hosts and storage in a different VSAN.

B-Series switches also support the use of the MP Router, allowing selected hosts and storage in one fabric or LSAN to communicate with hosts and storage in another fabric or LSAN.

Common server access

Common server access allows for simultaneous connectivity to different disk storage system families from the same server and, in some cases, the same HBA.

In addition, certain storage solutions utilizing multiple storage types, for example, disk and tape, may also specify support for common server access. In those cases, see the specific storage solution documentation for the supported common access configurations and rules.

Common server, separate HBAs

Common server access using separate HBAs is supported for XP, VA, MSA and EVA/EMA/ESA12000, EMA16000, MA/RA8000, and MA6000 storage systems for these specific configurations. Contact your HP storage representative for a configuration review prior to deployment.

Table 49: Common server, separate HBAs

Operating system	Storage systems	Notes
IBM AIX	EVA 3000/5000 EVA 4000/6000/8000 XP	Separate AutoPath for XP or MPIO for XP, and Secure Path for EVA HBAs and zones. Specific minimum versions of Auto Path and Secure Path are required. This is supported with B-Series and M-Series switches only.
Linux	EVA 3000/5000 EVA 4000/6000/8000 EMA/ESA12000, EMA16000, MA/RA8000 MSA 1000/1500 XP/VA	Connection to a common server with different HBA models requires separate HBA zones for XP/VA, MSA, and EVA/EMA/ESA12000, EMA16000, MA/RA8000.
Sun Solaris	EVA 3000/5000 EVA 4000/6000/8000 XP	Connection to a common server requires separate VxVM DMP for XP and Secure Path for EVA, HBAs, and zones. Specific minimum versions of VxVM DMP and Secure Path are required.

Common server, common HBAs

Simultaneous access to different storage system types from the same Server/HBA is supported when the storage system families listed use common HBA model numbers, common driver versions, and common multipathing software versions. HP supports the following storage types and operating systems for common server, common HBA access.

Table 50: Common server, common HBAs

Operating system	Storage systems	Notes
HP-UX	EVA 4000/6000/8000 XP12000, XP128/1024, XP48/512, XP256, VA7410/7110, VA7400/7100 MSA1000 (Single Controller)	PvLinks
	EVA 4000/6000/8000 XP12000, XP128/1024, XP48/512, XP256, VA7410/7110, VA7400/7100 EVA 3000/5000 EMA/ESA12000, EMA16000, MA/RA8000	Secure Path v3.0f or later
	XP12000, XP128/1024, XP48/512, XP256, VA7410/7110, VA7400/7100	Veritas DMP
OpenVMS Tru64	EVA 3000/5000 EVA 4000/6000/8000 XP12000, XP128/1024, XP48/512, XP256 EMA/ESA12000, EMA16000, MA/RA8000	Native Multi-path driver EMA/ESA12000, EMA16000, MA/RA8000 requires HSG80 platform kit v8.7 or later
IBM AIX	EVA 4000/6000/8000 XP12000, XP128/1024, XP48/512, XP256	MPIO Auto Path v5.4.2
	EVA 3000/5000 EMA/ESA12000, EMA16000, MA/RA8000	Secure Path v2.0d SP 2 or later

Table 50: Common server, common HBAs

Operating system	Storage systems	Notes
Linux	EVA 3000/5000 EVA 4000/6000/8000 MSA 1000/1500 XP12000, XP128/1024, XP48/512, XP256, VA7410/7110, VA7400/7100	QLogic multi-pathing
	EVA 3000/5000 MSA 1000/1500	Secure Path
NetWare	EVA 3000/5000 EMA/ESA12000, EMA16000, MA/RA8000 MSA 1000	
Windows 2000, 2003	EVA 3000/5000 EVA 4000/6000/8000 EMA/ESA12000, EMA16000, MA/RA8000 XP MSA 1000	Requires v3.0F or later platform kit. MSA1000/1500, EVA3000/5000, XP requires Secure Path v2.0c SP1 or SP2. EVA4000/6000/8000, XP requires Full Feature MPIO. EVA4000/6000/8000 requires updated HBA driver smart components, refer to the HP SAN Infrastructure web page at: http://h18006.www1.hp.com/storage/saninfrastructure.html See the *Emulex Host Bus Adapters for Windows 32-bit Systems Release Notes, HP StorageWorks Booting 32-bit Windows Systems from a Storage Area Network, or HP StorageWorks Booting 64-bit Windows Systems from a Storage Area Network* for more information.
Sun Solaris	EVA 4000/6000/8000 XP	MPxIO
	EVA 3000/5000 EVA 4000/6000/8000 EMA/ESA12000, EMA16000, MA/RA8000	Secure Path UX v3.0d or later

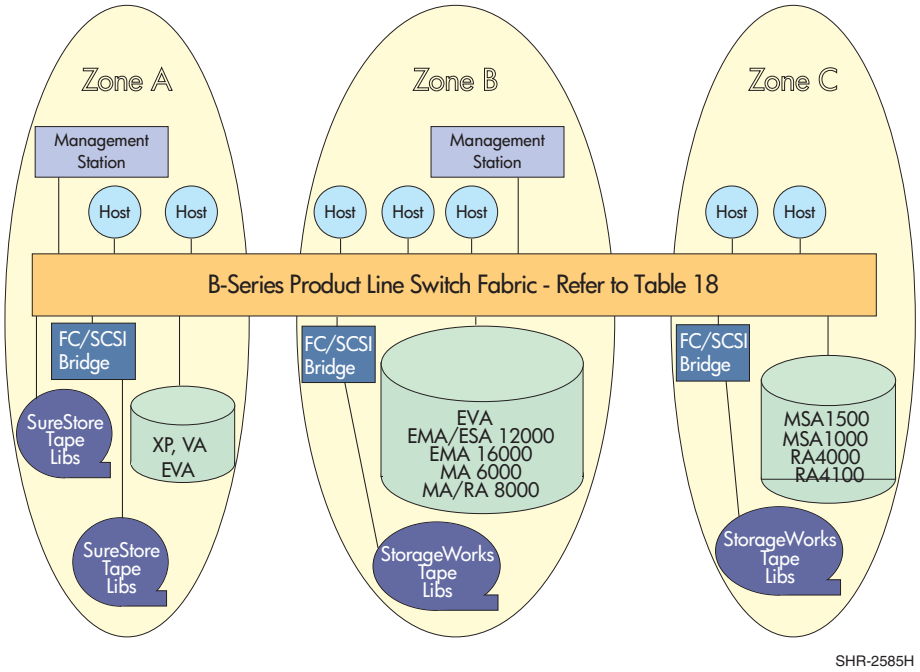

Figure 38: HP StorageWorks SAN using B-Series switches

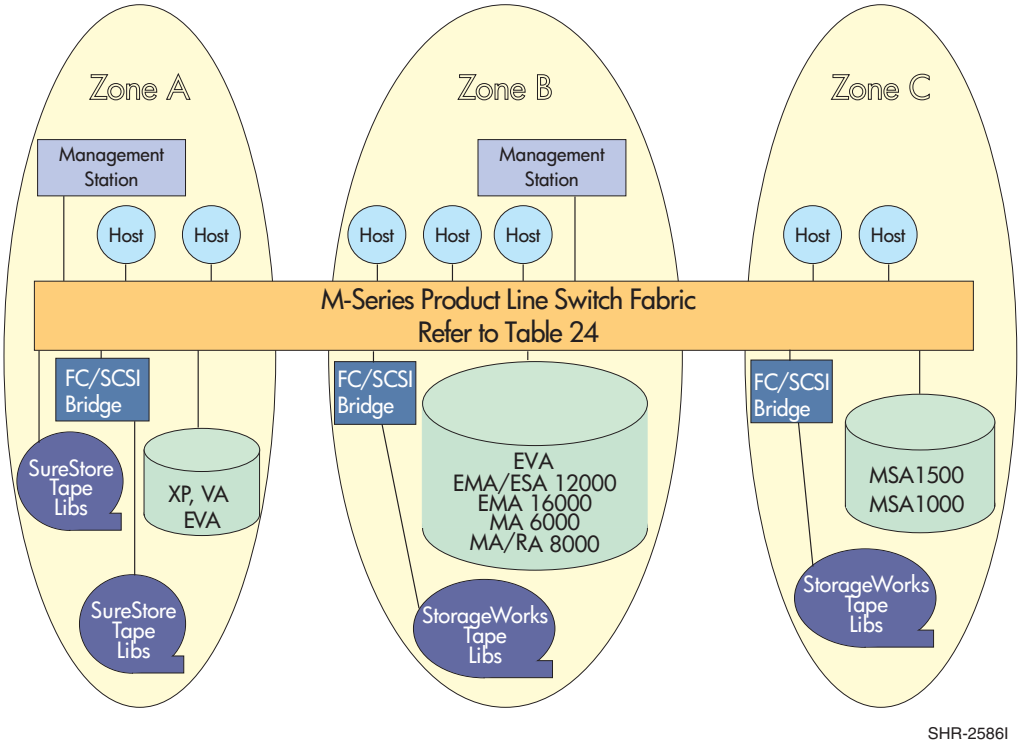

Figure 39: HP StorageWorks SAN using M-Series switches

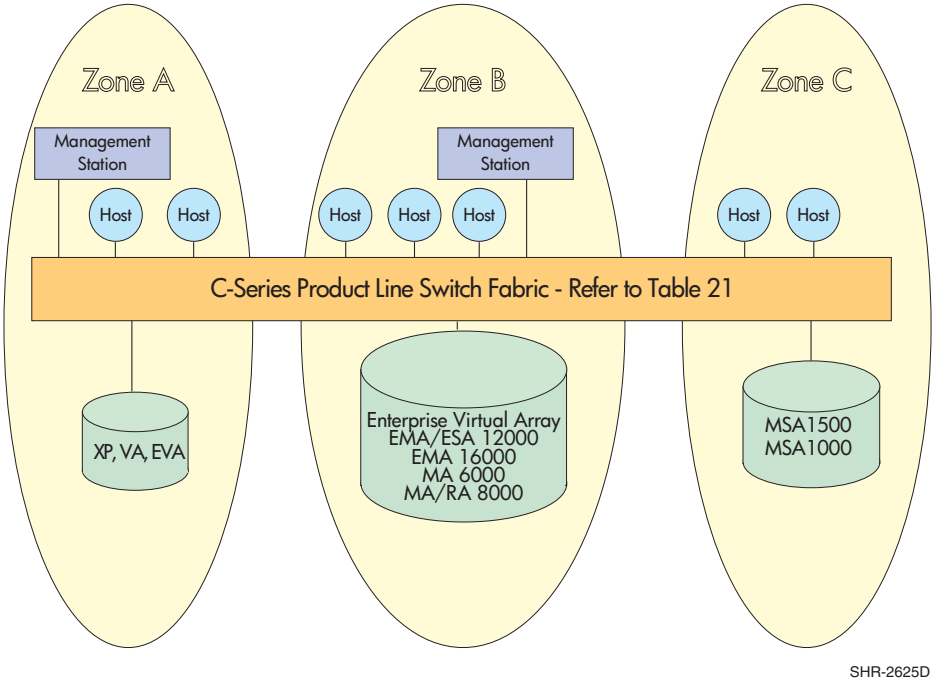

Figure 40: HP StorageWorks SAN using C-Series switches

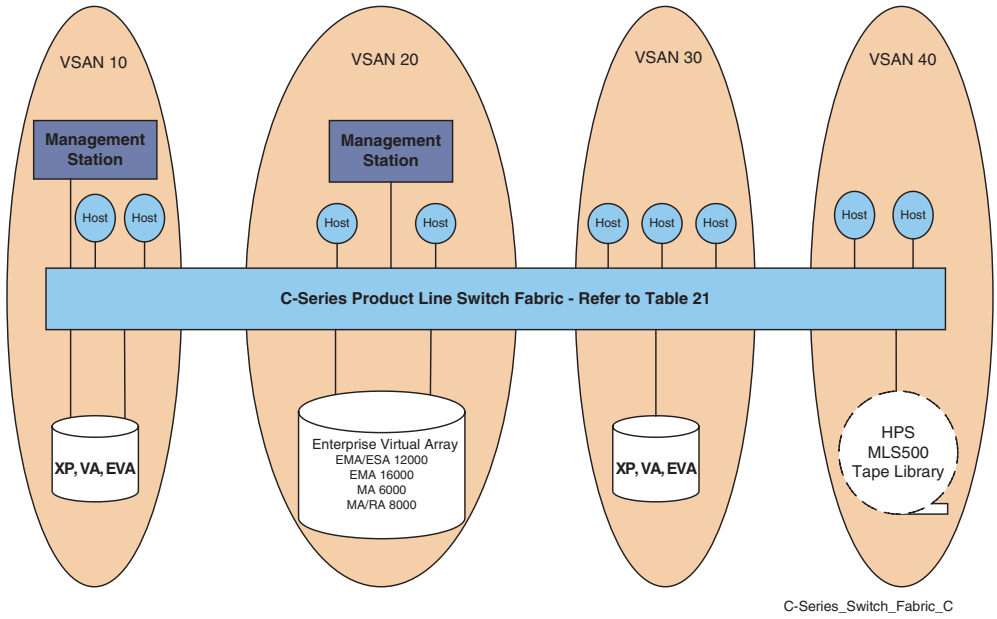

Figure 41: C-Series based SAN with VSANs

Specific platform/operating system rules – HP XP and VA storage systems

This section defines the rules and guidelines surrounding the design of SAN infrastructures for XP and VA arrays. For additional information on operating system, HBA, driver, firmware, or software support, contact your HP field representative.

These rules are subdivided into the following categories for ease of reference:

Legacy SAN support

Separate SANs are required if the environment consists of a mixture of legacy switches, devices, discontinued arrays, and current products. Legacy products include:

- Discontinued storage devices
- Discontinued switches F16 (HP first generation 16-port switch), SANBox-16 (Ancor/Qlogic)
- Storage devices, such as FC60, 12H

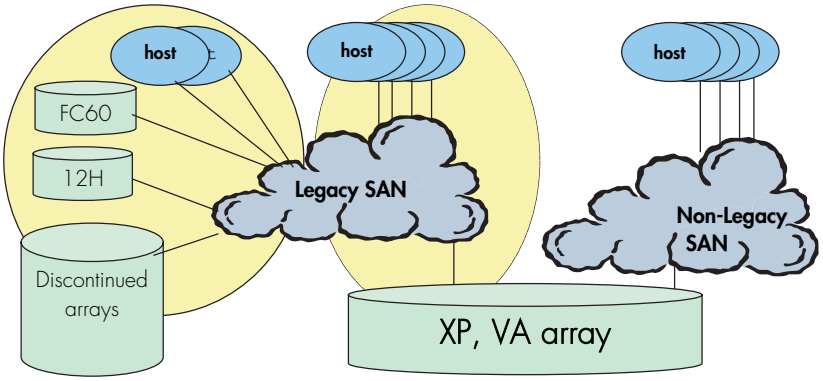

Figure 42: Legacy SAN support

High-availability/mission-critical SAN support

High-availability environments, such as HP Service Guard running on HP-UX servers, require dual fabric SAN configurations for achieving NSPOF and meeting customer expectations of "no infrastructure or application downtime" for mission-critical applications. This is true for other operating systems supporting similar high-availability solutions.

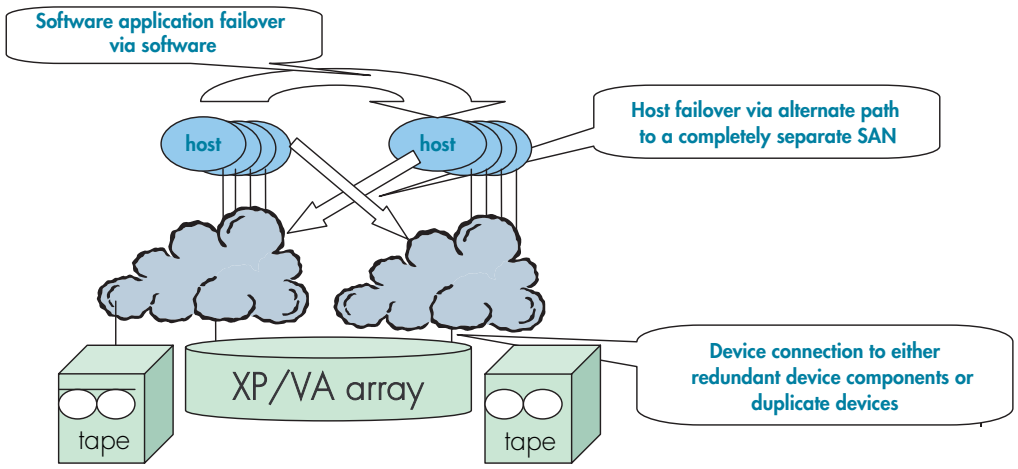

Figure 43: High availability SAN with XP/VA

SANs consisting of a single B-Series Core switch, C-Series Director switch, or M-Series Director switch can be supported for Level 3 high-availability (Figure 44). See "Data availability" on page 59, for more information.

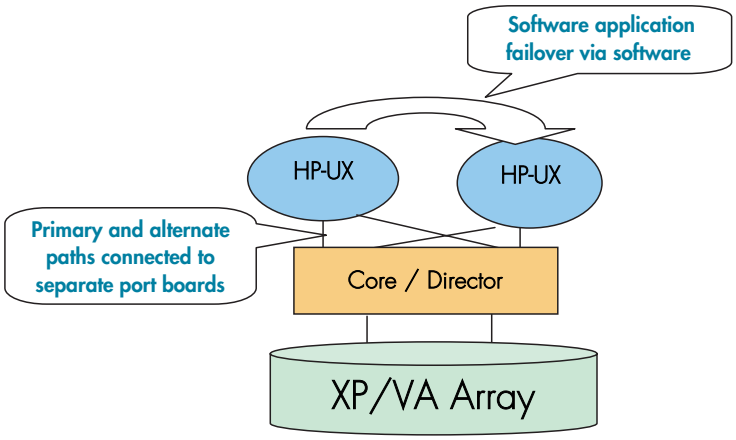

Figure 44: Software application failover

XP and VA with multiple operating systems in a shared switch fabric

- Multiple OS's and multiple clusters can be supported on the same switch/fabric with appropriate zoning:
  - Host zones must contain homogeneous operating system types.
  - Overlapping storage port zones are supported if more than one operating system needs to share an array port.
- Heterogeneous operating systems may share an XP array port with the appropriate host group/mode settings (see XP array documentation); All others must use a dedicated XP array port.
- Secure Manager XP and Secure Manager VA are required for LUN isolation with multiple hosts connected through a shared array port

Table 51: Zoning requirement for OSs sharing the same fabric with XP/VA storage

Platform OR OS type	HP-UX	Linux	Windows	Tru64 UNIX	OpenVMS	Sun Solaris	IBM AIX	Novell NetWare	SGI IRIX
HP-UX	Yes*	Zoning Required	Zoning Required	Zoning Required	Zoning Required	Zoning Required	Zoning Required	Zoning Required	Zoning Required
Linux	Zoning Required	Yes*	Zoning Required	Zoning Required	Zoning Required	Zoning Required	Zoning Required	Zoning Required	Zoning Required
Windows	Zoning Required	Zoning Required	Yes*	Zoning Required	Zoning Required	Zoning Required	Zoning Required	Zoning Required	Zoning Required
Tru64 UNIX	Zoning Required	Zoning Required	Zoning Required	Yes*	Zoning Required	Zoning Required	Zoning Required	Zoning Required	Zoning Required
OpenVMS	Zoning Required	Zoning Required	Zoning Required	Zoning Required	Yes*	Zoning Required	Zoning Required	Zoning Required	Zoning Required
Sun Solaris	Zoning Required	Zoning Required	Zoning Required	Zoning Required	Zoning Required	Yes*	Zoning Required	Zoning Required	Zoning Required
IBM AIX	Zoning Required	Zoning Required	Zoning Required	Zoning Required	Zoning Required	Zoning Required	Yes*	Zoning Required	Zoning Required
Novell NetWare	Zoning Required	Zoning Required	Zoning Required	Zoning Required	Zoning Required	Zoning Required	Zoning Required	Yes*	Zoning Required
SGI IRIX	Zoning Required	Zoning Required	Zoning Required	Zoning Required	Zoning Required	Zoning Required	Zoning Required	Zoning Required	Yes*

*Yes indicates these OS's can be part of the same zone in a fabric with XP/VA.

The above table indicates that OS types do not mix in the same zone, however, they can selectively share the storage ports that support this feature across zones. Storage ports can be overlapped in multiple zones.

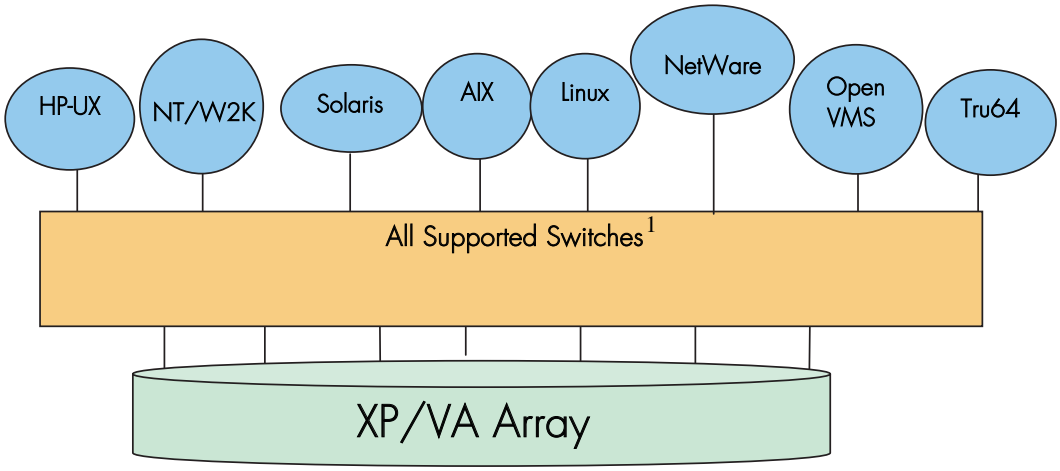

Figure 45: XP/VA with multiple OS's on a shared fabric[1]

1. Contact your HP representative for model and firmware versions for supported switches. Third party switches like ED-5000 and Inrange FC 9000 are also supported with some XP arrays, limited to a maximum two switch configurations.

XP/VA and tape with multiple OS's shared switch fabric

- Overlapping zones supported with disk and tape.
- Separate or common HBAs for disk and tape connections.
- HP recommends a dedicated tape HBA connection for servers with backups requiring more than 4 DLT8000 tape drives or 2 Ultrium (LTO) tape drives.
- Secure Manager XP and Secure Manager VA are required for LUN isolation with multiple hosts connected through a shared array port.
- Contact your HP representative for more information on tape support.

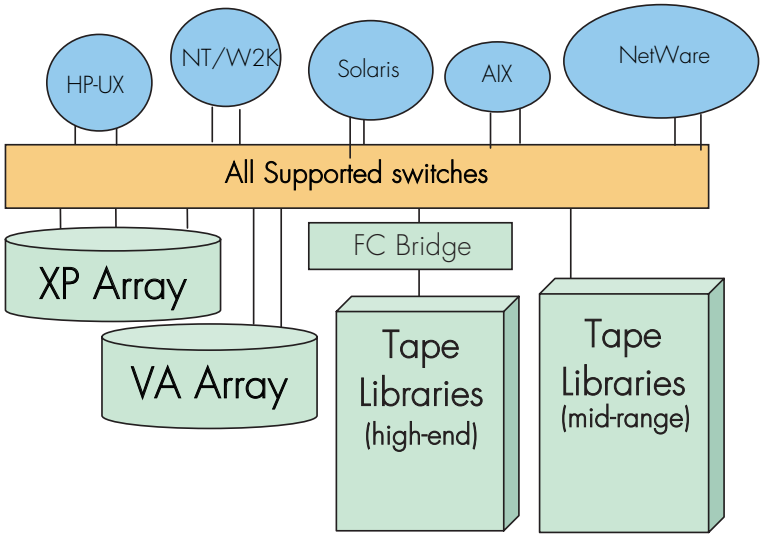

Figure 46: XP/VA with multiple OS's and tapes on a shared fabric, fabric only

Heterogeneous storage support

These configuration rules apply for heterogeneous SAN storage with the XP/VA:

- Zone the storage ports to isolate from all other storage vendor zones; no overlapping zones containing multiple storage ports
- Storage ports may be accessed from heterogeneous operating system types and multiple clusters for HA and non-HA configurations; overlapping zones are supported
- Secure Manager required for LUN isolation with multiple hosts connected through a shared array port
- Other vendor array zones governed by their vendor's configuration guidelines.
- Shared HBAs or hosts across 3rd party storage vendors are NOT supported.
- Supports connection to a common server with XP/VA and EVA storage systems. See "Common server access" on page 135.

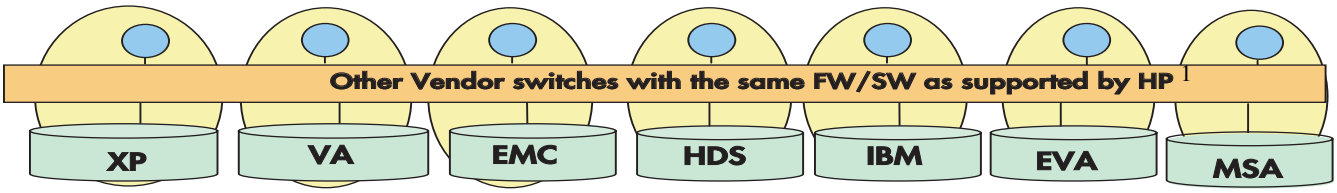

Figure 47: Heterogeneous storage support[1]

Secure manager support

HP SureStore E Secure Manager XP definition**:**

- An array based, LUN security and configuration tool
- Provides the ability to limit access between hosts and array LUNs
- Use host world-wide names to identify host access per LUN

HP SureStore E Secure Manager XP advantages:

- Provides LUN security at the array level to secure data, irrespective of switch port or direct connect between host and the XP array
- Provides consolidated and consistent data access management independent of switch vendor
- Improves boot performance during ioscan by limiting the visibility between host and targets

Note: Product cost must be compared to cost of additional array ports and switch cost requirements

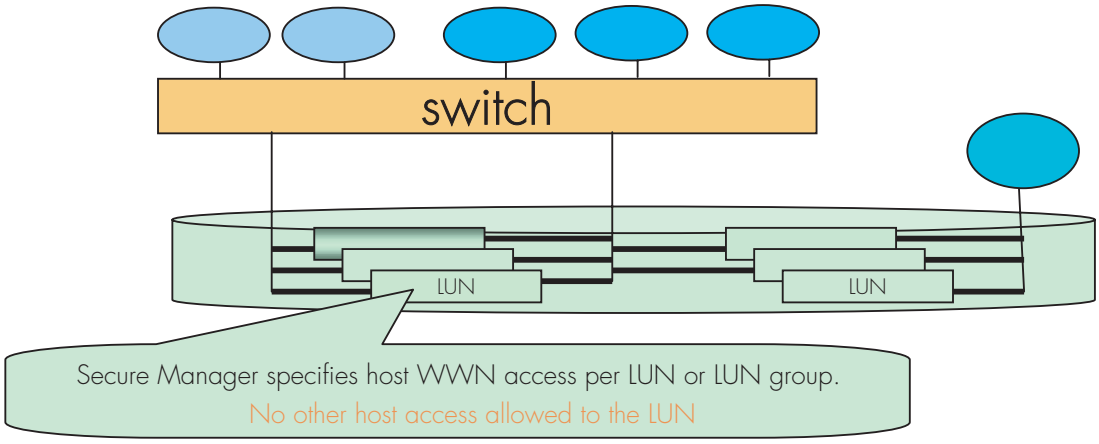

Figure 48: Secure Manager for XP support

1. There is support for third-party switches such as the ED-5000 and the Inrange FC 9000 with some XP arrays, limited to a maximum two-switch configuration. Contact your HP representative for specific details.

Note: LUNs may be shared across array ports. Therefore, limiting a host's visibility to a switch port does not limit its access to LUNs. Array-based configuration management, such as Secure Manager, is the only way to ensure data security.

Fabric boot support for XP/VA

XP and VA LUNs can be booted from the SAN using both SAN B-Series product line switches and SAN M-Series product line switches (contact your HP representative for exceptions.)

Note: SAN boot through the B-Series MP Router is not supported.

Booting from the SAN has dependencies that include PDC code, firmware, HBA, OS version/type, platform speed, and Fibre Channel port speed.

For more information, contact your HP field representative.

Table 52 describes a high level of boot support. This table does not cover different versions and flavors within each OS and switch type. A Yes in a column indicates at least one combination of array, HBA, OS type, and switch is supported as a bootable configuration.

Table 52: XP/VA SAN boot by operating system

OS/ array	HP-UX	Linux	Win-dows	Tru64 UNIX	Open-VMS	Sun Solaris	AIX	Net-Ware	SGI-IRIX	B-Series switches	M-Series switches	C-Series switches
XP-1024/128	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
XP-512/48	Yes	Yes	Yes	Yes	Yes	No	No	Yes	Yes	Yes	Yes	Yes
VA-7410	Yes	Yes	Yes	N/A	N/A	No	No	No	No	Yes	Yes	Yes
VA-7400	Yes	Yes	Yes	N/A	N/A	No	No	No	No	Yes	Yes	Yes
VA-7110	Yes	Yes	Yes	N/A	N/A	No	No	No	No	Yes	Yes	Yes
VA-7100	Yes	Yes	Yes	N/A	N/A	No	No	No	No	Yes	Yes	Yes
XP12000	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	N/A	Yes	Yes	Yes

Note: Tru64 and OpenVMS XP12000 boot requires a minimum of Alpha Server console v6.9.

Specific platform/operating system rules – EVA3000/5000 (VCS v3), EVA4000/6000/8000 (XCS v5), EMA/ESA12000, EMA16000, MA/RA8000, MA6000 (ACS 8.7, 8.8) storage systems, B-Series and M-Series switches

This section defines the rules and guidelines related to specific platforms/operating systems for EVA and EMA/ESA/MA/RA8000 storage systems, when used with B-Series and M-Series switches and the B-Series Multi-protocol Router. For operating system storage attachment, HBA attachment, current HBA/driver/FW revision support, multipathing software versions, and specific VCS minor release and ACS version patch level support, contact an HP storage representative.

HP-UX 11.0, 11iV1, 11iV2

- Zoning is required when HP-UX is used in a Heterogeneous SAN with other operating systems.
- Supports MC/ServiceGuard. Contact an HP storage representative for specific version support.

XCS v5, VCS v3

EVA3000/4000/5000/6000/8000:

- Supports Multiple-Bus Failover mode. A multipathing driver is required for Multiple-Bus failover if configured with two or more paths.
- Supports connection of single HBA servers, see the white paper white paper "Connecting Single HBA Servers to the Enterprise Virtual Array without Multipathing Software at:

  ftp://ftp.compaq.com/pub/products/storageworks/whitepapers
- Supports connection to a common server with EVA and XP/VA storage systems. See "EVA3000/4000/5000/6000/8000 configuration rules." on page 173.
- Does not support L-Port attachment.
- When using Continuous Access on EVA5000 with VCS v3.00 or EVA 3000/5000 with v3.01 and v3.02 or EVA4000/6000/8000 with XCS v5.02, HP recommends that all servers and storage systems configured for Continuous Access EVA be in a High Availability SAN of two fabrics (Figure 72), and contain only Continuous Access EVA supported operating systems. If servers running non-supported operating systems exist within the SAN, they should be excluded from the Continuous Access management environment via zoning or equivalent partitioning tools using virtual fabrics. In addition, that SAN (or management zone within a larger SAN) is limited to a maximum of 16 EVAs. Additional operating system specific zones within the larger Continuous Access environment zone may be required as defined in Table 53. See the *Continuous Access EVA Planning Guide* and the *Continuous Access EVA Release Notes* for additional details and configuration limitations.
- See the Continuous Access EVA Planning Guide and the Continuous Access EVA Release Notes for additional details and configuration limitations.
- 256 HBAS per EVA. Additional operating system specific zones within the larger Continuous Access environment zone may be required as defined in Table 53. See the *Continuous Access EVA Planning Guide* for additional details and configuration limitations.

ACS 8.7, 8.8 - HP-UX 11.0, 11iV1

EMA/ESA12000, EMA16000, MA/RA8000, and MA6000 storage systems:

- Supports Transparent failover mode and Multiple-Bus failover mode. Multiple-Bus failover mode is supported for HP-UX version 11.0 and 11iV1 using the Secure Path multipathing driver. Supports multipathing high-availability configuration implemented in separate fabrics or in a single fabric with zoned paths.
- Supports L-Port attachment with Fibre Channel SAN switches that support the QuickLoop feature (Table 10).
- Requires ACS 8.7P or later for DRM support,
- Requires all servers and storage systems configured for DRM to be in a zone or group of zones that excludes all non DRM supported operating systems. Multiple DRM zones are supported to reduce the size and/or complexity of a particular DRM solution instance. Multiple zones may be required due to multiple operating systems as defined in Table 53.

OpenVMS

- Supports Multiple-Bus failover mode. Multipathing driver is embedded in the operating system
- Supports multipathing high-availability configuration implemented in separate fabrics or a single fabric
- Zoning required when used in a Heterogeneous SAN with HP-UX, IBM AIX or Linux
- Supports booting over the SAN Fabric. See "Booting from the SAN" on page 169 for more information.

XCS v5 - OpenVMS 7.3-2
VCS v3 - OpenVMS 7.3-2, 8.2 (Alpha), 8.2 (i64)

- EVA3000/4000/5000/6000/8000
- Supports connection of single HBA servers, see the white paper "Connecting Single HBA Servers to the Enterprise Virtual Array without Multipathing Software" (SingleHBAforEVA_F.pdf) at:

  ftp://ftp.compaq.com/pub/products/storageworks/whitepapers
- When using Continuous Access on EVA5000 with VCS v3.00 or EVA 3000/5000 with v3.01 and v3.02 or EVA4000/6000/8000 with XCS v5.02, HP recommends that all servers and storage systems configured for Continuous Access EVA be in a High Availability SAN of two fabrics (Figure 72), and contain only Continuous Access EVA supported operating systems. If servers running non-supported operating systems exist within the SAN, they should be excluded from the Continuous Access management environment via zoning or equivalent partitioning tools using virtual fabrics. In addition, that SAN (or management zone within a larger SAN) is limited to a maximum of 16 EVAs. Additional operating system specific zones within the larger Continuous Access environment zone may be required as defined in Table 53. See the *Continuous Access EVA Planning Guide* and the *Continuous Access EVA Release Notes* for additional details and configuration limitations.

ACS 8.7, 8.8

EMA/ESA12000, EMA16000, MA/RA8000, MA6000 Storage Systems:

- All servers and storage systems configured for DRM must be in a zone or group of zones that excludes all non DRM supported operating systems. Multiple DRM zones are supported to reduce the size and/or complexity of a particular DRM solution instance. Multiple zones may be required due to multiple operating systems as defined in Table 53 on page 163.

Tru64 UNIX

- Supports multipathing high-availability configuration implemented in separate fabrics or a single fabric
- Zoning is required when used in a Heterogeneous SAN with HP-UX, IBM AIX or Linux
- Supports booting over the Fabric. See "Booting from the SAN" on page 169 for more information.

VCS v3 – Tru64 UNIX 5.1, 5.1A, 5.1B

EVA3000/5000:

- Supports Multiple-Bus Failover mode. The Multipathing driver is embedded in the operating systems.
- Supports connection of single HBA servers, see the white paper "Connecting Single HBA Servers to the Enterprise Virtual Array without Multipathing Software" at:

  ftp://ftp.compaq.com/pub/products/storageworks/whitepapers
- When using Continuous Access on EVA5000 with VCS v3.00 or EVA 3000/5000 with v3.01 and v3.02, HP recommends that all servers and storage systems configured for Continuous Access EVA be in a High Availability SAN of two fabrics (Figure 72), and contain only Continuous Access EVA supported operating systems. If servers running non-supported operating systems exist within the SAN, they should be excluded from the Continuous Access management environment via zoning or equivalent partitioning tools using virtual fabrics. In addition, that SAN (or management zone within a larger SAN) is limited to a maximum of 16 EVAs. Additional operating system specific zones within the larger Continuous Access environment zone may be required as defined in Table 53. See the *Continuous Access EVA Planning Guide* and the *Continuous Access EVA Release Notes* for additional details and configuration limitations.

ACS 8.7, 8.8 – Tru64 UNIX 4.0F, 4.0G, 5.1, 5.1A, 5.1B

EMA/ESA12000, EMA16000, MA/RA8000, MA6000 Storage Systems:

- Tru64 UNIX version 4.0F/4.0G supports Transparent failover mode only.
- Tru64 UNIX 5.1and 5.1A support Transparent and Multiple-Bus failover mode. A multipathing driver is embedded in the V5.1/V5.1A/V5.1B operating systems.
- Zoning is required when a SAN is configured for multiple TruCluster products with Tru64 UNIX 4.0F/4.0G. Each TruCluster configured with Tru64 UNIX 4.0F/4.0G must be in its own zone.
- All servers and storage systems configured for DRM (supported on Tru64 5.1, 5.1A, and 5.1B only), must be in a zone or group of zones that excludes all non DRM supported operating systems. Multiple DRM zones are supported to reduce the size and/or complexity of a particular DRM solution instance. Multiple zones may be required due to multiple operating systems as defined in Table 53 on page 163.

IBM AIX

- Supports HACMP/ES Clusters. Contact an HP storage representative for specific version support.
- Requires zoning when used in a Heterogeneous SAN with other operating systems.
- Supports Multiple-Bus Failover mode. Secure Path multipathing driver is required for Multiple-Bus failover. Supports multipathing high availability configuration implemented in separate fabrics or in a single fabric with zoned paths.

XCS v5 - AIX 5.2, 5.3
VCS v3 - AIX 4.3.3, 5.1, 5.2, 5.3

EVA3000/4000/5000/6000/8000

- When using Continuous Access on EVA5000 with VCS v3.00 or EVA 3000/5000 with v3.01 and v3.02 or EVA4000/6000/8000 with XCS v5.02, HP recommends that all servers and storage systems configured for Continuous Access EVA be in a High Availability SAN of two fabrics (Figure 72), and contain only Continuous Access EVA supported operating systems. If servers running non-supported operating systems exist within the SAN, they should be excluded from the Continuous Access management environment via zoning or equivalent partitioning tools using virtual fabrics. In addition, that SAN (or management zone within a larger SAN) is limited to a maximum of 16 EVAs. Additional operating system specific zones within the larger Continuous Access environment zone may be required as defined in Table 53. See the *Continuous Access EVA Planning Guide* and the *Continuous Access EVA Release Notes* for additional details and configuration limitations.

ACS 8.7, 8.8

For EMA/ESA12000, EMA16000, MA/RA8000, MA6000 Storage Systems:

- Supports Transparent and Multiple-Bus failover mode. Secure Path multipathing driver is required for Multiple-Bus failover.
- All servers and storage systems configured for DRM must be in a zone or group of zones that excludes all non DRM supported operating systems. Multiple DRM zones are supported to reduce the size and/or complexity of a particular DRM solution instance. Multiple zones may be required due to multiple operating systems as defined in Table 53 on page 163.

Secure Path for IBM AIX

- When using Multiple-Bus Failover and Secure Path for IBM AIX, zoning is required to limit each IBM server HBA access to one controller port per controller because typical installations in a heterogeneous SAN utilize more then one controller port cable connection per controller. See the *Secure Path for IBM AIX Installation and Reference Guide,* for more information. Zoning is required for the AIX servers for instances where the storage system is being shared with other heterogeneous servers that require cabled access to more than one controller port per controller.

Linux

XCS v5 - Red Hat EL 2.1, EL 3 (32-bit, 64-bit), EL 4, SLES 9
VCS v3 - Red Hat 7.2, EL 3 Advanced Server 2.1, SuSE SLES 7, SLES 8, United Linux 1.0

EVA3000/4000/5000/6000/8000:

- When using Continuous Access on EVA5000 with VCS v3.00 or EVA 3000/5000 with v3.01 and v3.02 or EVA4000/6000/8000 with XCS v5.02, HP recommends that all servers and storage systems configured for Continuous Access EVA be in a high-availability SAN of two fabrics (Figure 72), and contain only Continuous Access EVA supported operating systems. If servers running non-supported operating systems exist within the SAN, they should be excluded from the Continuous Access management environment via zoning or equivalent partitioning tools using virtual fabrics. In addition, that SAN (or management zone within a larger SAN) is limited to a maximum of 16 EVAs. Additional operating system specific zones within the larger Continuous Access environment zone may be required as defined in Table 53. See the *Continuous Access EVA Planning Guide* and the *Continuous Access EVA Release Notes* for additional details and configuration limitations.
- Supports Lifekeeper Clusters. Contact an HP storage representative for specific version support.
- Supports connection of single HBA servers, see the white paper "Connecting Single HBA Servers to the Enterprise Virtual Array without Multipathing Software" at: http://h18006.www1.hp.com/storage/arraywhitepapers.html
- Zoning is required when used in a Heterogeneous SAN with other operating systems.

ACS 8.7, 8.8 - Red Hat 7.2, Advanced Server 2.1, 7.1, 7.2 (Alpha Server), SuSE 7.2, SuSE SLES 7

EMA/ESA12000, EMA16000, MA/RA8000, MA6000 Storage Systems:

- Supports Lifekeeper Clusters. Contact an HP storage representative for specific version support.
- Supports Transparent failover mode.
- Requires zoning when used in a Heterogeneous SAN with other operating systems.
- Supports connection of single HBA servers, see the white paper "Connecting Single HBA Servers to the Enterprise Virtual Array without Multipathing Software" at:

  http://h18006.www1.hp.com/storage/arraywhitepapers.html
- Multipathing with Secure Path is supported on Red Hat Advanced Server 2.1 and SuSE SLES 7 only.

ACS 8.7, 8.8 - Secure Path for Linux, Red Hat Advanced Server 2.1, SLES 7

- Supports Multiple-Bus failover mode.
- Zoning is required when used in a Heterogeneous SAN with other operating systems.

Microsoft Windows 2000 Server, Advanced Server w/SP2, SP3, SP4 for VCS3.x only, Windows NT 4.0 w/SP6a, Windows 2003 Server (32-bit, 64-bit)

- Supports MSCS for EMA/ESA12000, EMA16000, and MA/RA8000 system in a 2-node configuration only. Contact an HP storage representative for specific version support.
- Supports multipathing high-availability configuration implemented in separate fabrics or in a single fabric with zoned paths.
- Requires zoning when used in a heterogeneous SAN with HP-UX, IBM AIX or Linux.
- Supports BL server models with Windows 2000 and Windows 2003 only.
- Supports EVA4000/6000/8000 on Windows 2003 and Windows 2000, SP4 only.

Note: In an environment with multiple Windows 2000 or Windows 2003 clusters sharing a storage array, HP recommends that each cluster and its storage be configured in a separate zone. If the same storage array is used by multiple clusters, the storage ports may be in an overlapping zone.

XCS v5, VCS v3

EVA3000/4000/5000/6000/8000:

- Supports Multiple-Bus Failover mode. Secure Path multipathing driver is required for Multiple-Bus failover if configured with two or more paths.
- Supports connection of single HBA servers, see the white paper "Connecting Single HBA Servers to the Enterprise Virtual Array without Multipathing Software at:

  http://h18004.www1.hp.com/products/storageworks/enterprise/documentation.html
- When using Continuous Access on EVA5000 with VCS v3.00 or EVA 3000/5000 with v3.01 and v3.02 or EVA4000/6000/8000 with XCS v5.02, HP recommends that all servers and storage systems configured for Continuous Access EVA be in a high-availability SAN of two fabrics (Figure 72), and contain only Continuous Access EVA supported operating systems. If servers running non-supported operating systems exist within the SAN, they should be excluded from the Continuous Access management environment via zoning or equivalent partitioning tools using virtual fabrics. In addition, that SAN (or management zone within a larger SAN) is limited to a maximum of 16 EVAs. Additional operating system specific zones within the larger Continuous Access environment zone may be required as defined in Table 53. See the *Continuous Access EVA Planning Guide* and the *Continuous Access EVA Release Notes* for additional details and configuration limitations.

ACS 8.7, 8.8

EMA/ESA12000, EMA16000, MA/RA8000, MA6000 Storage Systems:

- Supports MS Windows 2003 Enterprise edition (32 bit) and Standard edition (32 bit) with Multiple-Bus failover only.
- Supports Transparent failover mode and Multiple-Bus failover mode. Secure Path multipathing driver is required for Multiple-Bus failover.
- All servers and storage systems configured for DRM must be in a zone or group of zones that excludes all non DRM supported operating systems. Multiple DRM zones are supported to reduce the size and/or complexity of a particular DRM solution instance. Multiple zones may be required due to multiple operating systems as defined in Table 53 on page 163.
- Supports booting over the Fabric. See "Booting from the SAN" on page 169 for more information.
- Extended Configurations with Microsoft Windows NT 4.0.

  If you configure greater than 4 (up to 8) servers (assuming one Fibre Channel HBA per server) for access to a single controller host port on an MA6000, MA/RA8000, EMA/ESA12000, and EMA16000 storage system, and 1 or more of those servers is Windows based, select the "Extended Configuration" check box in the StorageWorks Windows NT/Windows 2000 Platform Kit Fibre Channel Software Setup utility custom installation setup for each Windows server. Select this option to adjust registry settings for your KGPSA host bus adapter to operate in an "Extended Configuration" environment.

Note: The default for this option is checked, so be sure to uncheck this option when you have 4 or fewer servers configured for access to a single controller host port.

Microsoft Windows 2000 Datacenter

- Supports MSCS. Contact an HP storage representative for specific version support.
- Supports multipathing high-availability configuration implemented in separate fabrics or in a single fabric with zoned paths.
- Requires zoning when used in a heterogeneous SAN with HP-UX, IBM AIX or Linux.

VCS v3

EVA3000/5000:

- VCS v3 is supported on WS2003 32/64 bit
- Supports Multiple-Bus Failover mode. Secure Path multipathing driver is required for Multiple-Bus failover.
- When using Continuous Access on EVA5000 with VCS v3.00 or EVA 3000/5000 with v3.01 and v3.02, HP recommends that all servers and storage systems configured for Continuous Access EVA be in a high-availability SAN of two fabrics (Figure 72), and contain only Continuous Access EVA supported operating systems. If servers running non-supported operating systems exist within the SAN, they should be excluded from the Continuous Access management environment via zoning or equivalent partitioning tools using virtual fabrics. In addition, that SAN (or management zone within a larger SAN) is limited to a maximum of 16 EVAs. Additional operating system specific zones within the larger Continuous Access environment zone may be required as defined in Table 53. See the *Continuous Access EVA Planning Guide* and the *Continuous Access EVA Release Notes* for additional details and configuration limitations.

ACS 8.7, 8.8

EMA/ESA12000, EMA16000, MA/RA8000, MA6000 Storage Systems:

- Supports Transparent failover mode and Multiple-Bus failover mode. Secure Path multipathing driver is required for Multiple-Bus failover.
- Supports heterogeneous operating system shared access to a single storage system when using ACS 8.7. Heterogeneous operating system shared access is not supported with ACS 8.6.

Secure Path for Windows

EMA/ESA12000, EMA16000, MA/RA8000, MA6000 Storage Systems:

- Supports using single host bus adapter with Secure Path and Multiple-Bus failover. See the white paper, *Using Secure Path for Servers with Single Host Bus Adapter (HBA)*, 14JK-0301A-WWEN, available at:

  http://h18006.www1.hp.com/products/sanworks/library/whitepapers/14JK-0301A-WWEN.html
- One instance of the Secure Path Manager can support multiple managed entities called profiles. For 4.x versions of Secure Path, a single profile can consist of up to 128 servers total, standalone or clustered, connected to and sharing up to 128 storage systems. For 3.x versions of Secure Path, a single profile can consist of up to 8 standalone servers

connected to and sharing up to 8 storage systems, or up to 8 clustered servers connected to and sharing up to 8 storage systems. For 3.x versions, you cannot manage both standalone and clustered servers in the same profile.

- Secure Path configurations utilizing 4 active controller ports connected to the same server or servers offer the flexibility to use the 4 active ports for either increased total LUN count, or increased PATH accessibility to a lesser number of LUNs. See "High-availability configuration considerations" on page 199 for more information.
- Provides for dynamic port I/O load distribution in non-clustered servers when configured for maximum paths.
- Distribute units equally across both controllers for proper static load balancing using the Unit Preferred Path parameter to assign units to a specific controller at initial boot.
- SSP/LUN level masking - Storagesets (LUNs) must be enabled for access from all server or clustered server paths using the storage LUN presentation or Unit Connection Name parameter feature.
- For Windows NT or Windows 2000, when using Secure Path in single or dual fabric configurations with both Multiple-bus Failover and Transparent Failover storage systems, the Transparent Failover storage systems must be in a different fabric zone and not be accessed by servers running Secure Path multipathing software.

Novell NetWare

- Supports NetWare Clusters. Contact an HP storage representative for specific version support.
- Zoning required when used in a Heterogeneous SAN with Sun, HP-UX, IBM AIX, or Linux.

VCS v3 – NetWare 5.1, 6, 6.5

EVA3000/5000:

- Supports Multiple-Bus Failover mode. Secure Path multipathing driver is required for Multiple-Bus failover if configured with two or more paths. Multiple-Bus Failover configurations support clusters with a maximum of 12 nodes.
- Supports connection of single HBA servers, see the white paper "Connecting Single HBA Servers to the Enterprise Virtual Array without Multipathing Software" at:

  http://h18006.www1.hp.com/storage/arraywhitepapers.html
- When using Continuous Access on EVA5000 with VCS v3.00 or EVA 3000/5000 with v3.01 and v3.02, HP recommends that all servers and storage systems configured for Continuous Access EVA be in a high-availability SAN of two fabrics (Figure 72), and contain only Continuous Access EVA supported operating systems. If servers running non-supported operating systems exist within the SAN, they should be excluded from the Continuous Access management environment via zoning or equivalent partitioning tools using virtual fabrics. In addition, that SAN (or management zone within a larger SAN) is limited to a maximum of 16 EVAs. Additional operating system specific zones within the larger Continuous Access environment zone may be required as defined in Table 53. See the *Continuous Access EVA Planning Guide* and the *Continuous Access EVA Release Notes* for additional details and configuration limitations.

ACS 8.7, 8.8 – NetWare 4.2

EMA/ESA12000, EMA16000, MA/RA8000, MA6000 Storage Systems supported using Transparent failover mode only.

ACS 8.7, 8.8 - NetWare 5.1 SP7, 6.0 SP4 and 6.5 SP1.1

EMA/ESA12000, EMA16000, MA/RA8000, MA6000 Storage Systems:

- Supports NetWare Clusters.
- Supports Transparent failover mode and Multiple-Bus failover mode. Secure Path multipathing driver is required for Multiple-Bus failover.
- Supports multipathing high-availability configuration implemented in separate fabrics or in a single fabric with zoned paths.
- All servers and storage systems configured for DRM must be in a zone or group of zones that excludes all non DRM supported operating systems. Multiple DRM zones are supported to reduce the size and/or complexity of a particular DRM solution instance. Multiple zones may be required due to multiple operating systems, as defined inTable 53 on page 163.

Sun Solaris

- Supports Sun Clusters, Veritas Clusters, and VxVM Veritas clusters. Contact an HP storage representative for specific version support.
  - — Each cluster must be in its own zone.
- Zoning required when used in a Heterogeneous SAN with NetWare, HP-UX, IBM AIX, or Linux
- Supports multipathing high-availability configuration implemented in separate fabrics or in a single fabric with zoned paths.

XCS v5 - Sun Solaris 8, 9
VCS v3 - Sun Solaris 2.6, 7, 8, 9

EVA3000/4000/5000/6000/8000:

- Supports Multiple-Bus Failover mode. Secure Path multipathing driver is required for Multiple-Bus failover if configured with two or more paths.
- Supports connection of single HBA servers, see the white paper “Connecting Single HBA Servers to the Enterprise Virtual Array without Multipathing Software” at:

  ftp://ftp.compaq.com/pub/products/storageworks/whitepapers
- Supports Sun Clusters and Veritas Clusters. Contact an HP storage representative for specific version support.
  - — A cluster must be in its own zone.
- When using Continuous Access on EVA5000 with VCS v3.00 or EVA 3000/5000 with v3.01 and v3.02 or EVA4000/6000/8000 with XCS v5.02, HP recommends that all servers and storage systems configured for Continuous Access EVA be in a high-availability SAN of two fabrics (Figure 72), and contain only Continuous Access EVA supported operating systems. If servers running non-supported operating systems exist within the SAN, they should be excluded from the Continuous Access management environment via zoning or equivalent partitioning tools using virtual fabrics. In addition, that SAN (or management zone within a larger SAN) is limited to a maximum of 16 EVAs. Additional operating

system specific zones within the larger Continuous Access environment zone may be required as defined in Table 53. See the *Continuous Access EVA Planning Guide* and the *Continuous Access EVA Release Notes* for additional details and configuration limitations.

ACS 8.7, 8.8 – Sun Solaris 2.6, 7, 8

EMA/ESA12000, EMA16000, MA/RA8000, MA6000 Storage Systems:

- Supports Transparent failover mode and Multiple-Bus failover mode. Secure Path multipathing driver is required for Multiple-Bus failover.
- Supports Sun Clusters and Veritas Clusters. Contact an HP storage representative for specific version support.
  - A cluster must be in its own zone.

All servers and storage systems configured for DRM must be in a zone or group of zones that excludes all non DRM supported operating systems. Multiple DRM zones are supported to reduce the size and/or complexity of a particular DRM solution instance. Multiple zones may be required due to multiple operating systems, as defined in Table 53 on page 163.

Specific platform/operating system rules – EVA3000/5000 (VCS v3), EVA4000/6000/8000 (XCS v5), EMA/ESA12000, EMA16000, MA/RA8000, MA6000 (ACS 8.7, 8.8) storage systems, C-Series switches

This section defines the rules and guidelines related to specific platforms/operating systems for EVA and EMA/ESA/MA/RA8000 storage systems, when used with C-Series switches and C-Series Fibre Channel switch inter-VSAN routing functionality. For supported server models, cluster version support, operating system storage attachment, HBA attachment, current HBA/driver/FW revision support, multipathing software versions, and specific VCS and ACS version patch level support, see you HP representative for more information.

HP-UX 11.0, 11iV1, 11iV2

- Zoning is required when used in a Heterogeneous SAN with other operating systems.
- Service Guard Clusters supported

XCS v5, VCS v3

EVA3000/4000/5000/6000/8000:

- Supports Multiple-Bus Failover mode. A multipathing driver is required for Multiple-Bus failover.

ACS 8.7, 8.8

- Supports Transparent failover mode and Multiple-Bus failover mode. Secure Path multipathing driver is required for Multiple-Bus failover.

Microsoft Windows 2000 Server, Advanced Server SP3, SP4, Windows 2003 (32-bit, 64-bit)

- Zoning is required when used in a Heterogeneous SAN with other operating systems.

XCS v5, VCS v3

EVA3000/4000/5000/6000/8000:

- Supports Multiple-Bus Failover mode. Secure Path multipathing driver is required for Multiple-Bus failover.
- Supports Microsoft Cluster Server.
- Supports EVA4000/6000/8000 on Windows 2003 and Windows 2000, SP4 only.

ACS 8.7, 8.8

- Supports Transparent failover mode and Multiple-Bus failover mode. Secure Path multipathing driver is required for Multiple-Bus failover.
- Microsoft Cluster Server supported (32-bit only)

OpenVMS

XCS v5 - OpenVMS 7.3-2
VCS v3 - OpenVMS 7.3-2, 8.2 (Alpha), 8.2 (i64)

EVA3000/4000/5000/6000/8000:

- Supports Multiple-Bus Failover mode
- OpenVMS Clusters supported

ACS 8.7, 8.8

- Supports Multiple-Bus failover mode
- OpenVMS Clusters supported

Tru64 UNIX 5.1A, 5.1B

- Zoning is required when used in a Heterogeneous SAN with other operating systems.

VCS v3 - Tru64 UNIX 5.1A, 5.1B

EVA3000/5000:

- Supports Multiple-Bus Failover mode
- TruCluster Server supported

ACS 8.7, 8.8

- Supports Multiple-Bus failover mode
- TruCluster Server supported

IBM AIX

- Zoning is required when used in a Heterogeneous SAN with other operating systems.

XCS v5 - AIX 5.2, 5.3
VCS v3 - AIX 4.3.3, 5.1, 5.2, 5.3

EVA3000/4000/5000/6000/8000:

- Supports Multiple-Bus Failover mode
- Supports HACMP/ES Clusters

ACS 8.7, 8.8

- Supports Multiple-Bus failover mode
- Supports HACMP/ES Clusters

Linux

- Zoning is required when used in a Heterogeneous SAN with other operating systems

XCS v5 - Red Hat EL 2.1, EL 3 (32-bit, 64-bit), EL 4, SLES 9
VCS v3 - Red Hat AS 2.1 (32-bit, 64-bit), SuSE 8, 9 (32-bit, 64-bit)

EVA3000/4000/5000/6000/8000:

- Supports Multiple-Bus Failover mode
- Supports Clusters

ACS 8.7, 8.8

- Supports Multiple-Bus failover mode
- Supports Clusters

Sun Solaris

- Zoning is required when used in a Heterogeneous SAN with other operating systems

XCS v5 - Sun Solaris 8, 9
VCS v3 - Sun Solaris 2.6, 7, 8, 9

EVA3000/4000/5000/6000/8000:

- Supports Multiple-Bus failover mode
- Supports Veritas Cluster Server and Sun Clusters

ACS 8.7, 8.8

- Supports Multiple-Bus failover mode
- Supports Veritas Cluster Server and Sun Clusters

Novell NetWare

- Zoning is required when used in a Heterogeneous SAN with other operating systems.

VCS v3 - NetWare 5.1, 6, 6.5

EVA3000/5000:

- Supports Multiple-Bus failover mode
- Supports Novell Cluster Services

Novell NetWare 4.2, 5.1, 6, 6.5

ACS 8.7, 8.8

- Supports Multiple-Bus failover mode
- Supports Novell Cluster Services

Specific platform/operating system rules – XP128/1024, XP48/512, XP12000 and C-Series switches

This section defines the rules and guidelines related to specific platforms/operating systems for XP128/1024, XP48/512, and XP12000 storage systems, when used with C-Series switches and C-Series Fibre Channel switch inter-VSAN routing functionality. For operating system storage attachment, HBA attachment, current HBA/driver/FW revision support, multipathing software versions, and specific XP version patch level support, see you HP representative for more information.

HP-UX 11.0, 11iV1, 11iV2

- Zoning is required when used in a Heterogeneous SAN with other operating systems.
- Supports Service Guard clusters.
- XP256 supported (1.2.1a and 1.2.1b FW only).

Red Hat Linux AS/ES 2.1 (32-bit, 64-bit), AS/ES/WS 3 (32-bit, 64-bit), SuSE Enterprise Server 7 (32-bit), 8 and 9 (32-bit, 64-bit)

- Zoning is required when used in a Heterogeneous SAN with other operating systems.
- Single-path support only.
- Supports clusters (32-bit, 64-bit).

Windows Server 2003 32-bit Enterprise and Standard Edition 64-bit Datacenter and Enterprise Edition, 2000 with SP3, SP4

- Zoning is required when used in a Heterogeneous SAN with other operating systems.
- Supports Microsoft Cluster Server

Sun Solaris 2.6, 7, 8, 9

- Zoning is required when used in a Heterogeneous SAN with other operating systems.
- Supports Veritas Cluster Server and Sun Clusters.

IBM AIX 4.3.3, 5.1, 5.2

- Zoning is required when used in a Heterogeneous SAN with other operating systems.
- Supports HACMP clusters.

OpenVMS 7.3-1, 7.3-2, 7.2-2. 8.2

- Zoning is required when used in a Heterogeneous SAN with other operating systems.
- Supports OpenVMS Clusters.

Tru64 UNIX 5.1A, 5.1B

- Zoning is required when used in a Heterogeneous SAN with other operating systems.
- Supports TruCluster Server.

Novell NetWare 5.1, 6, 6.5

- Zoning is required when used in a Heterogeneous SAN with other operating systems.
- Supports with Novell Cluster Services.

Specific platform/operating system rules – VA7400, VA7410, VA7100, VA7110, C-Series switches

This section defines the rules and guidelines related to specific platforms/operating systems for the VA7400,VA7410, VA7100, and VA7110 storage systems, when used with C-Series switches and C-Series Fibre Channel switch inter-VSAN routing functionality. For operating system storage attachment, HBA attachment, current HBA/driver/FW revision support, multipathing software versions, and specific VA version patch level support, see you HP representative for more information.

HP-UX 11.00, 11iV1, 11iV2

- Zoning is required when used in a Heterogeneous SAN with other operating systems.
- Supports Service Guard clusters.

Red Hat Linux Red Hat Advanced Server 2.1, AS/ES/WS 3 (32-bit, 64-bit), SuSE Enterprise Server 7 (i386), SuSE 8 and 9 (32-bit, 64-bit),

- Zoning is required when used in a Heterogeneous SAN with other operating systems.
- Supports clusters (32-bit, 64-bit).

Windows 2000 Server, Advanced Server SP3, SP4, Windows 2003 (32-bit, 64-bit), NT SP6a

- Zoning is required when used in a Heterogeneous SAN with other operating systems.
- Supports Microsoft Cluster Server.

Novell NetWare 5.1, 6, 6.5

- Zoning is required when used in a Heterogeneous SAN with other operating systems.
- Supports Novell Cluster Services.

Heterogeneous SAN platform interoperability for EVA3000/4000/5000/6000/8000 and EMA/ESA12000, EMA16000, MA/RA8000, MA6000 storage systems

For the EVA3000/4000/5000/6000/8000 and EMA/ESA12000, EMA16000, MA/RA8000, and MA6000 heterogeneous SAN platform interoperability is defined in Table 53 on page 163. A *Yes* in the table indicates that the listed platforms can be configured for shared access to the same storage system. *Zoning Required* indicates the platforms listed must be configured in different fabric zones in order to co-exist in the same physical SAN or share the same EVA3000/4000/5000/6000/8000 or EMA/ESA12000, EMA16000, MA/RA8000, MA6000. For C-Series switches, each operating system type must be in a separate zone or a separate VSAN.

For EMA/ESA12000, EMA16000, MA/RA8000, MA6000 storage systems there are three levels of heterogeneous interoperability rules: platform zoning rules, controller SCSI-modes, and controller failover modes.

The platform zoning rules define which platforms or operating systems must be in different fabric zones in order to coexist in the same physical SAN. See Table 53 on page 163.

The controller SCSI-mode and controller failover rules define which platforms or operating systems can be configured for shared access to a single shared storage system based on controller SCSI-mode and failover mode compatibility. See Table 54 on page 164, and Table 55 on page 165 for details.

Table 56 on page 166 combines (and to some extent repeats) the information from the other tables into a single table that can be quickly referenced to determine the settings and rules for mixing all possible combinations of any two platforms.

Platform zoning rules

Table 53 summarizes the zone compatibility for different platforms in a SAN using B-Series or M-Series switches. Platforms in the same columns can coexist in the same zone. For C-Series switches, each operating system type must be a separate zone. If a SAN Appliance resides in the fabric, it must be configured in a separate zone from all operating systems. See "Storage Management Appliance rules and recommendations"on page 176 for more information.

Table 53: SAN/Platform Zoning Requirements for EVA3000/4000/5000/6000/8000 and EMA/ESA12000, EMA16000, MA/RA8000, MA6000 Storage Systems (B-Series and M-Series switches)

Platform or operating system	HP-UX	OpenVMS	Tru64 UNIX	IBM AIX	Linux	Microsoft Windows	Novell NetWare	Sun Solaris
HP-UX	Yes	Zoning Required	Zoning Required	Zoning Required	Zoning Required	Zoning Required	Zoning Required	Zoning Required
OpenVMS	Zoning Required	Yes	Yes	Zoning Required	Zoning Required	Yes	Yes	Yes
Tru64 UNIX	Zoning Required	Yes	Yes	Zoning Required	Zoning Required	Yes	Yes	Yes
IBM AIX	Zoning Required	Zoning Required	Zoning Required	Yes	Zoning Required	Zoning Required	Zoning Required	Zoning Required
Linux	Zoning Required	Zoning Required	Zoning Required	Zoning Required	Yes	Zoning Required	Zoning Required	Zoning Required

Table 53: SAN/Platform Zoning Requirements for EVA3000/4000/5000/6000/8000 and EMA/ESA12000, EMA16000, MA/RA8000, MA6000 Storage Systems (B-Series and M-Series switches) (Continued)

Platform or operating system	HP-UX	OpenVMS	Tru64 UNIX	IBM AIX	Linux	Microsoft Windows	Novell NetWare	Sun Solaris
Microsoft Windows	Zoning Required	Yes	Yes	Zoning Required	Zoning Required	Yes	Yes	Zoning Required
Novell NetWare	Zoning Required	Yes	Yes	Zoning Required	Zoning Required	Yes	Yes	Zoning Required
Sun Solaris	Zoning Required	Yes	Yes	Zoning Required	Zoning Required	Zoning Required	Zoning Required	Yes

Note: The following table summarizes:

- NetWare and Sun platforms are incompatible in the same zone.
- HP-UX, IBM AIX, and Linux platforms are each incompatible in zones with all other platforms.
- Sun and Windows are incompatible in the same zone.

Zone 1	Zone 2	Zone 3	Zone 4	Zone 5
NetWare OpenVMS Tru64 UNIX Microsoft Windows	OpenVMS Microsoft Windows	Linux	HP-UX	IBM AIX

Compatible controller SCSI-modes and controller failover modes

Table 54 summarizes information about compatible SCSI modes, and Table 55 summarizes information about supported storage system failover modes for all platforms for EMA/ESA12000, EMA16000, MA/RA8000, MA6000 storage systems using ACS 8.7/8.8.

Table 54: Compatible SCSI modes for EMA/ESA12000, EMA16000, MA/RA8000, MA6000 storage systems using ACS 8.7, 8.8

SCSI-2 CCL	SCSI-2 no CCL	SCSI-3
HP-UX	HP-UX	HP-UX
		OpenVMS
Tru64 UNIX 4.0F, 4.0G	Tru64 UNIX 4.0F, 4.0G	
Tru64 UNIX 5.1, 5.1A, 5.1B	Tru64 UNIX 5.1, 5.1A, 5.1B	Tru64 UNIX 5.1, 5.1A, 5.1B
IBM AIX	IBM AIX	IBM AIX
		Linux
	Microsoft Windows	Microsoft Windows
Novell NetWare	Novell NetWare	Novell NetWare
Sun Solaris	Sun Solaris	Sun Solaris

Table 55: Compatible failover modes for EMA/ESA12000, EMA16000, MA/RA8000, MA6000 storage systems using ACS 8.7, 8.8

Transparent	Multiple-bus
HP-UX	HP-UX
	OpenVMS
Tru64 UNIX 4.0F, 4.0G, 5.1, 5.1A, 5.1B	Tru64 UNIX 5.1, 5.1A, 5.1B
IBM AIX	IBM AIX
Linux	
SuSE SLES 7 (ProLiant x86)	
Microsoft Windows	Microsoft Windows
Novell NetWare	Novell NetWare
Sun Solaris	Sun Solaris

Combined shared access interoperability table

Table 56 combines the information from the previous tables into a single table. The table can be used to determine controller settings for a single EMA/ESA12000, EMA16000, MA/RA8000, MA6000 Storage System using ACS 8.7 or 8.8, that is shared between two or more platforms and operating systems.

Table 56: Platform interoperability for single shared EMA/ESA12000, EMA16000, MA/RA8000, MA6000 storage systems – ACS 8.7, 8.8

Platform or operating system	HP-UX MC/ ServiceGuard Clusters	OpenVMS Clusters	Tru64 UNIX 4.0F, 4.0G Trucluster Software Products V1.6	Tru64 UNIX 5.1, 5.1A, 5.1B TruCluster Server Version 5.1, 5.1A, 5.1B	IBM AIX	Linux	Microsoft Windows MSCS	Novell NetWare 5.1, 6, 6.5 Clusters 1.01, 1.06, 1.07	Sun Solaris Sun Clusters VERITAS Clusters
HP-UX MC/ Service-Guard Clusters	Fabric attachment Transparent or Multiple-Bus FABRIC SCSI-2 or SCSI-3 FC-AL attachment, Transparent or Multiple-Bus LOOP_HARD SCSI-2	Fabric attachment With Zoning Multiple-Bus FABRIC SCSI-3	Fabric attachment With Zoning Transparent FABRIC SCSI-2	Fabric attachment With Zoning Transparent or Multiple-Bus FABRIC SCSI-2 or SCSI-3	Fabric attachment With Zoning Transparent or Multiple-Bus FABRIC SCSI-2 or SCSI-3	Fabric attachment With Zoning Transparent or Multiple-Bus FABRIC SCSI-3	Fabric attachment With Zoning Transparent or Multiple-Bus FABRIC SCSI-2 No CCL or SCSI-3	Fabric attachment With Zoning Transparent or Multiple-Bus (5.1, 6) FABRIC SCSI-2 or SCSI-3	Fabric attachment With Zoning Transparent or Multiple-Bus FABRIC SCSI-2 or SCSI-3
OpenVMS Clusters	Fabric attachment With Zoning Multiple-Bus FABRIC SCSI-3	Multiple-Bus FABRIC SCSI-3	Requires two storage systems	Multiple-Bus FABRIC SCSI-3	With Zoning Multiple-Bus FABRIC SCSI-3	With Zoning Multiple-Bus FABRIC SCSI-3	Multiple-Bus FABRIC SCSI-3	5.1, 6: Multiple-Bus FABRIC SCSI-3 4.2: Requires two storage systems	Multiple-Bus FABRIC SCSI-3
Tru64 UNIX 4.0F, 4.0G Trucluster Software Products V1.6	Fabric attachment With Zoning Transparent FABRIC SCSI-2	Requires two storage systems	Transparent FABRIC SCSI-2	Transparent FABRIC SCSI-2	With Zoning Transparent FABRIC SCSI-2	Requires two storage systems	Transparent FABRIC SCSI-2 No CCL	Transparent FABRIC SCSI-2	Transparent FABRIC SCSI-2
Tru64 UNIX 5.1, 5.1A, 5.1B TruCluster Server Version 5.1, 5.1A, 5.1B	Fabric attachment With Zoning Transparent or Multiple-Bus FABRIC SCSI-2 or SCSI-3	Multiple-Bus FABRIC SCSI-3	Transparent FABRIC SCSI-2	Transparent or Multiple-Bus FABRIC SCSI-2 or SCSI-3	With Zoning Transparent or Multiple-Bus FABRIC SCSI-2 or SCSI-3	With Zoning Transparent or Multiple-Bus FABRIC SCSI-3	Transparent or Multiple-Bus FABRIC SCSI-2 No CCL or SCSI-3	Transparent or Multiple-Bus (5.1, 6) FABRIC SCSI-2 or SCSI-3	Transparent or Multiple-Bus FABRIC SCSI-2 or SCSI-3

Table 56: Platform interoperability for single shared EMA/ESA12000, EMA16000, MA/RA8000, MA6000 storage systems – ACS 8.7, 8.8

Platform or operating system	HP-UX MC/ ServiceGuard Clusters	OpenVMS Clusters	Tru64 UNIX 4.0F, 4.0G Trucluster Software Products V1.6	Tru64 UNIX 5.1, 5.1A, 5.1B TruCluster Server Version 5.1, 5.1A, 5.1B	IBM AIX	Linux	Microsoft Windows MSCS	Novell NetWare 5.1, 6, 6.5 Clusters 1.01, 1.06, 1.07	Sun Solaris Sun Clusters VERITAS Clusters
IBM AIX	Fabric attachment With Zoning Transparent or Multiple-Bus FABRIC SCSI-2 or SCSI-3	With Zoning Multiple-Bus FABRIC SCSI-3	With Zoning Transparent FABRIC SCSI-2	With Zoning Transparent or Multiple-Bus FABRIC SCSI-2 or SCSI-3	Transparent or Multiple-Bus FABRIC SCSI-2 or SCSI-3	With Zoning Transparent or Multiple-Bus FABRIC SCSI-3	With Zoning Transparent or Multiple-Bus FABRIC SCSI-2 No CCL or SCSI-3	With Zoning Transparent or Multiple Bus (5.1, 6) FABRIC SCSI-2 or SCSI-3	With Zoning Transparent or Multiple-Bus FABRIC SCSI-2 or SCSI-3
Linux	Fabric attachment With Zoning Transparent or Multiple-Bus FABRIC SCSI-3	With Zoning Multiple-Bus FABRIC SCSI-3	Requires two storage systems	With Zoning Transparent or Multiple-Bus FABRIC SCSI-3	With Zoning Transparent or Multiple-Bus FABRIC SCSI-3	Transparent or Multiple-Bus FABRIC SCSI-3	With Zoning Transparent or Multiple-Bus FABRIC SCSI-3	With Zoning Transparent or Multiple-Bus FABRIC SCSI-3	With Zoning Transparent or Multiple-Bus FABRIC SCSI-3
Microsoft Windows MSCS	Fabric attachment With Zoning Transparent or Multiple-Bus FABRIC SCSI-2 No CCL or SCSI-3	Multiple-Bus FABRIC SCSI-3	Transparent FABRIC SCSI-2 No CCL	Transparent or Multiple-Bus FABRIC SCSI-2 No CCL or SCSI-3	With Zoning Transparent or Multiple-Bus FABRIC SCSI-2 No CCL or SCSI-3	With Zoning Transparent or Multiple-Bus FABRIC SCSI-3	Transparent or Multiple-Bus FABRIC SCSI-2 No CCL or SCSI-3	Transparent or Multiple-Bus (5.1, 6) FABRIC SCSI-2 No CCL or SCSI-3	With Zoning Transparent or Multiple-Bus FABRIC SCSI-2 No CCL or SCSI-3

Table 56: Platform interoperability for single shared EMA/ESA12000, EMA16000, MA/RA8000, MA6000 storage systems – ACS 8.7, 8.8

Platform or operating system	HP-UX MC/ ServiceGuard Clusters	OpenVMS Clusters	Tru64 UNIX 4.0F, 4.0G Trucluster Software Products V1.6	Tru64 UNIX 5.1, 5.1A, 5.1B TruCluster Server Version 5.1, 5.1A, 5.1B	IBM AIX	Linux	Microsoft Windows MSCS	Novell NetWare 5.1, 6, 6.5 Clusters 1.01, 1.06, 1.07	Sun Solaris Sun Clusters VERITAS Clusters
Novell NetWare 5.1, 6, 6.5 Clusters 1.01, 1.06	Fabric attachment With Zoning Transparent or Multiple-Bus (5.1, 6) FABRIC SCSI-2 or SCSI-3	5.1, 6: Multiple-Bus FABRIC SCSI-3 4.2: Requires two storage systems	Transparent FABRIC SCSI-2	Transparent or Multiple-Bus (5.1, 6) FABRIC SCSI-2 or SCSI-3	With Zoning Transparent or Multiple Bus (5.1, 6) FABRIC SCSI-2 or SCSI-3	With Zoning Transparent or Multiple-Bus FABRIC SCSI-3	Transparent or Multiple-Bus (5.1, 6) FABRIC SCSI-2 No CCL or SCSI-3	Transparent or Multiple-Bus (5.1, 6) FABRIC SCSI-2 or SCSI-3	With Zoning Transparent or Multiple-Bus (5.1, 6) FABRIC SCSI-2 or SCSI-3
Sun Solaris Sun Clusters VERITAS Clusters	Fabric attachment With Zoning Transparent or Multiple-Bus FABRIC SCSI-2 or SCSI-3	Multiple-Bus FABRIC SCSI-3	Transparent FABRIC SCSI-2	Transparent or Multiple-Bus FABRIC SCSI-2 or SCSI-3	With Zoning Transparent or Multiple-Bus FABRIC SCSI-2 or SCSI-3	With Zoning Transparent or Multiple-Bus FABRIC SCSI-3	With Zoning Transparent or Multiple-Bus FABRIC SCSI-2 No CCL or SCSI-3	With Zoning Transparent or Multiple-Bus (5.1, 6) FABRIC SCSI-2 or SCSI-3	Transparent or Multiple-Bus FABRIC SCSI-2 or SCSI-3

Booting from the SAN

Table 57 indicates the platforms and operating systems that are able to boot from SAN storage.

Note: SAN boot through the B-Series MP Router is not supported.

Table 57: EVA3000/4000/5000/6000/8000, EMA/ESA12000, EMA16000, MA/RA8000, and MA6000 SAN Boot by Operating System

Platform/operating system	Comments
HP-UX 11iv1, 11iv2	EVA5000/3000. Requires Secure Path v3.0d or later. Contact an HP storage representative for specific HBA requirements for boot support.
Microsoft Windows 2000 Server, Advanced Server, Windows NT 4.0 MSCS MS Windows Server 2003 Standard Edition and Enterprise Edition (32-bit)	EVA5000/EVA3000 EMA/ESA12000,EMA16000, MA/RA8000, MA6000, MSA1000 Contact an HP storage representative for specific HBA requirements for boot support. See *Booting 32-bit Windows from a storage area network* for more information
OpenVMS Clusters	EVA5000/EVA3000, EMA/ESA12000, EMA16000, MA/RA8000, MA6000, MSA1000, Contact an HP storage representative for specific HBA requirements for boot support. Requires use of *wwidmgr,* SRM console firmware v6.5 (minimum)
Tru64 UNIX TruCluster Software Products	
Linux (32-bit/64-bit)	See *Booting 32-Bit Linux systems from a storage area network and Booting 64-bit Linux systems from a storage area network application notes* for more information.

Continuous Access EVA or DRM supports the replication of the boot disk providing the page and swap files are not on that disk. For optimal performance, HP recommends that you place both disks on internal server storage. Support is limited to the operating systems listed in Table 57.

SAN storage system rules

9

This chapter describes rules related to specific storage systems. It describes the following topics:

For additional information, see the individual product-specific documentation or contact an HP storage representative. See “About this guide” on page 23 for a list of related documentation.

HP XP and VA configuration rules

1. XP12000/1024/128/512/48/256 storage systems are supported in all fabric topology configurations described in this guide. The SANs using the Fibre Channel switches are listed in Table 10 on page 83 (B-Series), Table 20 on page 99 (C-Series), Table 28 on page 107 (M-Series), and the B-Series MP Router, except when otherwise listed.
2. VA storage systems are supported in all fabric topology configurations described in this guide. The SANs using the Fibre Channel switches are listed in Table 10 on page 83 (B-Series), Table 20 on page 99 (C-Series), Table 28 on page 107 (M-Series), and the B-Series MP Router.
3. VA storage systems that are shared between MSCS (Windows) and MC/Service Guard (HP-UX) clusters are supported. Requires proper assignment and securing of LUNs to the individual clusters.
4. Zoning may be required when configuring these storage system types in the same physical SAN, or for access from the same server with other storage system types. See "Common server access" on page 135, for supported common server access configurations.

EVA3000/4000/5000/6000/8000 configuration rules

1. The EVA3000/4000/5000/6000/8000 Storage Systems are supported in all Fabric topology configurations using XCS v5 or VCS v3, as described in this guide. Contact an HP storage representative for specific XCS or VCS version support details. The EVA3000/4000/5000/6000/8000 Storage Systems are compatible in SANs using the Fibre Channel switches listed in Table 10 on page 83, Table 20 on page 99, and Table 28 on page 107, and the B-Series MP Router.
2. For SANs with more than 1024 HBAs, an HSV controller must be zoned to see a maximum of 1024 HBAs. It may be necessary to add a zone to a SAN to satisfy the 1024 HBA limit.
3. The supported platforms and operating systems are listed in "Specific platform/operating system rules – EVA3000/5000 (VCS v3), EVA4000/6000/8000 (XCS v5), EMA/ESA12000, EMA16000, MA/RA8000, MA6000 (ACS 8.7, 8.8) storage systems, B-Series and M-Series switches" on page 147.
4. Shared access and heterogeneous platform zoning requirements are listed in Table 56.
5. Supports Multiple-Bus Failover mode only. Generally, Multiple-Bus Failover requires a minimum of 2 Fibre Channel HBAs and native operating system or layered multi-path driver functionality. See the white paper listed below for exceptions.
6. EVA3000/5000 supports connection of single HBA servers. See the white paper *Connecting Single HBA Servers to the Enterprise Virtual Array without Multipathing Software* at:

   ftp://ftp.compaq.com/pub/products/storageworks/whitepapers and also

   http://storage.inet.Compaqcorp.net/Document_Storage/whitepapers/new_library/SingleHBA_for_EVA-HP3_121002.pdf

Note: Servers without multi-pathing software are *not* supported by Continuous Access EVA.

7. Overlapping zones are supported with disk and tape.
8. Overlapping storage port zones are supported when multiple operating systems require sharing an array port.
9. See "Storage Management Appliance rules and recommendations" on page 176 for information about configuring the Storage Management Appliance to manage EVA3000/4000/5000/6000/8000s and EMA/ESA12000, EMA16000, MA/RA8000, MA6000 storage systems in the same SAN fabric.
10. Zoning may be required when configuring these storage system types in the same physical SAN, or for access from the same server with other storage system types. See "Common server access" on page 135 for specific configurations where common server access is supported without the need for zoning or separate HBAs.
11. SSP/LUN level masking – Use storage system LUN presentation to enable/disable LUN access to specific hosts.
12. All host table entries must have the proper operating system type parameter set based on the platform type accessing the assigned LUNs.

Note: Shared access between different servers to the same storage unit (LUN) requires specific application software (for example, cluster software) to ensure proper data preservation.

EVA3000/4000/5000/6000/8000 maximums

Table 58 on page 174 lists the maximum connections supported by EVA storage systems as well as the maximum supported storage limits for each hardware platform or operating system. The maximums shown are for access to a single EVA with dual redundant controllers. If the connection requirements for the number of servers in a particular SAN exceed the maximum, deploy multiple storage systems within the SAN.

Note: This section specifies general EVA limits. Specific solution subset configurations such as high-availability clusters or applications such as Continuous Access (see "SAN/Continuous Access EVA integration" on page 193) may impose lower-level limits on connectivity for the solution. In these instances, the solution limits must be adhered to as specified by the solution configuration documentation.

- Maximum of 1024 Host Bus Adapters (HBA).
- Maximum of 512 LUNs per EVA3000/5000.
- Maximum of 1024 LUNs per EVA4000/6000/8000.
- Maximum of 256 Hosts: A Host is defined to contain one or more HBAs.
- The total number of LUN Presentations for all LUNs must not exceed 8192.
  A LUN Presentation is defined as the number of Hosts a LUN is presented to, irrespective of how many adapters might be in any given Host.
  For example, if a LUN is presented to 8 Hosts, that LUN has 8 LUN Presentations.
  If a LUN is presented to 2 Hosts, that LUN has 2 LUN Presentations).

  Example:

LUNs #001 thru #032 are presented to a 8 Node Cluster	= 0256 LUN Presentations
LUNs #033 thru #064 are presented to a 8 Node Cluster	= 0256 LUN Presentations
LUNs #065 thru #096 are presented to a 8 Node Cluster	= 0256 LUN Presentations
LUNs #097 thru #128 are presented to a 8 Node Cluster	= 0256 LUN Presentations
LUNs #129 thru #160 are presented to a 8 Node Cluster	= 0256 LUN Presentations
LUNs #161 thru #192 are presented to a 4 Node Cluster	= 0128 LUN Presentations
LUNs #193 thru #200 are presented to a single host	= 0008 LUN Presentations
Total LUN Presentations	= 1416 LUN Presentations

- When all LUNS are presented to all Hosts, this simple rule applies:
  The number of LUNs times the number of Hosts must not exceed 8192.

Table 58: SAN/platform storage maximums - EVA3000/4000/5000/6000/8000

Platform or operating system	Host bus adapters per server	Active controller ports (targets) per HBA	LUNs per HBA target
See **Reference Notes**	1	2	3, 4, 5
HP-UX	16	4 (2 storage systems)	128
OpenVMS	26	128 (32 storage systems)	511 (9999)
Tru64 UNIX	64	254 (64 storage systems)	255
IBM AIX	16	4 (2 storage systems)	32
Linux	4	4 (2 storage systems)	32 (128)
Microsoft Windows 2000, 2003	8	4 (2 storage systems)	255
Microsoft Windows NT	8	4 (2 storage systems)	8/64
Novell NetWare	4	4 (2 storage systems)	128/16
Sun Solaris	16	4 (2 storage systems)	128

Reference notes

1. The maximum number of HBAs supported per server is dependent on the specific server model.
2. Column 3 represents the total number of active storage controller ports supported per HBA. For Tru64 UNIX and OpenVMS this column typically represents the total number of active controller ports per HBA when accessing all ports of a storage system or all ports on multiple storage systems. For all other platforms, this column typically represents 2 ports per storage system, or a total of 4 ports across 2 storage systems, see "High-availability configuration considerations" on page 199. Zoning may be required to limit the number of active targets presented to each HBA to the stated maximum.
3. Microsoft Windows NT supports 8 LUNs per HBA target with Large LUN feature disabled and 64 LUNs per HBA target with Large LUN feature enabled. Windows 2000 and Windows 2003 support Large LUN by default with up to 255 LUNs per HBA target.
4. Sun configurations configured with the same HBA accessing both an EVA and an EMA/ESA12000, EMA16000, MA/RA8000, or MA6000, support a maximum of 64 LUNs per HBA target for the EMA/ESA12000, EMA16000, MA/RA8000, or MA6000 storage system.
5. Novell NetWare single path configurations support a maximum of 128 LUNs per host. For multipathing configurations using Secure Path, a maximum of 16 LUNs per host are supported.

EVA3000/4000/5000/6000/8000 Microsoft Windows cluster maximums

The maximum number of nodes which are part of a Microsoft cluster attached to one Enterprise Virtual Array may not exceed a total amount of 32 nodes. For example, the following configurations are all valid:

- Sixteen 2-node Windows 2000
- Four 8-node Windows Server 2003
- Two 2-node Windows 2000, three 4-node Windows Server 2003, and two 8-node Windows Server 2003

While using Continuous Access, an Enterprise Virtual Array may not exceed these figures, even after a failover has been performed. The limits are based on numbers of nodes/hosts regardless of actively presented LUNs.

Additional standalone servers of any supported type, or non-Windows clustered servers, up to the limit per EVA specified for the particular clustered operating system, may be added up to the total EVA limit of 256 servers.

Storage Management Appliance rules and recommendations

Whenever a Storage Management Appliance (SMA) is placed in a fabric it is recommended that a dedicated storage management zone be created. This zone is specifically for the SMA and the elements it is to monitor and manage.

The SMA communicates with the EVA (HSV controller) and EMA/ESA12000, EMA16000, MA/RA8000, MA6000 storage systems (HSG controller) in-band, that is, within the Fibre Channel fabric itself. It is not necessary or recommended to include either the switch WWNs or server HBA WWNs in this zone. Management communication to these devices from the SMA is done out-of-band or outside the fabric via TCP/IP.

For example, create a zone called SANAPP_1_ZONE that would contain the SMA host bus adapter WWN and the WWNs of all the HSG or HSV controllers managed by this SMA. Because fabric devices can be in multiple zones, this will have no effect on other zones containing the same HSG and HSV controller WWNs.

1. Command View EVA v4.0 supports EVA4000/6000/8000 using XCS v5 or later and EVA3000/5000 using VCS v3.014 or later.
2. Any EVA storage system can only have one active SMA, general purpose server (GPS), or management station (dedicated server) managing it. Any standby SMA, GPS, or management station can be powered on, but the Command View EVA or Continuous Access user interface must not control the storage system. For further information, see the HP StorageWorks Continuous Access EVA Operations Guide available on the HP web.
3. An SMA, GPS, or management station is required to manage EVA3000/4000/5000/6000/8000s or when managing EMA/ESA12000, EMA16000, MA/RA8000, MA6000 storage systems when using the HSG element manager.
4. Multiple SMAs, GPSs, or management stations per fabric are allowed as long as only one is active at a time. In a Continuous Access EVA configuration, both source (initiator) and destination (target) storage systems should be managed by the same SMA, GPS, or management station.
5. The HSV and HSG Element Managers can operate in a dual fabric configuration.
6. For SANs with more than 1024 HBAs, an HSV controller must be zoned so that it can see no more than 1024 HBAs. It may be necessary to add a zone to a SAN to satisfy the 1024 HBA limit.

See “Heterogeneous server rules” on page 127, for rules about mixing specific platforms in a Heterogeneous SAN without the need for fabric zoning.

EMA/ESA12000, EMA16000, MA/RA8000, MA6000 configuration rules

1. These storage systems are supported in all SAN Fabric topology configurations described in this guide. These storage systems are supported in SANs using the Fibre Channel switches listed in Table 10 on page 83, Table 20 on page 99, Table 28 on page 107, and SANs using the B-Series MP Router.
2. Limit the number of connections visible to each storage system to a maximum of 96 by using fabric zoning. (This is the maximum supported limit for ACS 8.7 and 8.8).
3. See "Specific platform/operating system rules – EVA3000/5000 (VCS v3), EVA4000/6000/8000 (XCS v5), EMA/ESA12000, EMA16000, MA/RA8000, MA6000 (ACS 8.7, 8.8) storage systems, B-Series and M-Series switches" on page 147 for supported platforms and operating systems
4. Shared access and heterogeneous platform zoning requirements are listed in Table 56 on page 166. The heterogeneous platform and operating system mix in the SAN determines the appropriate controller topology attachment, SCSI mode, and Command Console LUN settings for shared storage systems.
5. Overlapping zones are supported with disk and tape.
6. Overlapping storage port zones are supported if more than one operating system needs to share an array port.
7. Single or dual redundant controller configurations are supported. For dual redundant controllers, the available failover modes are Transparent and Multiple-Bus. Multiple-Bus failover requires native operating system or layered multi-path driver functionality.

Note: Windows 2003 is not supported with Transparent failover mode.

8. All host connection table entries must have the proper operating system type parameter set based on the platform type accessing the assigned LUNs.
9. See "Storage Management Appliance rules and recommendations" on page 176 for information about configuring the SMA to manage EVA5000/EVA3000s and EMA/ESA12000, EMA16000, MA/RA8000, MA6000 storage systems in the same SAN fabric.
10. Zoning may be required when configuring these storage system types in the same physical SAN, or for access from the same server with other storage system types. See "Common server access" on page 135 for specific configurations where common server access is supported without the need for zoning or separate HBAs.
11. SSP/LUN level masking – Use storage system Selective Storage Presentation to enable/disable LUN access to specific connections. Use the unit offset feature to provide needed LUN numbering for host connections. The default LUN numbering for Transparent Failover mode is 0 to 99 for controller port 1, and 100 to 199 for controller port 2. For Multiple-bus Failover mode, the default LUN numbering is 0 to 199 on all controller ports.

Note: Shared access between different servers to the same storage unit (LUN) requires specific application software (cluster software) to ensure proper data preservation.

12. F-Port fabric attachment to the SAN is available through all Fibre Channel switches listed in Table 10 on page 83, Table 20 on page 99, and Table 28 on page 107. Controller setting is FABRIC topology.
13. FL-Port fabric loop attachment to the SAN with QuickLoop is available through certain Fibre Channel SAN switch models. See Table 10 on page 83, and the specific Fibre Channel SAN switch model documentation for more information. Controller port topology set to LOOP_HARD.
14. All controller ports must be set to the same topology type.

Maximum paths or maximum LUNs

For EMA/ESA12000, EMA16000, MA/RA8000, or MA6000 storage systems, use the HSG60/80 controller Unit Offset feature to maximize path accessibility or to maximize the number of LUNs.

- **For Maximum Controller Path Accessibility to the same set of LUNs**

  Use a common unit offset value for all 4 controller ports. Access to a common set of LUNs through all 4 controller host ports is provided by using the same unit offset value on all controller host port connections for each server. For example, set the unit offset value for connections on all 4 controller ports to zero (0) for a given server. The server will be capable of accessing one set of LUNs beginning with LUN 0 (LUN 0 is the Command Console LUN if set to SCSI-3 mode) from all 4 controller host ports. This method provides for the highest number of paths to a given set of LUNs.

- **For Maximum LUN Count**

  Use distinct controller port unit offsets for each port pair. Access one set of LUNs with controller port 1 of each controller and access a different set of LUNs with controller port 2 of each controller. For example, set the unit offset value for connections on controller port 1 of each controller to zero, and then set a unit offset value for connections on controller port 2 of each controller to 100 for a given server. The server will be capable of accessing one set of LUNs beginning with LUN 0 (LUN 0 is the Command Console LUN if set to SCSI-3 mode) through controller port 1 on each controller, and a second set of LUNs beginning with LUN 100 through controller port 2 of each controller. This method provides the highest number of LUNs accessed through a reduced number of paths. It also allows for the highest number of servers. See Figure 49 on page 182.

EMA/ESA12000, EMA16000, MA/RA8000, and MA6000 maximums

Table 59 lists the maximum supported storage limits for each hardware platform or operating system with access to MA6000 storage systems with dual redundant HSG60 controllers, and EMA/ESA12000, EMA16000, or MA/RA8000 storage systems with dual redundant HSG80 controllers. If the maximums listed are below the requirements for the number for servers required, you can deploy multiple storage systems within the SAN.

Table 59: Platform maximums - MA6000, MA/RA8000, EMA/ESA12000, EMA16000 storage systems using ACS 8.7, 8.8

Platform or operating system	Host bus adapters per server	Active controller ports (targets) per HBA	LUNs per HBA target	Port maximums HBAs per active controller port		Storage system maximums HBAs per storage system	
Controller failover mode				TF	MB	TF	MB
HP-UX 11.0, 11i v1	16	4 (2)	8/128	8	8	16	32
OpenVMS 7.2-2, 7.3, 7.3-1, 7.3-2	26	128 (32)	128 (10000)	N/A	24	N/A	48
Tru64 UNIX 4.0F, 4.0G	32	4 (4)	8	4	N/A	8	N/A
Tru64 UNIX 5.1, 5.1A, 5.1B	64	254 (63)	128 (255)	48	24	96	48
IBM AIX 4.3.3, 5.1	8	8 (4)	32	12	12	24	24
Red Hat Linux 7.2, (Proliant x86) Red Hat Linux 7.1, 7.2 (Alpha)	2	4 (2)	64	4	N/A	8	N/A
Advanced Server 2.1 (BL20P, BL40P, Proliant x86)	2	4 (2)	64	4	4	8	16
SuSE Linux 7.2 (ProLiant x86)	2	4 (2)	64	4	N/A	8	N/A
SuSE SLES 7(ProLiant x86)	2	4 (2)	64	4	4	8	16
Microsoft Windows 2000 Server, Advanced Server SP2, SP3 Windows 2000 Datacenter	8	4 (2)	255	8	16 (See Figure 51)	16	32 (See Figure 51)
MS Windows Server 2003 Standard Edition and Enterprise Edition (32-bit)	8	6 (3)	255	N/A	16 (See Figure 51)	N/A	32 (See Figure 51)
Microsoft Windows NT 4.0 SP6a	8	4 (2)	8/64	8	8 (See Figure 50)	16	32 (See Figure 50)
Novell NetWare 5.1, 6, 6.5	4	4 (2)	32	8	8	16	32
SUN Solaris 2.6, 7 & 8 (32/64 bit)	16	4 (2)	64	8	8	16	32
Reference Notes	1	2	3, 4, 5, 6	7, 8, 9, 10		7, 9, 10	

Reference notes

1. The maximum number of HBAs and LUNs supported per server depends on the specific server model.
2. Column 3 represents the total number of active storage controller ports supported per HBA. For Tru64 UNIX and OpenVMS this column typically represents the total number of active controller ports per HBA when accessing all ports of a storage system or ports on multiple storage systems. For most other platforms this column typically represents 2 ports per storage system, or a total of 4 ports across 2 storage systems. Windows 2003 is supported for access to 6 ports across 3 storage systems. IBM AIX is supported for access

to 8 ports across 4 storage systems.Use of zoning may be required to limit the number of active targets (controller ports) presented to each HBA to the maximums stated for each platform in this column. OpenVMS 7.2-2 or later is required.

3. Values in this column are reduced by one if the command console LUN is enabled.
4. Microsoft Windows NT supports 8 LUNs per HBA target with Large LUN feature disabled, and 64 LUNs per HBA target with Large LUN feature enabled. Windows 2000 supports Large LUN by default (LUNs 0–199), Secure Path for Windows supports usage of LUNs 0–63.
5. The Tru64 5.1, 5.1A, 5.1B, operating system maximum is 255 LUNs per target. For OpenVMS, the operating system maximum is 10,000 LUNs. The single storage system maximum is 128 LUNs.
6. For HP-UX, 128 LUNs per HBA target when the connection operating system type is set to HP_VSA.
7. A Storage Management Appliance or a general purpose server with management software requires 2 connection table entries per fabric. However, these connection table entries do not affect the total number of servers or HBAs supported if there are available entries in the table. For example, Windows supports a maximum of 16 servers regardless of using an SMA. The SMA executes management commands through its Fibre Channel connection, therefore it is not counted when determining the total number of servers allowed on a single storage system from an I/O load perspective.
8. The maximum number of HBAs that can be configured for access to an active controller port. Assumes 1 HBA per server for single path using controller transparent failover, or 2 HBAs per server for multipathing using controller multiple-bus failover. For transparent failover, the limit is specified by controller port pair–1 active and 1 standby controller port. For multiple-bus failover, the limit is specified per single active port.
9. The maximum specified for each platform are the result of one or more of following conditions:
   - A qualification limit.
   - Command flow queuing characteristics of specific HBA drivers.
   - Connection table size in the array controller software in conjunction with the number of HBA to controller port paths.

   For maximum server or HBA connectivity using controller transparent failover, limit the number of active HBA-to-controller-port paths to one per server (Figure 49.) The use of zoning is required to limit the number HBAs visible to each active controller port.

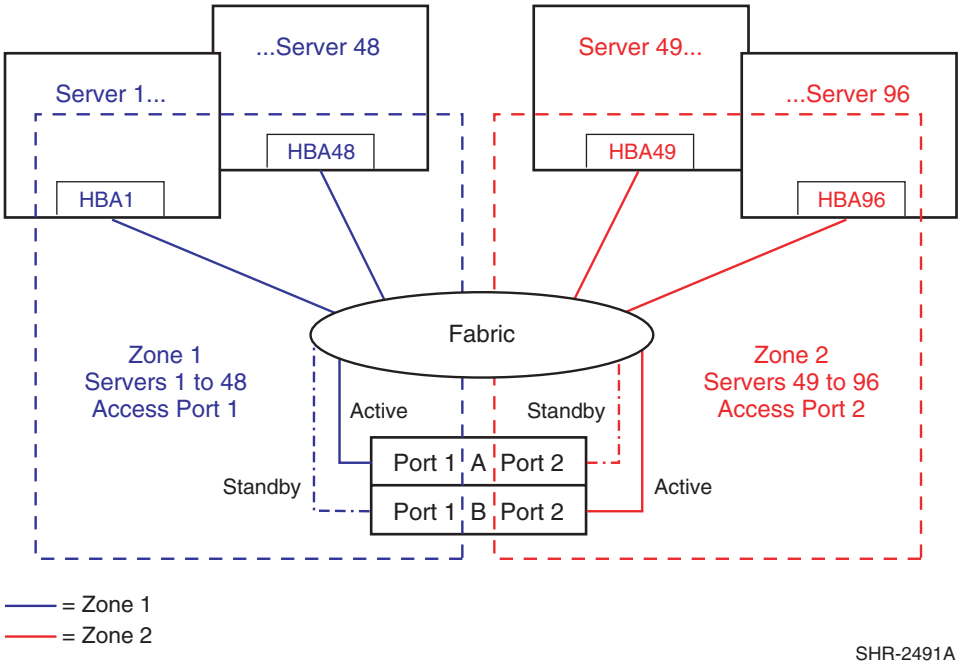

Figure 49: Maximum server example for Tru64 UNIX 5.x with transparent failover using 96 connections and one path per server

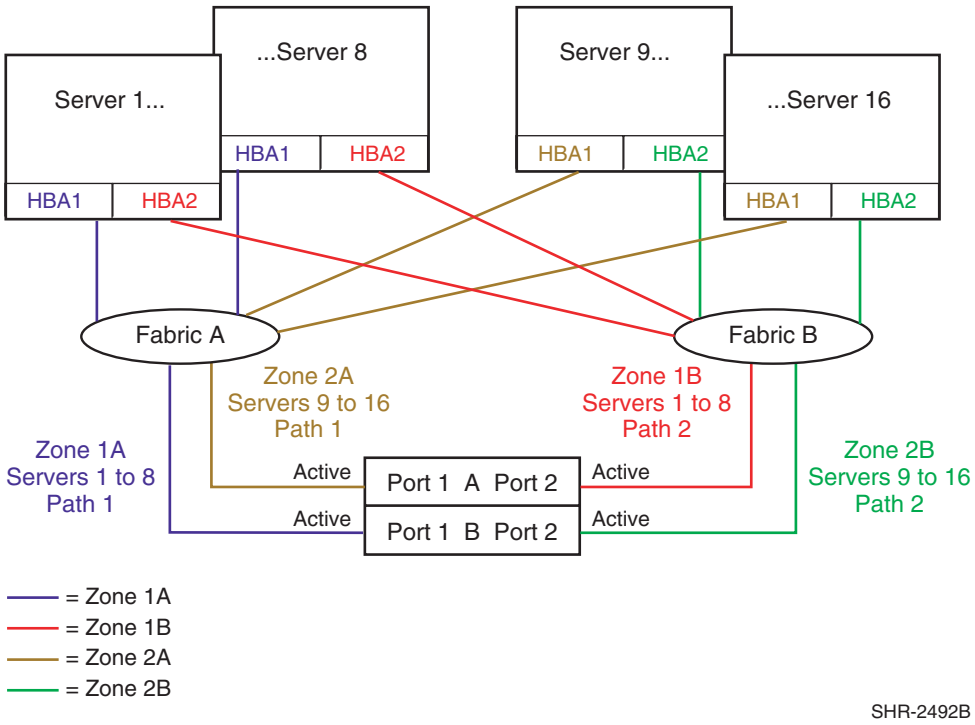

Figure 50: Maximum server example for Windows NT using 16 servers with multiple-bus failover and two paths per server

Note: Zones for the figure above are visible on-screen or on a color printout.

For maximum server or HBA connectivity on most operating systems, when using controller multiple-bus failover, limit the number of active HBA-to-controller-port paths to two per server (Figure 50) The use of zoning is required to limit the number of HBAs visible to each active controller port.

For OpenVMS, Tru64 UNIX, and Windows 2000, the maximum server or HBA connectivity is available with up to four paths per server in multiple-bus failover mode. The maximum Windows 2000 configuration is shown in Figure 51.

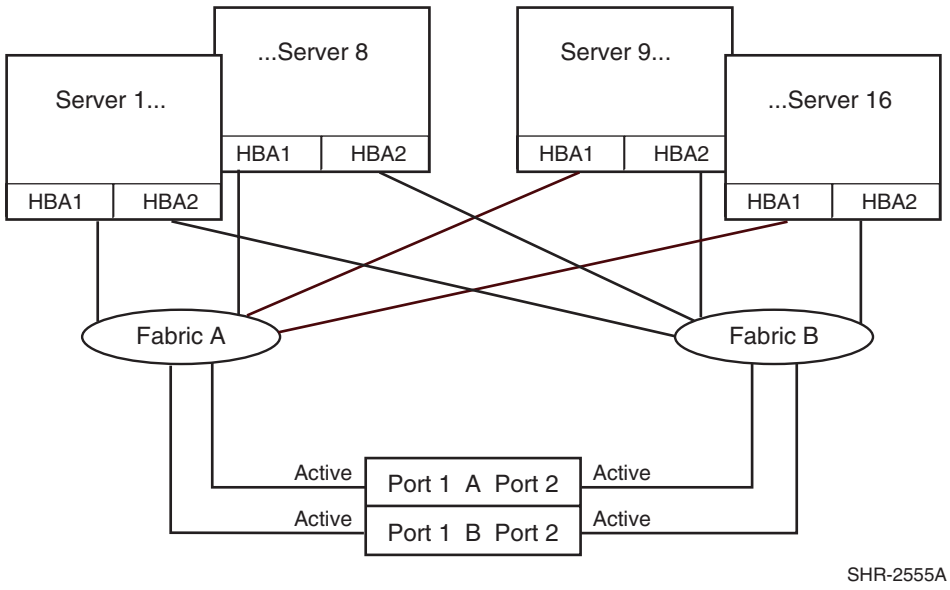

Figure 51: Maximum server example for Windows 2000 using 16 servers with multiple-bus failover and four paths per server

10. In a heterogeneous SAN environment where different platform or operating system types are sharing a single storage system, the maximum number of servers or HBAs supported is equal to the lowest maximum listed in these columns for the operating systems that are sharing the storage system. All platforms or operating systems listed are supported for shared access to the same storage system provided the rules listed in "Heterogeneous SAN platform interoperability for EVA3000/4000/5000/6000/8000 and EMA/ESA12000, EMA16000, MA/RA8000, MA6000 storage systems" on page 190 are followed. See Table 56 on page 166 for specific information about sharing a single storage system across multiple platform or operating system types.

Note: See "SAN/DRM integration" on page 195 for the maximum number of servers supported for storage systems configured for DRM.

Note: See "SAN/Continuous Access EVA integration" on page 193 for the maximum number of servers supported for storage systems configured for Continuous Access EVA.

Specific platform/operating system rules – MSA1500, MSA1000, RA4100, RA4000

This section defines the rules and guidelines related to specific platforms/operating systems for MSA1500, MSA1000, RA4100, and RA4000 storage systems. For operating system storage attachment, HBA attachment, current HBA/driver/FW revision support, and specific MSA and RA4x FW version support, contact an HP storage representative. Information about the latest MSA1500 and MSA1000 support is available at:

http://www.hp.com/go/msa1500cs

http://www.hp.com/go/msa1000

MSA1500 FW 4.82, B-Series and MP Router, C-Series, M-Series switches

The MSA1500 is supported on the B-Series Switches and MP Router, M-Series Switches, and C-Series Switches.

The MSA1500 is supported with these platforms:

- Red Hat Linux EL 2.1, EL 3 32-Bit/64-Bit
- SuSE Linux SLE8 32-Bit/64-Bit
- Windows 2003 32-Bit/IA64
- Windows 2000
- NetWare 5.1, 6.0, 6.5
- VMware ESX 2.5

For specific support information, see the *HP StorageWorks Modular Smart Array 1500 cs Compatibility Matrix* listed under the Technical Documentation link at:

http://www.hp.com/go/msa1500cs

Heterogeneous SAN platform interoperability for MSA1500 storage

This section specifies the rules for shared heterogeneous access to a single MSA1500 storage system. MSA1500 storage systems are supported for shared access with any combination of the operating systems listed.

Single MSA1500 shared heterogeneous access, any combination of:

- Windows 2003 32-Bit/64-Bit
- Windows 2000 SP3/SP4

Additionally, standalone servers and clustered servers are supported on the same MSA1500.

MSA1000 FW 4.38, B-Series switches and MP Router, C-Series and M-Series switches

- HP-UX 11i v1, 11i v2

MSA1000 FW 4.32, 2.38, B-Series and MP Router, C-Series, and M-Series switches

VMware ESX 2.1 and 2.5 are supported with the MSA1000 with firmware versions 4.32 and 2.38.

Linux Red Hat EL 3, Red Hat AS 2.1 (32-bit) (64-bit single-path only), SLES8 SP2a (32-bit)(64-bit single-path only), SLES 8/United Linux 1.0 32-bit and 64-bit

- Supports LifeKeeper Clusters (does not support 64-bit)
- Supports ServiceGuard
- Contact an HP storage representative for specific version support.

Windows Server 2003 Enterprise Edition (32-bit), 2000 Server and Advanced Server (SP3, SP4), Windows NT 4.0 SP6A (MSA FW 2.38), MSCS Clusters, Server 2003 (IA-64), Enterprise Edition (64-bit), Datacenter (64-bit)

- MSCS cluster support with up to 8 nodes

MSA1000 FW 4.32 (Alpha servers only), B-Series switches and MP Router, C-Series, and M-Series switches

OpenVMS 7.3-2, 7.3-1, 7.3, 7.2-2

- 7.3-2 requires DEC-AXPVMS-V732_Fibre_SCSI-V0300 (4.32 FW)
- 7.3-1 requires DEC-AXPVMS-V731_Fibre_SCSI-V0600 (4.32 FW)
- 7.3-1 requires DEC-AXPVMS-V731_FIBRE_SCSI-V0400
- 7.3 requires DEC-AXPVMS-V73_Fibre_SCSI-V0700 (4.32 FW)
- 7.3 requires DEC-AXPVMS-V73_FIBRE_SCSI-V0500
- 7.2-2 requires DEC-AXPVMS-V722_Fibre_SCSI-V0600 (4.32 FW)
- 7.2-2 requires DEC-AXPVMS-V22_FIBRE_SCSI-V0400
- OpenVMS Clusters versions 7.3-1, 7.3, and 7.2-2 are supported
- OpenVMS requires a dedicated MSA1000

Tru64 UNIX 5.1A, 5.1B

- 5.1A requires Patch Kit 6 and New Hardware Delivery Kit 6 (NHD7)
- 5.1B requires Patch Kit 4 and NHD7
- TruClusters versions 5.1A and 5.1B are supported
- Tru64 UNIX requires a dedicated MSA1000

Novell NetWare 5.1, 6.0, 6.5

- Supports Novell NetWare Clusters. Contact an HP storage representative for specific version support.

Note: For specific Secure path and Errata support please see the Compatibility guide: http://www.hp.com/go/MSA

Heterogeneous SAN platform interoperability for MSA1000 storage

This section specifies the rules for shared heterogeneous access to a single MSA1000 storage system. MSA1000 storage systems are supported for shared access with combinations of operating systems.

Single MSA1000 shared heterogeneous access:

Any combination of:

- Windows 2003 Enterprise Edition
- Windows 2000 SP3/SP4
- Windows NT 4.0 SP6a
- Novell NetWare 5.1, 6, 6.5
- Red Hat EL 3, Red Hat Linux 7.2 (2.38), Advanced Server 2.1 (4.32)
- SuSE SLES 7 single path
- SLES 8/United Linux 1.0 32-bit and 64-bit (4.32)
  SuSE is Homogeneous when operating in a Secure Path environment.

Additionally, standalone servers and clustered servers are supported on the same MSA1000.

Homogeneous SAN platform support for MSA1000 storage

MSA1000 supports homogeneous access with the following operating systems:

- Tru64 UNIX 5.1A, 5.1B
- OpenVMS 7.3-2, 7.3-1, 7.3, 7.2-2

Each operating system requires a dedicated MSA1000.

Additional rules:

- Tru64 UNIX supports standalone servers, or an up-to-four node cluster
- OpenVMS supports standalone servers, or an up-to-eight node cluster

MSA1000 configuration rules

- MSA1000 storage systems are supported in all SAN fabric topology configurations described in this guide and can be configured in a SAN directly using the switch models shown in Table 10 on page 83 (B-Series), Table 20 on page 99 (C-Series), or Table 28 on page 107 unless otherwise specified. "Specific platform/operating system rules – MSA1500, MSA1000, RA4100, RA4000" on page 184, lists the platforms and operating systems that are supported using these storage systems.
- MSA1000 storage systems with the MSA SAN Switch 2/8 are supported with B-Series product line switches only.
- MSA1000 storage systems are supported with B-Series, C-Series, and M-Series product line switches.
- For standalone server and cluster maximums per MSA1000, see the *MSA1000 Compatibility Guide* available at:

  http://www.hp.com/go/msa1000
- Multipathing with Linux, Novell NetWare, and Microsoft Windows is supported.
- Attaching non-Secure Path (single HBA) servers to an MSA1000 with dual controllers with servers with Secure Path (dual HBA) attached, is supported with Microsoft or Novell operating systems. In the event of a controller fail-over, (failure of active controller) the single path servers will lose access to their data on the MSA1000.

Note: If Secure Path for Linux is used on any node or cluster attached to an MSA1000, all nodes must also have Secure Path installed, regardless of operating systems. See http://www.hp.com/go/securepath for the latest Secure Path parameters.

- Use ACU to enable/disable LUN access to specific connections
- Zoning may be required to prevent access from servers to multiple storage system types when configuring these storage systems in the same physical SAN, or for access from the same server with other storage system types. See "Common server access", page 135 for specific configurations where common server access is supported without the need for zoning or separate HBAs.

  Storage solutions utilizing multiple storage types, such as disk and tape, may specify support for common server access. In those cases, see the specific storage solution documentation for the supported common access configurations and rules.

MSA1000 maximums

The following table lists the maximum configurations for MSA1000 systems.

Table 60: MSA1000 maximum configurations

Platform or operating system	Host bus adapters per server	Active controller ports (targets) per HBA	LUNs per HBA target	Port maximums HBAs per active controller port
Microsoft Windows Server 2003 (32/64-bit), Windows 2000 Server SP3/SP4, Advanced Server SP2 MS Windows Server 2003 Standard Edition and Enterprise Edition (32-bit) Windows NT 4.0 SP6a Red Hat EL 3, Professional v7.2, Advanced Server 2.1, SuSE SLES 7, SLES 8/United Linux 1.0 Novell NetWare 5.1, 6, 6.5 Clusters 1.01, 1.06	2	8	32	32
Tru64 UNIX 5.1A, 5.1B OpenVMS 7.3-2, 7.3-1, 7.3, 7.2-2	Server Dependent			

Heterogeneous SAN platform interoperability for RA4100/RA4000 storage systems

This section specifies the rules for shared access to a single RA4100 or RA4000 storage system. RA4100/4000 storage systems are supported for shared access with any combination of the following operating systems:

- Linux Red Hat 7.0, 7.1
- Linux SuSE 7.1
- Microsoft Windows 2000 Server, Advanced Server SP1
- Microsoft Windows NT 4.0, SP5, SP6a
- Novell NetWare 5.1

- RA4100/RA4000 systems can not be shared by more than one cluster when using Microsoft Windows NT 4.0 or Microsoft Windows 2000.
- RA4100/RA4000 systems owned by a Microsoft Windows NT or Microsoft Windows 2000 cluster can not be shared with a standalone server or server.

RA4100 and RA4000 configuration rules

- These storage systems can be configured in a SAN directly using the switch models shown in Table 10, and through the FC-AL Switch 8 cascaded to the other switch models listed. The section "Specific platform/operating system rules – MSA1500, MSA1000, RA4100, RA4000" on page 184 lists the platforms and operating systems that are supported using these storage systems.
- Supports all fabric rules for fabrics using the Fibre Channel switches listed in Table 10.
- The FC-AL Switch 8 is supported for cascaded attachment to the SAN through a single FL-Port on B-Series switches only.
- Use single or redundant controllers with Active/Passive controllers.
- Use ACU to enable/disable LUN access to specific connections.
- For RA4100/RA4000 SAN configurations with heavy I/O traffic, you must increase the fabric switch buffer capacity from the default value of 16 to 27.
- Servers accessing RA4100 or RA4000 storage systems must not have access to EVA, HP XP or VA, EMA/ESA12000, EMA16000, MA/RA8000, MA6000, or MSA1000 storage systems. Use zoning to prevent access from servers to multiple storage system types when configuring these storage systems in the same physical SAN.
- Zoning is required with multiple clusters. Each cluster must be in its own zone.

RA4100 and RA4000 maximums

Table 61 lists the maximum configurations for RA4100/RA4000 storage systems.

Table 61: RA4100 and RA4000 maximum configurations

Platform or operating system	Host bus adapters per server	Active controller ports (targets) per HBA	LUNs per HBA target	Port maximum HBAs per active controller port
Red Hat Linux 7.0, 7.1 SuSE Linux 7.1	1	1	32	32
Microsoft Windows 2000 Server, Advanced Server SP2 Windows NT 4.0 SP6a	2	1	32	32
Novell NetWare 5.1	2	1	32	32

SAN/Continuous Access EVA integration

The HP StorageWorks Continuous Access EVA solution is approved for use within a larger Heterogeneous Open SAN provided the following additional rules are followed: All Continuous Access EVA implementations require Level 4 NSPOF SANs using two separate fabrics. See "Data availability" on page 59. For additional information see the *Continuous Access EVA design reference guide* at:

http://h18006.www1.hp.com/products/storage/software/conaccesseva/index.html

The current Continuous Access solution supports a sub-set of those operating systems listed in this guide which limits the type of servers that may reside within the Continuous Access EVA management zone.

1. Shared usage of Continuous Access EVA configured storage systems by non-Continuous Access EVA configured servers (e.g., single HBA or an OS without multipathing support) or non-Continuous Access EVA supported operating systems is not supported.

   Refer to the HP Continuous Access QuickSpec available at http://welcome.hp.com/country/us/en/prodserv/storage.html for supported operating systems.

Note: Contact your HP storage representative for information on specific supported versions of clustering software and Secure Path.

2. Each Continuous Access EVA 3000/5000 solution may contain up to 16 EVAs, where each EVA is limited to at most 256 HBAs, which at 2 HBA per server, equates to 128 servers. Multiple Continuous Access EVA solutions may exist within the same SAN as long as no one solution exceeds the 16 array limit, and that limit is imposed by zoning. Any SANs running Continuous Access EVA have 7 hops between devices connected to B-Series switches, 7 hops with C-Series switches, and 3 hops with M-Series switches with the understanding that there are three links involved. There is the host to local storage link, the local storage to remote storage link, and the local host to remote storage link. Each of these links must not exceed 7 or 3 hops, depending on the switch family. All active/standby host-to-storage links as well as local-to-remote storage links must conform to the 7/7/3-hop limit.

Note: Contact your HP storage representative for information on supported configuration maximums for Continuous Access EVA 4000/6000/8000.

3. A single EVA 3000/5000 storage system may support up to 128 DR groups, and up to 128 copy sets, where a single DR group may contain up to 8 member copy sets. A single EVA 4000/6000/8000 model array may support up to 32 DR groups and up to 32 copy sets, where a single DR group may contain up to 8 member copy sets. On all storage systems, the limit is the total number of DR groups and copy sets that are either a source or destination. When replicating across storage systems with different limits, the lower limit also applies to the replication relationship, that is the storage system pair.
4. The Continuous Access EVA Link supports mixed heterogeneous SAN, DRM, Continuous Access EVA, and Host Based Shadowing traffic.

5. Two Storage Management Appliance Command View element mangers are required, one active and one either active in stand by mode or in powered off passive mode. The active appliance and Command View EVA can be used for initial setup of Continuous Access EVA storage. Management of the operational Continuous Access EVA environment is done through the Continuous Access user interface, and separate product also installed on the storage management appliances. See the Continuous Access EVA Planning Guide and the Continuous Access EVA Operations Guide for additional information.
6. Please see the Continuous Access EVA release notes for current information about the solution. The release notes are available at this URL: http://h18006.www1.hp.com/products/storage/software/conaccesseva/index.html

SAN/DRM integration

The HP Data Replication Manager for HSG80 (DRM) is approved for use within a larger Heterogeneous Open SAN provided the following additional rules are followed: All DRM implementations require Level 4 NSPOF SANs using two separate fabrics. See "Data availability" on page 59. Several special purpose DRM configurations are also supported as defined in the *DRM for HSG80 Design Guide* which is available on the Web at:

http://h18000.www1.hp.com/products/sanworks/drm/documentation.html

Each shared storage array must adhere to the DRM sharing rules as defined in the *DRM Design Guide*. These sharing rules may be more restrictive than those in this guide due to the requirements for DRM, for example, the operating system must support multiple-bus failover. In addition, the current DRM solution supports a sub-set of those operating systems listed in this guide.

1. Shared usage of the DRM configured storage systems by non-DRM configured servers (e.g., running in transparent failover) or non-DRM supported operating systems is not supported.
2. All servers sharing the same storage sub-system must share a compatible SCSI command mode as shown by a yes in the following table:

Table 62: Heterogeneous DRM operating systems

Operating system	Versions	SCSI-2	SCSI-3
HP OpenVMS	7.2-2, 7.3, 7.3-1,	No	Yes
HP Tru64 UNIX	5.1A, 5.1B	Yes	Yes
HP-UX	11.0, 11i v1	Yes	Yes
IBM AIX	4.3.3, 5.1,	Yes	Yes
Microsoft Windows NT	4.0 SP 6a	Yes	Yes
Microsoft Windows 2000 Server, Advanced Server	Server, Advanced Server, Datacenter	Yes	Yes
Novell NetWare	5.1, 6.0, 6.5	Yes	Yes
Sun Solaris	2.6, 7, 8, 9	Yes	Yes

3. Each DRM solution instance may contain up to 96 servers and 8 storage arrays. With large fabrics, multiple solution instances may exist, as long as each is in a separate zone on a SAN or a separate zone within a VSAN (as recommended in C-Series). In other cases the actual limit will be smaller due to restrictions imposed by the intersite link. DRM supports a limit of 7 hops between devices connected to B-Series switches, 7 hops with C-Series switches, and 3 hops with M-Series switches with the understanding that there are three links involved: the host to local storage link, the local storage to remote storage link, and the local host to remote storage link. Each of these links must not exceed 7 or 3 hops, depending on the device. All active/standby host-to-storage links as well as local-to-remote storage links must conform to the 7/3-hop limit.
4. Each DRM solution instance may contain up to 12 servers per storage system per site provided both controllers of the storage system are using the “P” version of ACS configured in remote peer-to-peer replication mode. With 1 remote copy set per server, a maximum of 12 remote copy sets per pair of storage systems, and a maximum of 8 storage systems per site per instance, a single DRM instance can support up to 96 servers per site. If you increase the number of remote copy sets per server, you must reduce the total number of servers per storage system. For example, if you configure 2 remote copy sets per server, the maximum limit is 6 servers per storage system.

5. DRM over ATM configurations are supported for switches listed in Table 10 on page 83, with switch FW 2.1.9m only, with a limit of two Fibre Channel switches per fabric, for a total of 4 switches, one at each end of each fabric (2 fabrics, times 2 switches per fabric equals 4 switches). Cascaded switches are not supported. This no-cascaded switch restriction also includes non-support for the SAN Switch Integrated/32 or SAN Switch Integrated/64 port switches due to the fact that these switch models are made up internally of 6, 16-port switches that are cascaded together.
6. The DRM Link supports mixed heterogeneous SAN, DRM, and Host Based Shadowing traffic.
7. StorageWorks Command Console (SWCC) and the Storage Management Appliance (SMA) element manger can be used for initial setup of Data Replication Manager (DRM) storage sub-systems. However, neither of these tools should be used for DRM failover and failback operations. Therefore to prevent any potential inference by MA polling of the HSG80 when running DRM scripts, HP recommends that the MA be removed from all DRM zones before running the scripts.
8. Please see the DRM release notes for current information on any hop count restrictions between devices.

SAN/DRM/OpenVMS host based volume shadowing integration

OpenVMS servers implementing Host Based Volume Shadowing are supported integrated in a heterogeneous SAN with remote shadowset distances of up to 200 km over 1 Gbps direct fiber. The direct fiber long distance link supports mixed heterogeneous SAN, DRM, Continuous Access EVA, and OpenVMS Host Based Volume Shadowing traffic.

StorageWorks CSS 2105 storage system interoperability and integration

HP provides support for heterogeneous multi-vendor online storage interoperability on a common SAN. This support includes both the StorageWorks Centralized Shared Storage 2105 (CSS 2105) and the StorageWorks Enterprise RAID Array.

The initial integration support represents the first phase or level of interoperability. This level of support provides for:

1. Coexistence of HP and IBM storage systems in a common heterogeneous Open SAN. The HP and IBM storage systems operate in separate fabric zones within the same physical SAN.
2. Data migration support between HP and IBM storage systems using a shared server running either Windows 2000/NT, IBM AIX, or Sun Solaris.
3. Multi-path failover capabilities using HP Secure Path for the HP storage and the IBM Subsystem Device Driver on a single shared server running Windows NT, IBM AIX, or Sun Solaris. Each storage system is connected to the server using independent HBA pairs.
4. Simultaneous enterprise backup support from both the HP storage and IBM storage utilizing a single shared server and the HP Enterprise Backup Solution with VERITAS NetBackup to a common tape library for Windows 2000/NT.

See Technical Note *Compaq StorageWorks Centralized Shared Storage 2105 Interoperability* for additional information.

High-availability configuration considerations

Cabling scheme options

This section describes cabling scheme options for implementing high availability multi-path configurations for EVA5000/EVA3000, MA6000, MA/RA8000, EMA/ESA12000, and EMA16000 storage systems. Figure 52 and Figure 53 show cabling options when implementing a Level 4 high availability no-single-point-of-failure configuration. Figure 54 and Figure 55 show the cabling and associated zoning requirements when implementing a Level 3 high-availability configuration. See "Data availability" on page 59 for a description of the availability levels.

Note: DRM requires the high availability NSPOF configuration. DRM cabling is fully described in the DRM Design Guide, available at:
http://h18000.www1.hp.com/products/sanworks/drm/documentation.html

Note: Continuous Access EVA cabling is supported, as shown in Figure 52 and Figure 72 on page 298. The cabling is also described in the Continuous Access EVA design reference guide:
http://h18006.www1.hp.com/products/storage/software/conaccesseva/index.html

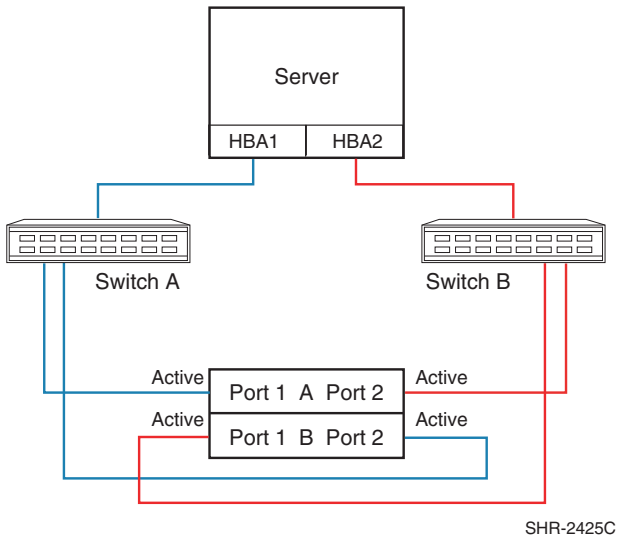

Figure 52: Cross-cable high-availability NSPOF configuration

Figure 52 shows the physical connections for a cross cable, high availability, no-single-point-of-failure configuration for storage systems using two separate fabrics. The advantage of this cabling scheme is that it is the same cabling scheme used for Continuos Access EVA. This allows you to upgrade to a Continuous Access solution without the need to re-cable the controller connections.

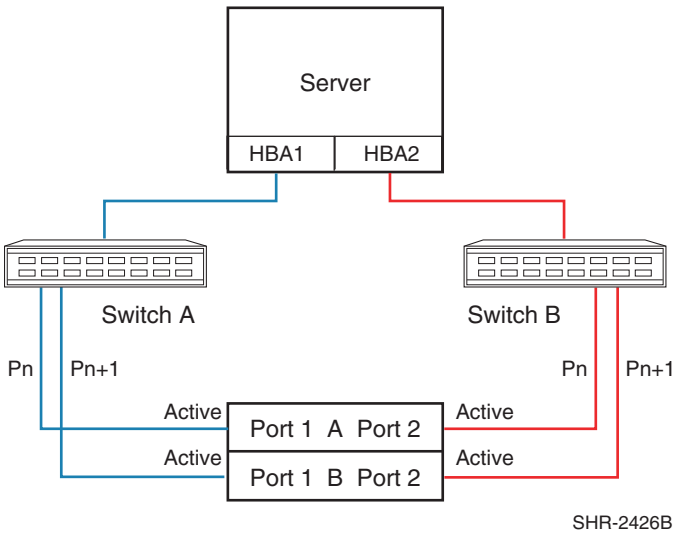

Figure 53: Straight-cable high-availability NSPOF configuration

Note: The straight-cable high-availability configuration is not supported with Continuous Access EVA.

Figure 53 shows the physical connections for a straight cable, high-availability, no-single-point-of-failure configuration. The advantage of this cabling scheme is that it is the same cabling scheme used in Transparent failover mode for MA6000, MA/RA8000, EMA/ESA12000, and EMA16000 storage systems. This allows you to migrate from Transparent failover mode to Multiple-Bus failover mode without the need to re-cable the controller connections.

Figure 54 and Figure 55 specify the logical path zoning that may be required for cross-cable and straight-cable configurations when implementing a level 3, single-fabric high-availability configuration. The requirement to zone separate logical paths in single-fabric, high-availability implementations is O/S and platform specific. The zoning specified enforces and effectively results in the same configuration as physically depicted in Figure 52 and Figure 53. Single-fabric cross-cable implementations require cross-port zoning, and straight-cable implementations require straight port zoning. In order to provide high availability, ensure each HBA is cabled to a different switch and configured for access to specific controller ports.

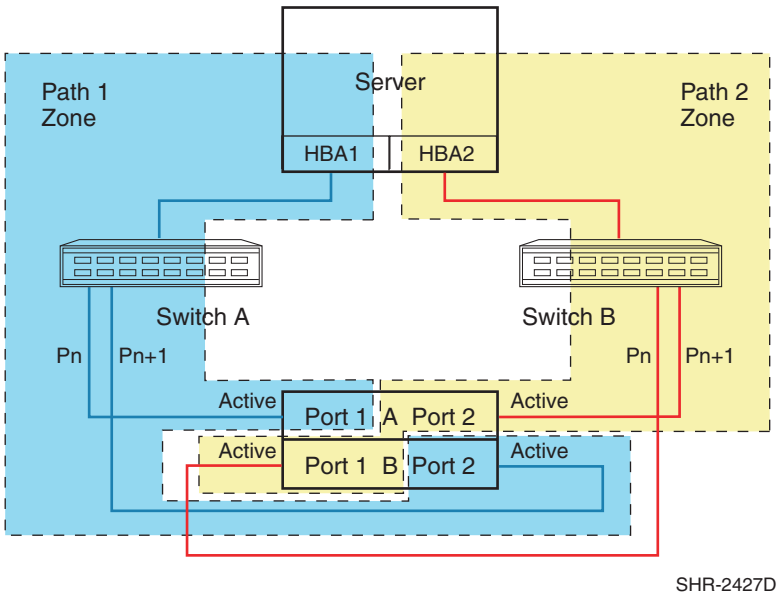

Figure 54: Cross-cable high availability single fabric zoned configuration

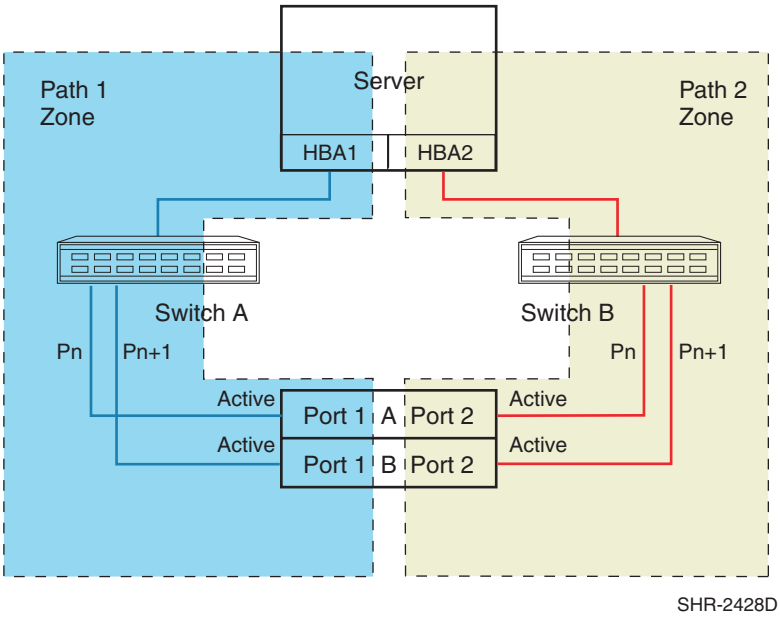

Figure 55: Straight-cable high-availability single fabric zoned configuration

Note: The straight-cable high-availability configuration is not supported with Continuous Access EVA.

For two or more high availability server configurations, it is suggested that the first adapter in each server be connected to the first (same) Fibre Channel switch, the second two adapters to the second switch, etc. For example:

- Server 1 Fibre Channel HBA 1 to Fibre Channel Switch 1 - Switch Port 1
- Server 1 Fibre Channel HBA 2 to Fibre Channel Switch 2 - Switch Port 1
- Server 2 Fibre Channel HBA 1 to Fibre Channel Switch 1 - Switch Port 2
- Server 2 Fibre Channel HBA 2 to Fibre Channel Switch 2 - Switch Port 2

HP highly recommends that the cabling scheme shown in each Secure Path multiple-bus configuration be followed as depicted. This is not required; however, it does aid in understanding logical to physical LUN and path mapping for maintenance purposes.

Cabling scheme options for dual channel HBAs

Dual channel HBAs are typically utilized in situations where the number of server PCI slots is limited. As such, most installations are configured as shown in either Figure 56 or Figure 57. Both configurations are implemented using a single PCI slot to provide access to either the same Targets/LUNs or a different set of storage Targets/LUNs through separate ports on the HBA.

Note: Each dual channel HBA is theoretically capable of twice the performance of a single channel HBA for a given single PCI slot.

Note: Target ranges are shown for example purposes; the number of storage controller Targets and LUNs associated with each Target accessible is operating system dependent.

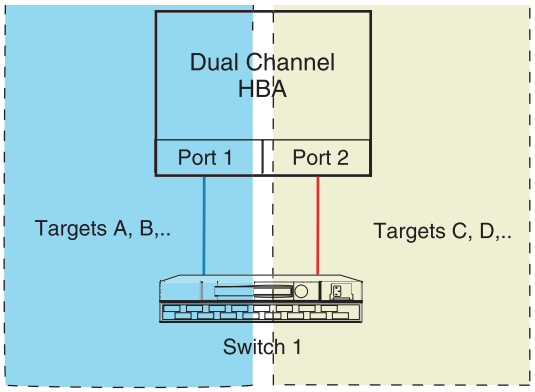

Figure 56: Single PCI slot with dual channel HBA and one switch

Figure 56 shows connectivity with both HBA paths connected to the same Fibre Channel switch.

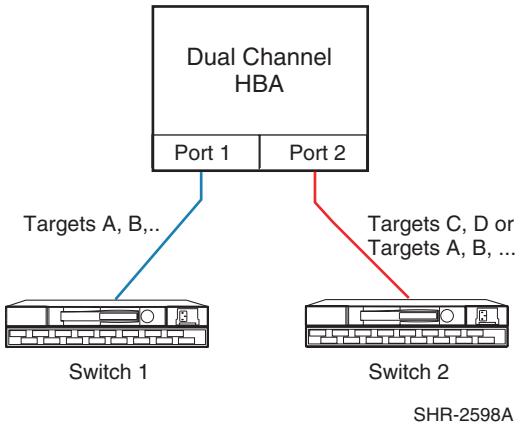

Figure 57: Single PCI slot with dual channel HBA and two switches

Figure 57 shows increased availability over Figure 56 in the event of a single switch failure. Availability to a specific set of Targets/ LUNs can be further increased by configuring access to the Targets (A, B.) on both paths as shown.

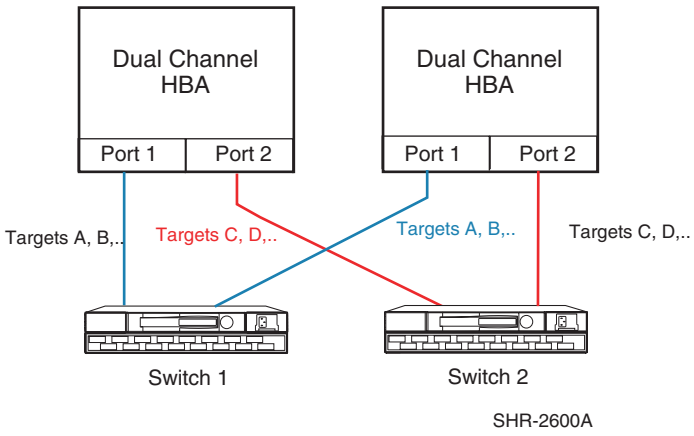

Figure 58: Two PCI slots with dual channel HBAs - NSPOF

Figure 58 shows an example of how an NSPOF solution can be implemented with two dual-channel HBAs. From an availability standpoint, this equates to the use of two single-channel HBAs. See "Data availability" on page 59 for further information.

Enterprise Backup Solution

10

One of the most significant benefits of a SAN is the ability to share the SAN infrastructure for both disk and tape. With a SAN backup solution, all the benefits of the SAN (such as, consolidated storage, centralized management, and increased performance) can be gained for the backup operations. Additionally, implementing a SAN backup solution lays the foundation for advanced data protection features such as serverless backup and backup to disk. HP's SAN backup solution is the HP StorageWorks Enterprise Backup Solution (EBS).

The first step in deploying an Enterprise Backup Solution is to design the backup SAN configuration. Consult the HP StorageWorks Backup Sizer Tool and the EBS Compatibility Matrix, which are available at:

http://h18000.www1.hp.com/products/storageworks/tapecompatibility.html

The HP StorageWorks Enterprise Backup Solutions Design Guide is the second step in configuring your Enterprise Backup Solution. This guide describes the EBS hardware configurations currently supported and how to efficiently and effectively provide shared tape library backup in a heterogeneous SAN environment.

http://h18004.www1.hp.com/products/storageworks/ebs/documentation.html

The third step in implementing your Enterprise Backup Solution is installing and configuring your backup application. Rules and recommendations for individual backup applications may be found in separate implementation guides.

For more information about EBS, see:

http://www.hp.com/go/ebs

Volume 4

SAN extension and bridging

SAN extension and bridging are presented in these chapters:

SAN extension

11

With the advent of extension technologies specifically developed for the transport of data, it is possible to consolidate, simplify, manage and integrate storage in Fibre Channel SAN fabrics within the enterprise to further exploit networking investments and lower the cost to manage global storage.

A SAN extension is considered an interswitch link (ISL) connection between two Fibre Channel switches over extended distances. Extended distances are considered to be:

- 150 meters for 4 Gb/s Fibre Channel ISLs
- 300 meters for 2 Gb/s Fibre Channel ISLs
- 500 meters for 1 Gb/s Fibre Channel ISLs
- Any distance between a pair of Fibre Channel over IP products

Whether it's called SAN Extension or SAN Bridging, HP seamlessly integrates these new technologies into the benefits of HP Fibre Channel SANs.

This chapter describes the current HP supported technologies and products available for heterogeneous SAN Extension, and the current HP provided products for Continuous Access and DRM SAN Extension. For additional information on disaster recovery SAN extension and a complete list of all supported third party extension products, please read the section "SAN/Continuous Access EVA integration" or "SAN/DRM integration" in Chapter 9, and see the *HP StorageWorks Continuous Access and Data Replication Manager SAN extensions reference guide* available at:

http://h18000.www1.hp.com/products/storageworks/san/documentation.html

This chapter covers the following major topics:

Why extend the SAN?

The growing need for storage data that is permeating the business community, coupled with the available bandwidth afforded by IP networks or WDM, for example, are making SAN extension an increasingly attractive option to grow the storage network. With SAN extension, end users can connect to data centers at opposite ends of a campus, metropolitan, and wide-area environment. The challenge is to do so at full-wire speed, with the same reliability and availability as the storage traffic within each data center.

Supported SAN extension technologies

HP supports the following technologies for Fibre Channel ISL SAN extension.

- Fibre Channel Long Distance Technologies
  - Long Wave Transceivers
  - Wavelength Division Multiplexing (WDM)
- IP Data Protocol Technologies
  - Fibre Channel over Internet Protocol (FCIP) using the SR2122-2 IP Storage Router
  - FCIP using the Cisco MDS 4/8-Port IP Storage module
  - FCIP using the Cisco MDS 14/2-Port Multiprotocol Services module
  - FCIP using the B-Series MP Router
  - FCIP using the SAN Valley SL700/SL1000 IP-SAN Gateway

Note: For additional SAN Extension products supported for DRM or Continuous Access see the *HP StorageWorks Continuous Access And Data Replication Manager SAN Extensions* available at: http://h18000.www1.hp.com/products/storageworks/san/documentation.html

Supported SAN bridging technology

- iSCSI to Fibre Channel Bridging using the SR2122 iSCSI Storage Router
- iSCSI to Fibre Channel Bridging using the SR2122-2 IP Storage Router
- iSCSI to Fibre Channel Bridging using the Cisco MDS 4/8-Port IP Storage module
- iSCSI to Fibre Channel Bridging using the Cisco MDS 14/2-Port Multiprotocol Services module

For information on bridging with iSCSI, see “iSCSI storage” on page 235.

Note: Not all technologies are supported by all HP Fibre Channel switch product lines. Please read each technology description for further details.

Fibre Channel long distance technologies

Long wave transceivers

Fibre Channel switches use two types or styles of fiber-optic transceivers that come in both short wave and long wave varieties. The 1-Gbps transceivers use “SC” style connectors that are known as Giga-Bit Interface Converters, or GBICs. The 2 Gbps transceivers use the “LC” style connectors that are known as Small Form Factor Pluggable transceivers, or SFPs. Long wave GBIC or SFP transceivers are required to go beyond the 500 meter limit for 1 Gbps and the 300 meter limit for 2 Gbps links respectively. There are long-wave optical transceivers that are capable of transmitting up to 100 km.

HP supports the following long wave transceivers:

- 10 km GBIC
- 100 km GBIC
- 10 km SFP
- 35 km SFP
- 100 km CWDM SFP (See “Wavelength division multiplexing” on page 211.)

Long wave transceivers are supported on HP B-Series, HP C-Series, and HP M-Series product lines. B-Series Fibre Channel switch products support 10 km and 100 km GBICs (certain switch models), 10 km and 35 km SFPs, and 100 km Coarse Wave Division Multi-plexing (CWDM) SFPs. The B-Series MP Router supports 10 km and 35 km SFPs. C-Series Fibre Channel switch products support 10 km SFPs and 100 km CWDM SFPs. M-Series Fibre Channel switch products support 10 km and 35 km SFPs.

Wavelength division multiplexing

Wavelength Division Multiplexing devices can be used to extend the distance between two Fibre Channel switches. These devices are transparent to the switches themselves and do not count as an additional hop. The only consideration that should be made to accommodate these devices is to have enough buffer-to-buffer credits in order to maintain line speed performance. Wavelength Division Multiplexing is supported for both 1 Gbps and 2 Gbps. This technology is ideally suited for metro data center deployments. When designing SAN extension across an optical ring, buffer-to-buffer credits become a very important consideration. In many WDM ring designs, the recovery path due to a link failure can be significantly longer distance than the primary path due to routing the traffic in the opposite direction around the ring. It is important to consider the distance over primary and recovery paths to ensure enough buffer-to-buffer credits exist for both so as not to impede performance during a ring fault event.

HP offers a CWDM technology solution which involves similar concepts as Dense Wave Division Multiplexing (DWDM) but is less expensive, less expandable (maximum 8 channels) and works over a distance of 100 km. CWDM allows up to eight 1 Gbps or 2 Gbps channels (or colors) to share a single fiber pair. Each channel uses a different color or wavelength transceiver. These channels are networked with a variety of wavelength specific add-drop multiplexers to enable an assortment of ring or point-to-point topologies.

See the individual switch product line WDM sections in this chapter for additional information about WDM support.

Note: HP supports the use of all WDM products as Fibre Channel ISLs provided the WDM equipment is configurable to 1 Gbps or 2 Gbps data rates, and does not implement time division multiplexing or any additional conversion method that alter the data links other than multiplexing different wavelengths.

The HP CWDM solution consists of the following components:

- 2-slot chassis for optical Add/Drop Multiplexers (OADMs)
- CWDM OADMs
  - Eight single-channel OADMs
  - Two 4-channel OADMs
  - 8-channel multiplexer/demultiplexer
- CSDM SFPs (1470nm, 1490nm, 1510nm, 1530nm, 1550nm, 1570nm,1590nm, 1610nm)

A typical CWDM installation will include (2) OADMs, matched pairs (same frequency) of CWDM SFPs, (4) or (8) single mode fiber optic cables, and a single long distance fiber optic cable. See the C-Series CWDM product documentation for more information at:

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

Maintaining performance beyond 5 or 10 km

A primary consideration with extended fabrics is maintaining the performance of the interswitch link - connection(s) between a pair of switches. The flow control mechanism for a Fibre Channel connection is buffer-to-buffer credits. The number of credits a port has is equal to the number of frames a port can transmit before getting an acknowledgement that the frame was received.

At the speed of light in a fiber-optic cable, it takes a full second for light to travel 200,000 km or 5 microseconds per km. If you calculate the time it takes a frame to travel 100 km and for the "RRDY" (buffer available - analogous to a frame acknowledgement) to travel back the same 100 km at 1 Gbps you need about 60 buffer-to-buffer credits to keep the link running at full speed. The rule-of thumb in Fibre Channel is that to sustain 1 Gbps of bandwidth for full 2148B frames approximately one buffer-to-buffer credit is required for every 2 km of distance between two interfaces on a link. For a 2 Gbps link, one buffer-to-buffer credit is required for every 1 km of distance between two interfaces on a link. For smaller frame sizes, the number of buffer-to-buffer credits that are required increases.

There are different limits on the extended link parameters as well as the maximum number allowed across all HP switch product lines. In addition, the commands to configure the buffer-to-buffer credits for each switch product line also vary. The following sections detail these limits and the procedures for configuring extended links for each of the HP switch product lines.

HP B-Series product line

Extended fabric limits using WDM

WDM is supported on both 1 Gbps and 2 Gbps switch models.

The maximum number of hops allowed in an B-Series product line Fabric is 7, with a maximum total distance of 200 km across the SAN between any two devices.

B-Series switches are supported with the HP CWDM solution up to a maximum segment distance of 100 km. The maximum supported distance for other vendor WDM equipment is based on the specific WDM product used. In all cases the maximum total distance or total of all segments combined is 200 km.

The B-Series Trunking feature is supported with WDM on trunk ports configured for the same speed and distance setting that are relatively close in distance (FW v4.4 and up.) For FW versions prior to 4.4, then trunking is supported up to a distance of 5 km with no more than 30m difference between trunk ports.

Note: HP supports up to 200 km across a WDM link at 1 Gbps Fibre Channel with reduced performance levels. The performance levels attained are dependent on the number of buffers available in the particular switch models used, and the specific application data transfer size. In configurations where the WDM link is 200 km, there can be no other long wave segments.

Extended fabric compatibility support

HP has three series of switches in the B-Series product line as listed below; these switches can be divided into two classes based on the internal ASIC technology used in the switch. The two classes are switches limited to 1 Gbps and those that are 2 Gbps capable.

- StorageWorks 1 Gbps SAN switch series with version 2.x installed,
- StorageWorks 2 Gbps SAN switch series with version 3.x installed, or
- HP StorageWorks SAN Switch 2/32, Core Switch 2/64, and SAN Director 2/128 switches with version 4.x installed.

An extended fabric link (a link >5 km at 2 Gbps or >10 km at 1 Gbps) can only exist between two switches of the same technology, meaning a B-Series 1 Gbps only switch can only have an extended fabric link to another B-Series 1 Gbps only switch. Likewise a B-Series 2 Gbps capable switch can only have an extended fabric link to another B-Series 2 Gbps capable switch regardless of the link speed.

ISL connections up to 10 km are supported between 1 Gbps only and 2 Gbps capable switches at the "L0" portcfglongdistance setting only.

“portcfglongdistance” settings

Extended Fabric optimizes the internal buffering algorithm for StorageWorks switches, which results in line speed performance of close to full Fibre Channel speed. The “portcfglongdistance” setting is used to configure the port with the appropriate amount of buffers based on the speed and distance of the extended link.

The possible settings are:

- L0: 10, 5, and 2.5 km at 1, 2, & 4 Gbps respectively
  No Extended Fabric license required
- L0.5 >10 km up to 25 km
  Extended Fabric license required
- L1: >25 km up to 50 km
  Extended Fabric license required
- L2: >50 km up to 64 km (V3.x) or 100 km (V4.x)
  Extended Fabric license required
- LE: E-Ports between 5 and 10 km
  No Extended Fabric license required

- LD: 400, 200, and 100 km at 1, 2, and 4 Gbps respectively (Model 4/32)
  100 km (V3.x) or 200 km (V4.x) at 1 Gbps
  50 km (V3.x) or 100 km (V4.x) at 2 Gbps
  Extended Fabric license required

These port settings modify the number of Buffer-To-Buffer credits a particular port is allocated and there are limited numbers of these credits available. Buffer-To-Buffer credits are allocated to a group of 4 ports or what is referred to as a "Quad". A quad consists of ports 0 through 3, 4 through 7, 8 through 11, 12 through 15 and so on.

The following table lists the configuration limits for a "Quad". The rows are mutually exclusive options for the specified fabric OS.

Table 63: Long distance port matrix

Fabric OS	Speed	Port A	Port B	Port C	Port D
HP StorageWorks FOS versions: 2.x	1 Gbps	L2	E/L1	LE/L0.5/Fx	Disabled
	1 Gbps	L2	L0.5	L0.5/LE/Fx	Disabled
	1 Gbps	L2	L0.5	LE/Fx	LE
	1 Gbps	L2	LE/Fx	LE/Fx	LE/Fx
	1 Gbps	E/L1/L0.5/ LE/Fx	E/L1/LE/L0.5/ Fx	E/L1/LE/L0.5/ Fx	E/L1/LE/L0.5/ Fx
	1 Gbps	LD	LD	LD	LD
HP StorageWorks FOS versions: 3.0, 3.0.1, 3.0.2, 4.0, 4.0.2	1 Gbps	L2	E/L1	Fx	Disabled
	1 Gbps	L2	Fx	Fx	Fx
	1 Gbps	E/Fx/L1	E/Fx/L1	E/Fx/L1	E/Fx/L1
HP StorageWorks FOS versions: 3.0, 3.0.1, 3.0.2, 4.0, 4.0.2	2 Gbps	L2	Disabled	Disabled	Disabled
	2 Gbps	L1	L1	Disabled	Disabled
	2 Gbps	L1	E	E/LE/Fx	Disabled
	2 Gbps	L1	LE/Fx	LE/Fx	Fx
	2 Gbps	E/LE/Fx	E/LE/Fx	E/LE/Fx	E/LE/Fx
HP StorageWorks FOS Version: 3.1, 4.1, and 4.2	1 Gbps	L2	E/L1	LE/L0.5/Fx	Disabled
	1 Gbps	L2	L0.5	LE/L0.5/Fx	Disabled
	1 Gbps	L2	L0.5	LE/Fx	LE
	1 Gbps	L2	LE/Fx	LE/Fx	LE/Fx
	1 Gbps	E/L1/L0.5/ LE/Fx	E/L1/L0.5/LE/ Fx	E/L1/L0.5/LE/ Fx	E/L1/L0.5/LE/ Fx
	1 Gbps	LD	LD	LD	LD

Table 63: Long distance port matrix (Continued)

Fabric OS	Speed	Port A	Port B	Port C	Port D
HP StorageWorks FOS Version: 3.1, 4.1, and 4.2	2 Gbps	L2	E	Fx	Disabled
	2 Gbps	L2	LE/Fx	LE/Fx	Disabled
	2 Gbps	L2	L0.5	Disabled	Disabled
	2 Gbps	L1	L1	Disabled	Disabled
	2 Gbps	L1	E	E/LE/Fx	Disabled
	2 Gbps	L1	LE/Fx	LE/Fx	Fx
	2 Gbps	L1	L0.5	LE/Fx	Disabled
	2 Gbps	L0.5	L0.5	L0.5	Disabled
	2 Gbps	L0.5	E/L0.5/LE/Fx	E/LE/Fx	Disabled
	2 Gbps	L0.5	E/L0.5/LE/Fx	LE/Fx	LE/Fx
	2 Gbps	L0.5	E/LE/Fx	E/LE/Fx	LE/Fx
	2 Gbps	E/LE/Fx	E/LE/Fx	E/LE/Fx	E/LE/Fx
	2 Gbps	LD	LD	LD	LD

Fx = Fabric port L0, LE, L0.5, L1, L2 = interswitch links

Fabric long distance bit setting

The Fabric Long Distance Bit needs to be set on all switches in the fabric when any pair of StorageWorks 1 Gbps SAN series switches has an extended link greater than 10 km (portcfglongdistance = L0.5, L1, or L2). This bit sets fabric wide parameters so that all switches know how to use the legacy method to calculate the number of buffer-to-buffer credits.

Whenever a pair or pairs of StorageWorks 2 Gbps SAN series or an HP StorageWorks core switch 2/64 switches have a port configured for LE, L0.5, L1, or L2 than the Fabric Long Distance Bit must be off. In other words you cannot have an extended link of greater than 10 km between a pair of StorageWorks 1 Gbps SAN series switches and an extended link greater than 5 km between a pair of StorageWorks 2 Gbps SAN series or an HP StorageWorks core switch 2/64 switches in the same fabric.

You can have extended links of up to 10 km (portcfglongdistance = L0) between a pair of StorageWorks 1 Gbps SAN series switches and up to 200 km extended link between a pair of StorageWorks 2 Gbps SAN series or an HP StorageWorks core switch 2/64 switches in the same fabric. Likewise, you can have an extended link over 10 km between StorageWorks 1 Gbps SAN series switches as long as there are no interswitch link (ISL) connections greater than 5 km between a pairs of StorageWorks 2 Gbps SAN series or an HP StorageWorks core switch 2/64 switches in the same fabric.

HP C-Series product line

Extended fabric limits using WDM

WDM is supported on all C-Series switches at both 1 Gbps and 2 Gbps speeds.

The maximum number of hops allowed in C-Series product line fabric is 7, with a maximum total distance of 200 km across the SAN between any two devices.

C-Series switches are supported with the HP CWDM solution up to a maximum segment distance of 100 km. The maximum supported distance for other vendor WDM equipment is based on the specific WDM product used. In all cases the maximum total distance or total of all segments combined is 200 km.

Note: HP supports up to 200 km across a WDM link at 1 Gbps Fibre Channel with reduced performance levels. The performance levels attained are dependent on the number of buffers available in the particular switch models used, and the specific application data transfer size. In configurations where the WDM link is 200 km, there can be no other long wave segments.

See the C-Series switch product documentation on the HP Storage web page for more information.

Extended fabric compatibility support

All C-Series switches are compatible from a functionality perspective. The long distance Fibre Channel connection can be formed between any two C-Series Fibre Channel switch models. All C-Series products support up to 255 buffer-to-buffer credits for extended distance configurations.

In C-Series, each port on the 16-port line card supports 255 buffer-to-buffer credits that are available on a per-port basis for an ISL when using either a single link or a link aggregated via Port Channel. A Port Channel forms a logical ISL and can bundle up to 16 x 2Gbps links to form a single 32 Gbps link. All links within the Port Channel must be the same speed.

The MDS 9216i and 14/2 products support 3,500 buffer-to-buffer credits.

HP M-Series product line

Extended fabric limits using WDM

WDM is supported on both 1 Gbps and 2 Gbps switch models.

The maximum number of hops allowed in a M-Series product line Fabric is 3, with a maximum total distance of 200 km across the SAN between any two devices.

Note: HP supports up to 200 km across a WDM link at 1 Gbps Fibre Channel with reduced performance levels. The performance levels attained are dependent on the number of buffers available in the particular switch models used, and the specific application data transfer size. In configurations where the WDM link is 200 km, there can be no other long wave segments.

Note: The HP CWDM solution is not supported on M-Series switch products.

HP StorageWorks Edge Switch 2/24 limits

The HP StorageWorks edge switch 2/24 has a fixed buffer-to-buffer credit setting and are limited in which ports can support links beyond the 500 meter limit for 1 Gbps and 300 meter limit for 2 Gbps links.

Ports 0 through 3 are capable of supporting distances up to 20 km at 1 Gbps and up to 10 km at 2 Gbps per second.

Ports 4 through 23 do not have enough buffer-to-buffer credits to support long wave SFP transceivers and are limited to the short wavelength SFP transceiver limits of 500 meters at 1 Gbps or 300 meters at 2 Gbps.

10-100km port setting

In order to maintain line speed performance of close to full Fibre Channel speed for extended lines over 10 km, it is necessary to configure the applicable ports for 10-100km setting. Using the High Availability Fabric Manager (HAFM), select the configure ports menu option and then click on the 10-100 km box for the applicable ports. This will increase the number of buffer-to-buffer credits from 16 to 60 for the selected port.

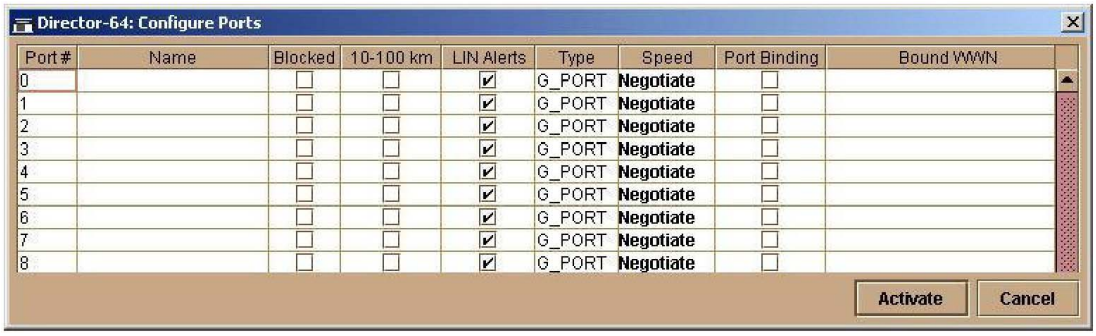

Figure 59: HAFM configure ports for 10-100 km setting

TCP/IP data protocol technologies

Fibre Channel over Internet Protocol (FCIP)

FCIP is a protocol that encapsulates Fibre Channel frames into IP packets and tunnels them through an existing IP network infrastructure to transparently connect two or more SAN fabrics together. The IP tunnel acts as a dedicated link to transmit the Fibre Channel data stream over the IP network, while maintaining full compatibility with the Fibre Channel SAN.

FCIP Gateways perform Fibre Channel encapsulation process into IP Packets and reverse that process at the other end.

Fibre Channel switches (B-series and M-series) connect to the FCIP gateways through an E_Port for SAN fabric extension to remote locations. C-series switches use plug-in modules for FCIP functionality.

A tunnel connection is set up through the existing IP network routers and switches across LAN/WAN/MAN.

This example shows a configuration that connects Fibre Channel SANs using an Internet Protocol (IP) intersite link.

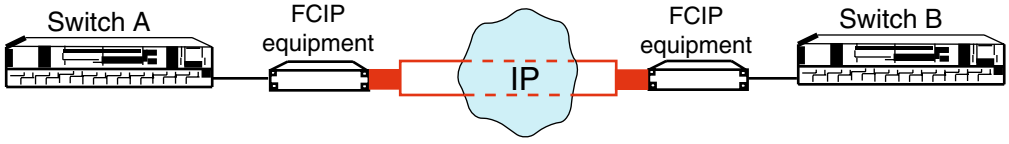

Figure 60: Connecting Fibre Channel SANs with an IP link

Using Internet Protocol over an IP-based network, FCIP can link sites over extended distances. Local SANs connect through an IP network to create an extended SAN. An FCIP gateway is used at each end of the intersite link. Each FCIP gateway box encapsulates received Fibre Channel frames into IP packets for transmission over the IP network. Similarly, the FCIP box extracts the original Fibre Channel frame from received IP packets and retransmits them to the destination Fibre Channel node. The FCIP boxes also handle IP-level error recovery.

Note: The gateways at either end of the link must be from the same gateway family to insure interoperability. See the *Continuous Access Data Replication Manager SAN extension reference guide* for details.

FCIP Products supported for heterogeneous SAN extension

The HP StorageWorks SR2122-2 IP Storage Router provides heterogeneous SAN FCIP extension support on HP Series B-Series, C-Series, and M-Series Fibre Channel switches.

The HP StorageWorks MP Router provides heterogeneous SAN FCIP extension support on HP B-Series Fibre Channel switches.

The Cisco MDS 4/8-Port IP Storage Module and MDS 14/2-Port Multiprotocol Services Module are supported with C-Series product line switches.

The SAN Valley SL700/SL1000 IP-SAN Gateways are supported with B-Series and M-Series product line switch models for heterogeneous SAN FCIP extension. Please read the manufacturer's documentation for further configuration details.

IP network considerations

Considerations relevant to using the existing IP network

The ability to use your existing network with FCIP depends on the type of storage I/O you plan to do and the traffic already existing on your current network. The key consideration is whether you have enough unused/available bandwidth from your network to continue the current network load, accommodate future growth, and handle FCIP SAN load demands.

Table 64: IP network issues to consider

Storage I/O type	Use existing IP network?	Factors
Mirrored I/O or continuous I/O throughput over the FCIP intersite link.	A separate network is recommended.	For peak performance for your current network, and for peak Storage I/O performance, a separate network is recommended.
Data Migration or Adhoc Data Updates	The use of your existing network may be possible.	Data migration is a one-time movement of data for upgrade or maintenance purposes. Ad hoc Data Updates is more of a 'burst' of data from one site to another for remote backups, database content delivery, etc. It is possible to use your existing network; however, the network performance may be significantly affected.

Network speeds

In general, the FCIP equipment supports Ethernet throughput of 10/100 Mb/s, and 1 Gbps (Gigabit Ethernet). The network connection should be selected to match the amount of data to be transferred.

The speed of light through fibre is approximately 200,000 km per second or 5 microseconds to travel one km.

Network distance considerations

The HSG80 controller uses SCSI protocol to manage the storage devices. Before a SCSI I/O can be transmitted, it must be encapsulated into Fibre Channel frames. Because of SCSI protocol, a minimum of 4 trips over the long-distance link is required.

These trips conceptually:

1. Tell the remote site you want to transmit data.
2. Wait for the acknowledgment from the remote site.
3. Send the data to the remote site.
4. Wait for the acknowledgment from the remote site.

When sending data over fiber, the one-way transmission time is approximately 5 microseconds per km of optical cable. Because a minimum of four trips is required for each SCSI data transfer, this translates to a total transmission delay per command of 20 microseconds per km, or about 32.2 microseconds per mile. For example, if a remote site is located 150 miles away from the local site, the total time will be 4,830 microseconds (4.83 milliseconds) for every data transfer. Because a typical I/O operation on a non-DRM configuration with write-back cache takes approximately 500 microseconds, long distances can have a significant effect on performance.

Note: The above calculations for a link of 150 miles do not include any latency induced by the FC-to-IP conversions, or latency of the routers and switches in the network.

Additional I/Os, either from additional LUNs on the same controller or from a different controller, will require additional bandwidth. Care must be taken to understand this principle. Adding bandwidth to a given link at a given distance will not increase the time it takes to complete an I/O operation. It will, however, allow you to add additional I/Os from different LUNs, thereby consuming the available bandwidth.

Conversely, if enough bandwidth is not given to a link, then the number of I/Os per second will decrease, possibly to the point of failure.

Note: The time it takes an I/O to complete an operation is more complex than the above example, and there are additional factors involved with this calculation. This discussion is an attempt to help you understand the importance that distance latency has on the time it takes to complete an I/O operation. A more complete example is in the *Continuous Access EVA performance estimator user guide* available through this link: Continuous Access documentation.

Network distance/latency example calculations

1. **1.0 MB Link**

   Link Bandwidth: 1.0 MB/s

   Write size: 8 KB

   Available bandwidth divided by size of I/O equals maximum I/Os per second:

   $$\frac{\text{1.0 MB/s}}{\text{8 KB per I/O}} = \textbf{125 I/Os per second}$$

2. **50 Miles of Latency**

   Distance: 50 miles (80 km)

   Latency: 8 µs/mile (5 µs/km)

   Write size: 8 KB

   Latency for 1 I/O per mile: 4 trips * 8 µs/mile = 32 µs per mile

   Latency for 1 I/O at 50 miles: 50 miles * 32 µs/mile = 1.6 ms per I/O

   Reciprocal of total latency indicates maximum I/Os:

   $$\frac{1.0}{\text{1.6 ms per I/O}} = \textbf{625 I/Os per second}$$

   I/Os multiplied by size of I/O = bandwidth used:

   625 I/O per second * 8 KB per I/O = **5 MB/s**

3. **150 Miles of Latency**

   Distance: 150 miles (241 km)

   Latency: 8 µs/mile (5 µs/km)

   Write size: 8 KB

   Latency for 1 I/O per mile: 4 trips * 8 µs/mile = 32 µs per mile

   Latency for 1 I/O at 150 miles: 150 miles * 32 µs/mile = 4.8 ms per I/O

   Reciprocal of total latency indicates maximum I/Os:

   $$\frac{1.0}{\text{4.8 ms per I/O}} = \textbf{208 I/Os per second}$$

   I/Os multiplied by size of I/O = bandwidth used:

   208 I/O per second * 8 KB per I/O = **1.6 MB/s**

In summary, when an IP Network is used in a situation where the local and remote sites are located many miles apart, the speed of light through fiber may cause unacceptable delays in the completion of an I/O transaction. Increasing the amount of available bandwidth cannot solve this problem. Careful consideration must be given to these factors when matching your needs and wants to a particular application.

IP network best practices

Most IP networks do not manage bandwidth to each individual connection. As traffic increases due to other demands on the network, bandwidth can be robbed from the FCIP Intersite Link. The following techniques can be used to minimize this effect:

- Create virtual private networks (VPNs) with Quality of Service (QoS) through premise routers for the FCIP circuit.
- Create separate physical and dedicated networks.
- Guarantee the bandwidth, latency, and latency jitter using a third-party router/QoS vendor.

As mentioned, distance has a dramatic effect on the amount of work that can be done across a link. Therefore, site planning should include:

- Using the shortest possible distance between remote sites.
- Minimizing the amount data transferred over the FCIP link.
- Designing a plan to add additional storage I/O that will not impact normal data traffic.
- Consider additional controller pairs to effectively use available bandwidth.

Cisco MDS 4/8-Port IP Storage and MDS 14/2-Port Multiprotocol Services Modules

The Cisco MDS 4/8-Port IP Storage and MDS 14/2-Port Multiprotocol Services modules extend the reach of Fibre Channel SANs by using open-standard, IP-based technology. The switches connect separated SAN islands using Fibre Channel over IP (FCIP). They integrate seamlessly into the C-Series MDS 9000 Family, and support the full range of features available on other switching modules, including VSANs, security, and traffic management.

The MDS 4/8-Port IP Storage and MDS 14/2-Port Multiprotocol Services Modules can be used in the C-Series 9500 and 9200 families of switches and have 4, 8 and 2 Gigabit Ethernet ports respectively.

- FCIP-FCIP transports Fibre Channel frames transparently over an IP network between two C-Series MDS 9000 Family switches. Figure 61 depicts the FCIP scenarios in which the IPS module is used.
- Simplifies data protection and business continuance strategies by enabling backup, remote replication, and disaster recovery over WAN distances using open-standard FCIP tunneling.
- Improves utilization of WAN resources for backup and replication by tunneling up to 3 virtual interswitch links (ISLs) on a single Gigabit Ethernet port.
- Reduces SAN complexity by eliminating the need to deploy and manage a separate remote connectivity platform.
- Preserves C-Series MDS9000 Family enhanced capabilities including VSANs, advanced traffic management, and security across remote connections.

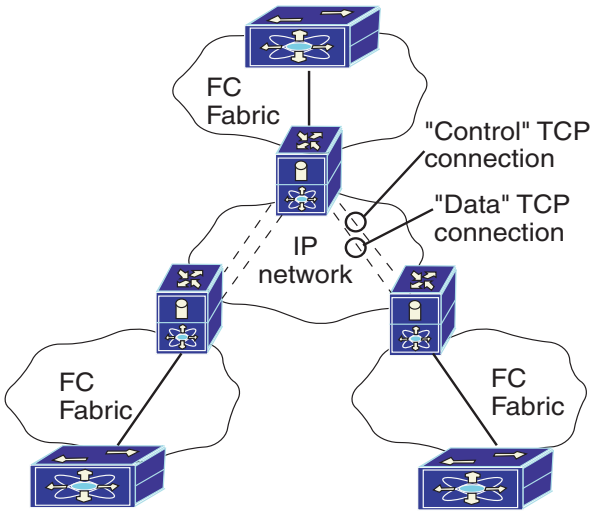

Figure 61: FCIP scenarios

Table 65: Supported SFPs

Optics	Media	Distance
1-Gbps—SX, LC SFP	50/125 micron multimode	500 m
1-Gbps—SX, LC SFP	62.5/125 micron multimode	200 m
1-Gbps—LX/LH, LC SFP	9/10 micron singlemode	10 km

Cisco MDS 4/8-Port IP Storage and MDS 14/2-Port Multiprotocol Services Module documentation

More information related to the use of the MDS IP 4/8-Port Services Module can be found at the following web site:

http://www.cisco.com/en/US/products/hw/ps4159/ps4358/products_data_sheet09186a00800c465b.html

Cisco MDS 4/8-Port IP Storage and MDS 14/2-Port Multiprotocol Services Module hardware and software support

This section lists the hardware and operating systems that are supported with the Cisco MDS 4/8-Port IP Storage Module.

IP network support

Network Protocols: TCP/IP IPv4, Ethernet 10Mbs/100Mbs/1000Mbs

Fibre Channel switch hardware support for iSCSI and FCIP with the IP Services Module

- C-Series MDS 9216/9216A/9216i Multilayer Fabric Switches
- C-Series MDS 9506 Multilayer Director Switch
- C-Series MDS 9509 Multilayer Director Switch

Storage array hardware support for FCIP

This section lists the storage array support when using the MDS 4/8-Port IP Storage Services Module for FCIP. See "Heterogeneous server rules" on page 127 for specific operating system support.

- Continuous Access Enterprise Virtual Array (EVA), VCS v3.0x

- Continuous Access XP48/1024
- Data Replication Manager for RA/MA8000, ESA/MSA12000, EMA16000, (HSG80)

HP StorageWorks MP Router

The B-Series HP StorageWorks MP Router offers FCIP SAN extension functionality and Fibre Channel routing capabilities. Individual ports on the MP Router can be configured as GigE ports for connectivity to an IP network for FCIP Tunneling services, or as Fibre Channel EX-Ports for connectivity to B-Series Fibre Channel fabrics for Fibre Channel Routing services. The two features can be combined to provide an FCIP configuration with very specific device connectivity without the need to fully merge the local and remote fabrics. This prevents unwanted access to all devices on both the local and remote fabrics.

HP StorageWorks MP Router documentation

The following MP Router product specific documentation is available on the external web at:

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

- HP StorageWorks Multi-protocol Router XPath OS Version 7.1.x Installation Guide
- HP StorageWorks XPath OS Version 7.1.x Advanced Web Tools User Guide
- HP StorageWorks XPath OS Version 7.1.x Command Reference Guide
- HP StorageWorks XPath OS Version 7.1.x Procedures User Guide
- HP StorageWorks XPath OS Version 7.1.x MIB Reference Guide
- HP StorageWorks XPath OS Version 7.1.x System Error Messages Reference Guide

HP StorageWorks MP Router - FCIP overview

Router FCIP enables a tunnel connection between Fibre Channel switches through an IP network. If the two MP Routers are configured for FCIP only, the switches and associated fabrics connected through the FCIP tunnel merge into a single Fibre Channel fabric. Using FCIP in conjunction with the Fibre Channel routing feature, device access can be limited to those devices requiring access such as local and remote storage arrays and associated servers used with HP's disaster tolerant Continuous Access storage products. LSANs can be defined that contain only the devices requiring access, avoiding a full merge of fabrics. See "Fibre Channel routing" on page 65 for information about Fibre Channel routing.

To deploy FCIP, two MP Routers are required, each acting as a local or remote router from the perspective of each respective fabric. Each MP Router is configured for FCIP and connected to the same IP network.

A Fibre Channel device in a local fabric needs no additional hardware or software to access other devices in the remote fabric via an MP Router deployed for FCIP, other than connectivity to an IP network.

With FCIP, peer systems transport Fibre Channel frames over an IP network. From the perspective of the SAN, the storage devices accessed through the peer systems appear to be part of one unified SAN.

Once configured, FCIP instances, or connections, on each system become active and establish their connectivity via the IP network. The storage devices in one SAN access the storage devices in the connected SAN using Fibre Channel frames, which are encapsulated in IP packets by the FCIP instance, and transmitted to the peer system. The peer FCIP instance

strips the IP packet data and passes only the Fibre Channel frames over the Fibre Channel interfaces to the storage devices. The peer systems are connected to each other through an IP network.

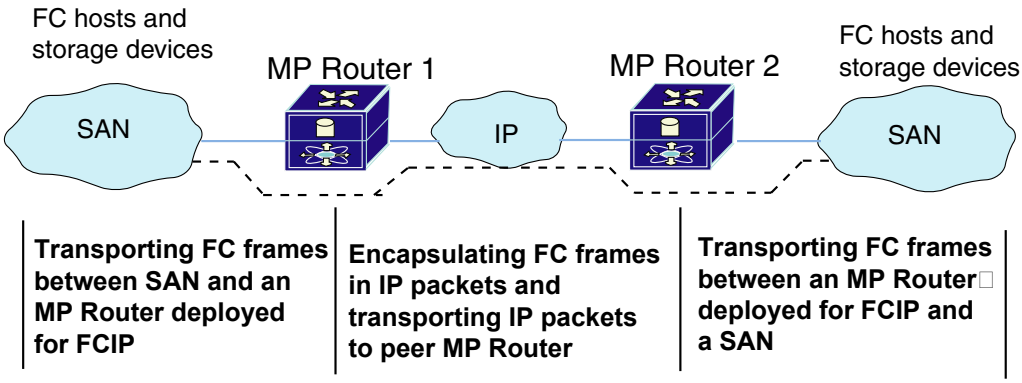

Figure 62: MP Routers connecting peer systems through an IP network

Note: Multiple FCIP links between any two MP Routers, either as multiple E_Ports or in combination with exchange-based trunking are not supported for EVA CA.

Note: All MP Router FCIP configurations require a minimum level of IP bandwidth for Continuous Access XP, Continuous Access EVA, and DRM replication. For specific IP bandwidth requirements, see the *CA/DRM SAN Extension reference guide* at:
http://h200006.www2.hp.com/bc/docs/support/SupportManual/c00244034/c00244034.pdf

In this example an FC host or FC device connects to one or more B-Series Fibre Channel switches within the SAN. The B-Series Fibre Channel switch or switches connect to the MP Router's Fibre Channel interface. Each MP Router connects to the IP network through one of its Gigabit Ethernet interfaces. Through the IP network each FCIP instance accesses its peer, thereby connecting the SANs.

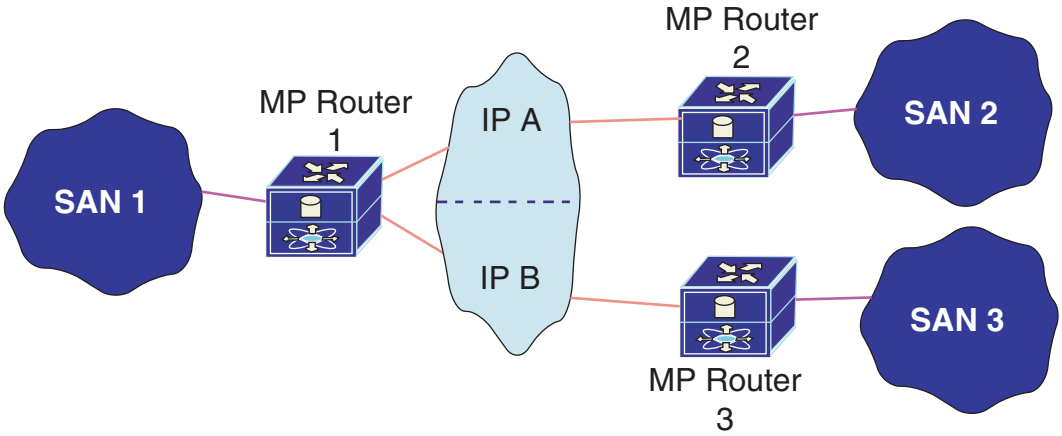

Figure 63: Sample configuration using two IP subnets

In this sample configuration, MP Router 1 must have one GigE connection into an IP Subnet for the IP A network and one GigE connection into a different IP subnet for the IP B network.

In this example configuration, an FC host or FC device connected to SAN 1 could have access to devices located in SANs 2 and/or 3.

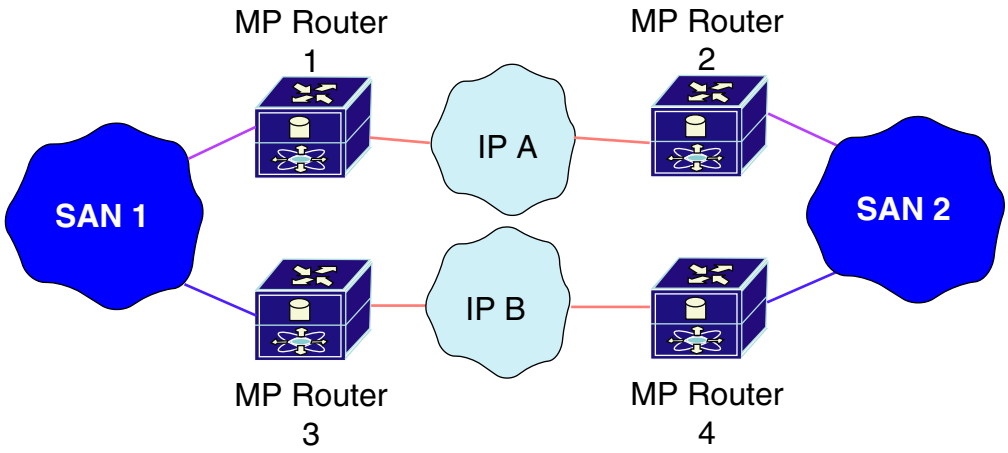

Figure 64: Fully redundant MP Routers with FCIP

This example shows a more reliable FCIP configuration, in which pairs of MP Routers provide full redundancy. In this configuration, loss of an MP Router or loss of connectivity through one of the IP networks can be tolerated with no loss of connectivity between the SANs.

Note: For multiple paths between SANs, multiple pairs of systems deployed for FCIP need to be connected to the Fibre Channel hosts or Fibre Channel devices. It is assumed that the multipath management is being done by an entity outside the MP Router (for example, by management applications on the Fibre Channel host or storage devices).

MP Router FCIP hardware and software support

This section lists the hardware, devices, and operating systems that are compatible with the MP Router for FCIP.

Note: For more information on support, contact an HP storage representative.

Storage array hardware support

This section lists the storage array support for FCIP when using the MP Router. Contact an HP storage representative for specific support information.

- Continuous Access for Enterprise Virtual Array
- Continuous Access for XP12000, XP128/1024, XP 48/512
- Data Replication Manager for RA/MA8000, ESA/MSA12000, EMA16000, HSG80

Fibre Channel switch hardware support

The MP Router is supported with the HP B-Series product line switches listed in “B-Series switches and fabric rules” on page 81.

Operating system support

The MP Router is supported for FCIP on the following operating systems, see the appropriate storage system type support documentation for specific version support:

- HP-UX
- OpenVMS
- Tru64
- IBM AIX
- Microsoft Windows
  - Windows 2000
  - Windows 2003
  - Windows NT 4.0
- NetWare
- Sun Solaris
- Red Hat Linux
- United Linux

HP StorageWorks SR2122-2 IP Storage Router

The HP StorageWorks SR2122-2 IP Storage Router offers FCIP SAN extension functionality and iSCSI to Fibre Channel Bridge capability within a single chassis.

The HP StorageWorks SR2122-2 IP Storage Router can be configured to run in Single-Mode (FCIP or iSCSI Routing only) or in Multi-mode (FCIP and iSCSI Routing concurrently).

IP SR2122 Storage Router documentation

Further SAN configuration documentation, including:

- *HP StorageWorks IP Storage Router 2122-2 Command Line Interface Reference Guide*
- *HP StorageWorks IP Storage Router 2122-2 User Guide*
- *HP StorageWorks IP Storage Router 2122-2 Getting Started Guide*

is available via the HP website at:

http://h18000.www1.hp.com/products/storageworks/san/documentation.html

HP StorageWorks SR2122-2 IP Storage Router - FCIP overview

Fibre Channel over IP (FCIP) enables SR2122-2 Storage Routers to provide connectivity between Fibre Channel hosts and Fibre Channel storage devices over an IP network. To deploy FCIP, two SR2122-2 Storage Routers are required. Each system is configured for FCIP and connected to a SAN. A Fibre Channel host or Fibre Channel device needs no additional hardware or software to access storage devices via an SR2122-2 Storage Router deployed for FCIP.

Note: See Table 44 on page 119 for a list of devices supported for FCIP heterogeneous SAN extension.

With FCIP, peer systems transport Fibre Channel frames over an IP network. From the perspective of the SANs, the storage devices accessed through the peer systems appear to be part of one unified SAN.

Once configured, FCIP instances, or connections, on each system become active and establish their connectivity via the IP network. The storage devices in one SAN access the storage devices in the connected SAN using Fibre Channel frames, which are encapsulated in IP packets by the FCIP instance, and transmitted to the peer system. The peer FCIP instance strips the IP packet data and passes only the Fibre Channel frames over the Fibre Channel interfaces to the storage devices.

The peer systems are connected to each other through an IP network.

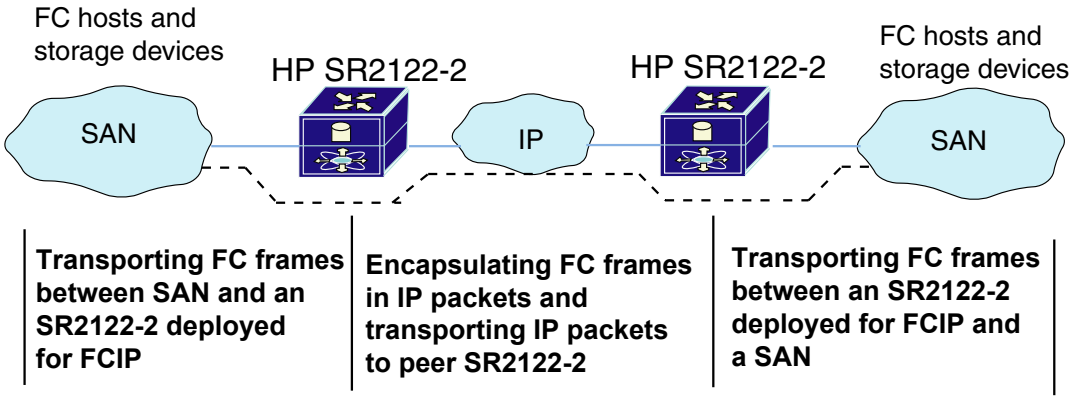

Note: SR2122-2 FCIP configurations require a minimum level of IP bandwidth. For specific IP bandwidth requirements, see the *CA/DRM SAN Extension reference guide* at: http://h200006.www2.hp.com/bc/docs/support/SupportManual/c00244034/c00244034.pdf

In this example, a Fibre Channel host or Fibre Channel device connects to one or more Fibre Channel interfaces of each peer SR2122-2 Storage Router deployed for FCIP. Each SR2122-2 connects to the IP network through one of its Gigabit Ethernet interfaces. Through the IP network, each FCIP instance accesses its peer, thereby connecting the SANs.

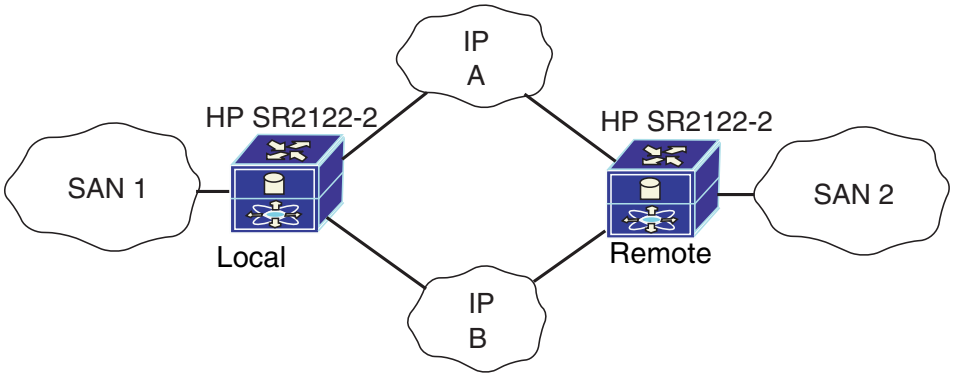

Note: In this sample configuration, only one side of each instance can be an SR2122-2 server. The opposite side must be an SR2122-2 client. (Use of the terms "server" and "client" in this context relates to the role of the SR2122-2 instance, not a physical server or client.)

In this example configuration, a Fibre Channel host or Fibre Channel device connects to one or more Fibre Channel interfaces of each peer SR2122-2 Storage Router deployed for FCIP, and each SR2122-2 connects to two separate IP networks through each of its Gigabit Ethernet interfaces. Through the IP network, each FCIP instance accesses the peer storage router deployed for FCIP, connecting the SANs. In this configuration, IP A and IP B are redundant paths, so that the loss of connectivity via either path does not cause a loss of connectivity between the SANs.

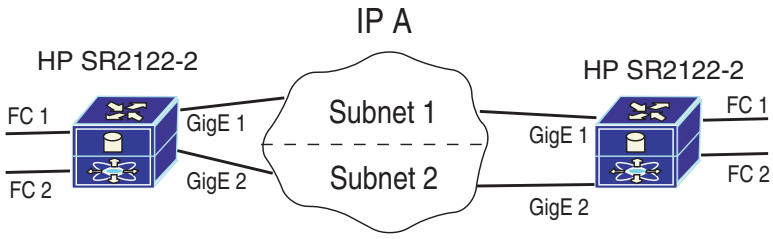

This next example uses two separate subnets in the same IP network for the two redundant IP pathes between the SR2122-2's.

Each SR2122-2 is capable of 2 FCIP instances running simultaneously. At most, there can only be one server per SR2122-2. To connect to another SR2122-2 a server/client relationship must exist.

For example:

- SR2122-2 *local* FC1/GE1 -> server connects to SR2122-2 *remote* FC1/GE1 -> client
- SR2122-2 *local* FC2/GE2 -> client connects to SR2122-2 *remote* FC2/GE1 -> server
- SR2122-2 *remote* FC1/GE1 -> client connects to SR2122-2 *local* FC1/GE1 -> server
- SR2122-2 *remote* FC2/GE2 -> server connects to SR2122-2 *local* FC2/GE1 -> client
- SR2122-2 *local* contains 1 server and 1 client, likewise for SR2122-2 *remote*. For each instance there is a server/client relationship.

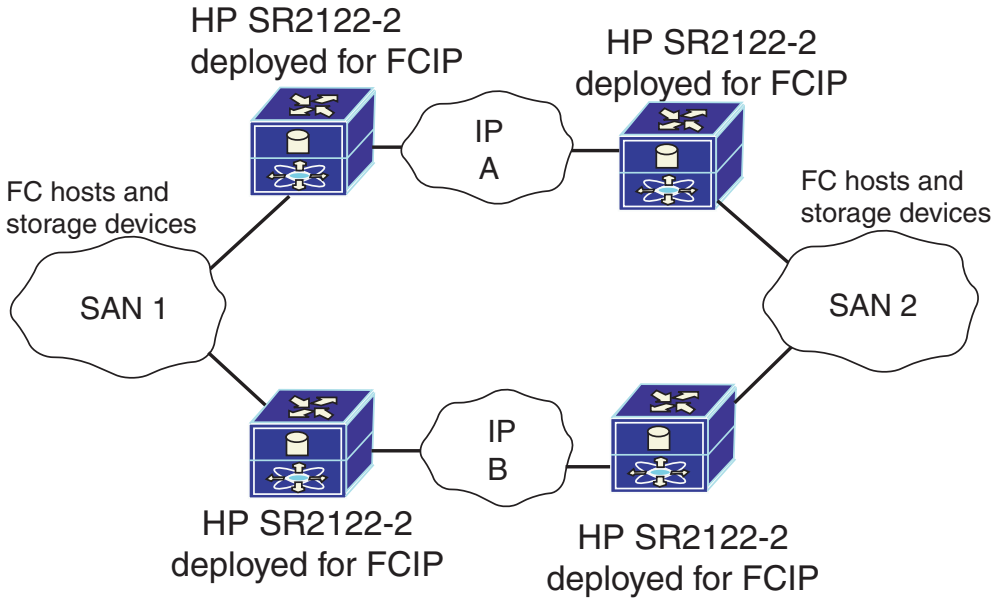

This example shows an even more reliable FCIP configuration, in which pairs of SR2122-2 Storage Routers provide full redundancy. In this configuration, loss of an SR2122-2 or loss of connectivity through one of the IP networks can be tolerated with no loss of connectivity between the SANs.

Note: For multiple paths between SANs, multiple pairs of systems deployed for FCIP need to be connected to the Fibre Channel hosts or Fibre Channel devices. It is assumed that the multipath management is being done by an entity outside the SR2122-2 (for example, by management applications on the Fibre Channel host or storage devices).

SR2122-2 FCIP hardware and software support

This section lists the hardware, devices, and operating systems that are compatible with the SR2122-2 for FCIP.

Note: For more information on support, contact an HP storage representative.

Storage array hardware support

This section lists the storage array support for FCIP when using the SR2122-2 IP storage router. Contact an HP storage representative for specific support information.

- Continuous Access for Enterprise Virtual Array
- Continuous Access for XP12000, XP128/1024
- Data Replication Manager for HSG80

Fibre Channel switch hardware support

The SR2122-2 is supported for FCIP with the HP B-Series, C-Series, and M-Series product line switches listed in "B-Series switches and fabric rules" on page 81, "C-Series switches and fabric rules" on page 97, and "M-Series switches and fabric rules", on page 105.

Note: The SR2122-2 does not connect with C-Series switches running firmware version 1.34a.

Operating system support

The SR2122-2 is supported for FCIP on the following operating systems:

CA EVA

- HP-UX 11i v1, 11i v2, 11.0
- OpenVMS 7.3-2, 7.3-1, 7.2-2
- Tru64 5.1a, 5.1b
- IBM AIX 5.2, 5.1, 4.3.3
- Microsoft Windows 2003 Advanced Server, Windows 2000 Server, Windows NT 4.0 SP6a
- NetWare 6.5, 6.0 SP3, 5.1 SP6
- Red Hat Linux Advanced Server 2.1, ES 3
- SuSE Linux SLES 7, 8, 9, United Linux 1.0

- Sun Solaris 2.6, 7, 8, 9

CA XP12000/1024/128

- HP-UX 11i v1
- OpenVMS 7.3-1, 7.2-2
- Tru64 5.1a
- IBM AIX 5.2
- Microsoft Windows 2003 Advanced Server, Windows 2000 Server
- Red Hat Linux Advanced Server 2.1
- SuSE Linux SLES 8
- Sun Solaris 8

DRM RA/MA8000, ESA/EMA12000, EMA16000 (HSG80)

- HP-UX 11i v1, 11.0
- OpenVMS 7.3-1, 7.3, 7.2-2
- Tru64 5.1a, 5.1b
- IBM AIX 5.1, 4.3.3
- Microsoft Windows 2000 Server, Windows NT 4.0 SP6a
- NetWare 6, 5.1
- Sun Solaris 2.6, 7, 8, 9

SR2122-2 FCIP configuration rules

SR2122-2 router rules

- IP network throughputs less than 10 Mb/sec are not supported.
- When the SR2122-2 is configured with HP data replication products, two separate long distance links must be implemented.
- If the router is to be configured for both FCIP and iSCSI, then the FCIP and iSCSI traffic is not supported through the same SR2122-2 GbE port.

Sample SR2122-2 configurations

This section provides a brief overview of three recommended host/storage configuration using the HP SR2122-2 IP Storage Router:

- FCIP Only
- FCIP with Local iSCSI Hosts
- FCIP with Remote iSCSI Hosts

Further SAN configuration documentation, including the

- HP StorageWorks IP Storage Router 2122-2 Command Line Interface Reference Guide
- HP StorageWorks IP Storage Router 2122-2 User Guide
- HP StorageWorks IP Storage Router 2122-2 Getting Started Guide

are available via the HP website at:

http://h18000.www1.hp.com/products/storageworks/san/documentation.html

SR2122-2 sample configuration - FCIP only

Two SAN islands may be joined into a single large, geographically dispersed SAN using the HP SR2122-2s as Fibre Channel to IP gateways to translate between Fibre Channel protocol and FCIP protocol.

FCIP protocol transmitted over a WAN network is used to extend the connection between the two SAN islands beyond the nominal 10 km maximum length for direct Fibre Channel.

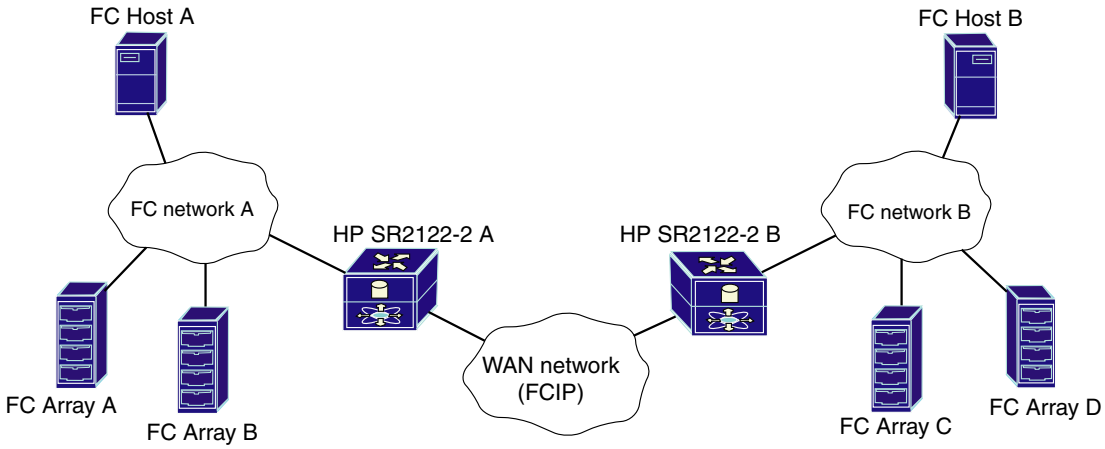

Figure 65: FCIP only

Disk LUNs at either site A or site B may be assigned either to local hosts or to remote hosts.

This basic configuration may also be used when Data Replication Manager or Continuous Access is employed to replicate disk data between the two sites. Because these data replication products use redundant Fibre Channel fabrics, two separate long distance links must be implemented. Although the two Fibre Channel fabrics could be routed through only two SR2122-2's, to avoid a single point of failure, a total of four SR2122-2 units should be included in this configuration.

As shown in Figure 65, a single Fibre Channel connection is required between the SR2122-2 and the Fibre Channel network at each site. The second Fibre Channel port on the SR2122-2 is not used. The iSCSI protocol is not used in this configuration.

SR2122-2 sample configuration - FCIP with local iSCSI hosts

One or more host servers may be connected to the extended SAN through a local IP network at site A using the iSCSI protocol. This connection uses the second Gigabit Ethernet port on the site A SR2122-2.

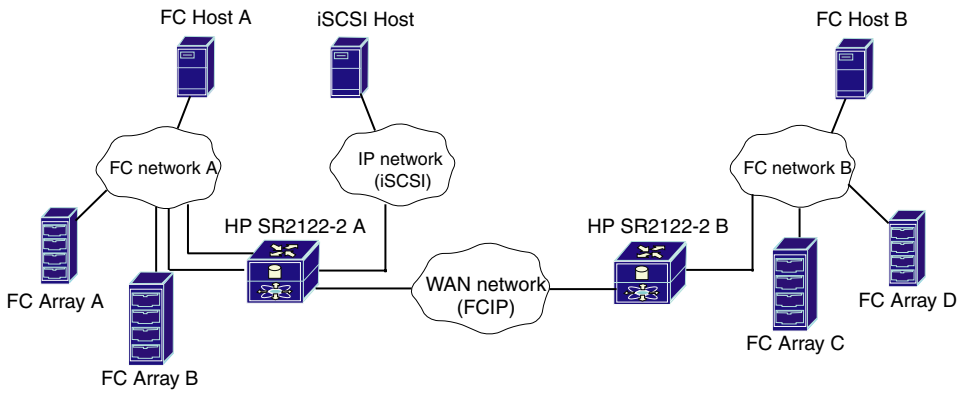

Figure 66: FCIP with local iSCSI hosts

For this configuration, two connections from the SR2122-2 to the Fibre Channel network at site A are required. One SR2122-2 Gigabit Ethernet (Fibre Channel) port is assigned to the FCIP connection, and the second is designated for the iSCSI connection.

Disk LUNs at site A may be assigned to the iSCSI hosts. The SR2122-2 translates the iSCSI I/O commands into Fibre Channel protocol commands. The iSCSI hosts at site A are also able to access the disk LUNs at site B. The iSCSI protocol I/O commands are converted to FCIP protocol in the site A SR2122-2 and transmitted to site B using FCIP. The iSCSI host applications must be able to tolerate the total latency incurred through the multiple protocol conversions plus the overall network delay to access disk LUNs at site B. A further expansion of this configuration would be to mirror the site A iSCSI configuration to include iSCSI hosts at site B. This would provide access to site B disk LUNs, as well as site A disk LUNs, through the SR2122-2 at site B.

SR2122-2 sample configuration - FCIP with remote iSCSI hosts

The FCIP configuration with local iSCSI hosts may be extended by locating the iSCSI hosts apart from either site A or site B. This configuration requires that the iSCSI IP network be connected to the large SAN through a WAN network as shown below.

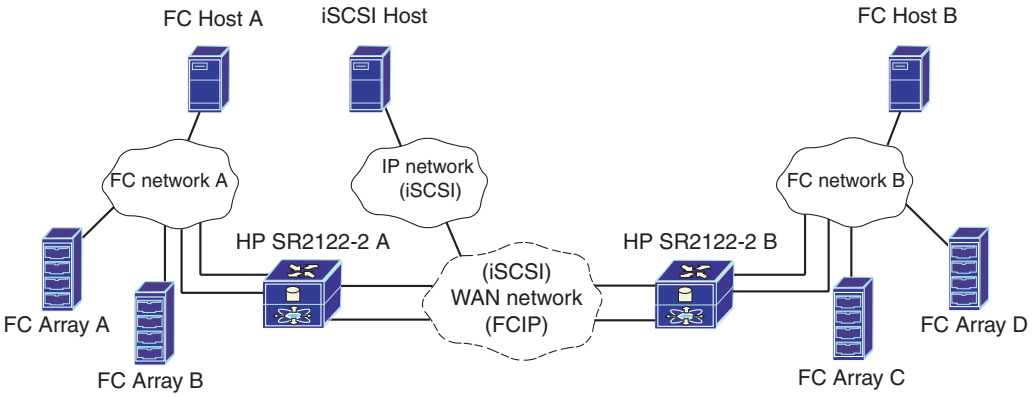

Figure 67: FCIP with remote iSCSI hosts

This configuration allows the iSCSI hosts to access disk LUNs at either site A or site B, providing maximum configuration flexibility.

Note: Within the WAN network the iSCSI protocol traffic is kept isolated from the FCIP protocol traffic and connects to the SR2122-2s through the second Gigabit Ethernet port on each gateway.

The SR2122-2 has the capability to rate-limit or “pace” the FCIP protocol traffic that it handles. This is accomplished using standard Fibre Channel flow control mechanisms that allows the user to limit the amount of FCIP traffic through the SR2122-2 so that it does not exceed the bandwidth allotted for this connection through the WAN network.

However, the iSCSI protocol traffic has no corresponding flow control mechanism. If the iSCSI protocol traffic and the FCIP protocol traffic are combined on a single network and if the combined traffic exceeds the available network bandwidth, the iSCSI protocol traffic can theoretically consume some or all of the bandwidth allotted to the FCIP connection. If that happens, both iSCSI and FCIP I/O commands are subject to failure due to dropped packets in the WAN network.

By isolating the iSCSI protocol from the FCIP protocol using separate network connections, it is possible to prevent a failure in the FCIP portion of the system due to oversubscription of the WAN connection.

iSCSI storage

12

Small Computer System Interface over IP (iSCSI) technology facilitates the creation of storage area networks (SANs) that include IP network technology. It is a TCP/IP-based protocol for establishing and managing connections between IP-based storage devices, hosts and clients.

This chapter discusses the fundamental concepts of iSCSI and then describes the current state of HP iSCSI support. It contains the following sections:

iSCSI overview

This section provides an overview of iSCSI protocol organized under these two topics:

- Features of iSCSI
- iSCSI concepts

Features of iSCSI

The iSCSI protocol enables universal access to storage devices and storage-area networks (SANs) over standard Ethernet-based TCP/IP networks.

These TCP/IP networks may be networks dedicated to storage or may be shared with traditional Ethernet applications.

The appeal of iSCSI for the storage market includes;

- Low cost
- Standardized network adapters and topology hardware
- Years of expertise developed in deploying IP networks

Note: Existing networks must be evaluated for their capacity to support iSCSI storage.

In any deployment of iSCSI, HP recommends using a dedicated Gigabit Ethernet network between iSCSI Initiators and Targets. This ensures adequate performance and provides data security. As an alternative, IPSec may be used to secure the connection in a shared network, but with lower performance.

Comparing iSCSI to Fibre Channel

There are many reasons to choose iSCSI over Fibre Channel when expanding a storage network. iSCSI can be more affordable, less complex, and can operate at acceptable speeds for almost all storage applications.

Other important advantages of iSCSI compared to Fibre Channel include:

- Lower deployment costs
- Familiar infrastructure - your entire IT staff may be comfortable with IP
- Fewer components - No HBA card required
- Support for servers and workstations of all sizes
- True “open architecture” design

iSCSI-enabled storage

HP provides the HP ProLiant Storage Server iSCSI Feature Pack, software that adds iSCSI target functionality to HP Storage Server (NAS) devices designed for small and medium business environments. These servers are typically used to provide storage for applications like Microsoft Exchange Server.

Bridging iSCSI to Fibre Channel

HP supports a wide array iSCSI bridges compatible with all HP supported Fibre Channel switches. This bridging technology enables storage administrators to route Fibre Channel storage to all servers in the environment, whether they reside on an IP fabric or FC fabric.

This iSCSI routing provides IP hosts with access to Fibre Channel storage devices as if the storage devices were directly attached to the hosts. Storage devices are not aware of each IP host.

Some key benefits of bridging include:

- Maximizing storage consolidation by presenting Fibre Channel storage to IP hosts
- Allow "stranded" servers without FC connectivity to take advantage of Enterprise Storage
- IP hosts are not restricted to FC distance limitations

HP's supported iSCSI-to-Fibre Channel Bridging technologies include:

- HP StorageWorks SR2122-2 IP Storage Router
- C-Series IPS Service Modules
- C-Series MDS 14/2-Port Multiprotocol Service Modules
- C-Series MDS 9216i Multilayer Fabric Switch
- B-Series MP Router

iSCSI concepts

iSCSI is a SCSI transport protocol that maps block-oriented (command descriptor block or CDB) storage data over TCP/IP networks. Key concepts related to iSCSI include:

- Sessions and Logins
- iSCSI names
- Discovery mechanisms
- Security
- Initiators and targets
- iSCSI HBA
- Bridging and routing
- iSCSI Boot

Sessions and Logins

For the initiator to communicate with the target, the initiator must establish a session with the target. An iSCSI session between initiator and target must be enabled through an iSCSI login process. It starts with a TCP/IP connection establishment. Then, the iSCSI Login Phase is used to negotiate any variable parameters between the two iSCSI entities and may invoke a security routine to authenticate allowable connectivity. If the iSCSI Login Phase is successful, the target will issue a login accept to the initiator; otherwise the login is rejected and the TCP connection is closed. The iSCSI Login negotiations are processed via text fields with proffered values for a series of parameters. Parameter keys and the corresponding range of valid values supported by each iSCSI entity are exchanged. When differences in values occur (e.g., the number of concurrent TCP connections that may be supported) between two devices, the lower of the two is used as the common session parameter. Text fields are also used to exchange names and aliases, as well as negotiable parameters such as

- the type of security protocol
- the maximum data payload size
- support for unsolicited data
- time-out values

As initiators establish iSCSI sessions with targets, session IDs are generated to uniquely identify individual conversations between specific iSCSI Nodes within the corresponding Network Entities. For example, an initiator logging on to a target would include its iSCSI name and an initiator session ID (ISID), the combination of which would be unique within its host Network Entity. A target, responding to the login request, would generate a unique target session ID (TSID), which likewise, in combination with its iSCSI name, gives that session a unique identity within its Network Entity. A single ISID/TSID session pair may have multiple TCP connections between them, per the results of login negotiation. Once login is completed, the iSCSI session leaves the Login Phase and enters the Full Feature Phase for normal SCSI transactions.

iSCSI names

All iSCSI initiators and target devices must be named with a unique identifier and an assigned address for network access. ISCSI initiators and targets are participants on an IP network, the clients and servers have a Network Entity identity which is equivalent to the IP addresses they are assigned. The Network Entity may contain one or more iSCSI Nodes. The iSCSI Node object identifies a SCSI device within a Network Entity that is accessible through the network. The Network Portal is the Combination of the Node's assigned IP address and the TCP Port number. Because a Network Entity may represent a gateway fronting multiple initiators or targets, the Network Entity object allows for multiple iSCSI Nodes. Each iSCSI Node is identified by a unique iSCSI name that may be up to 255-bytes long. While this may seem to be a fairly long identifier, the iSCSI protocol attempts to follow Internet conventions with human-readable names that can be parsed by a Domain Name Server (DNS) or other resource locator implementations. The 255-bytes name length ensures globally unique names that can be formatted for the convenience of the storage administrator.

The combination of IP address and TCP port generates a unique network address for an iSCSI device. The 255-byte iSCSI name provides a unique human-readable identity. The separation of iSCSI names and iSCSI addresses insures that a storage device will have a unique identity in the network regardless of its location in the network. While the IP address plus TCP port number will necessarily change if a device is moved onto a different network segment, the iSCSI name will travel with the device, allowing it to be rediscovered. A further benefit of iSCSI naming is that it is soft-assigned and remains independent of supporting hardware. This allows, for example, a device driver on a host platform to be assigned a single iSCSI name even if multiple storage NICs are used to attach the host to the network. Likewise, a target device could have multiple connections to the network for redundant pathing and yet be consistently identified as a single entity via the iSCSI name.

iSCSI naming rules creates a naming scheme that provides permanent and unique identity. There are 2 naming schemes: an iSCSI qualified name (iqn) or and enterprise unique identifier (eui).

Each iSCSI Node has an address consisting of the IP address, the TCP port number, and either the IQN or EUI name. IP addresses can be assigned by using common methods.

Discovery mechanisms

Static configuration

iSCSI initiators discover the targets using statically configured addresses that persist across reboot of initiators. This is recommended for simple SAN configurations. Administrators manually configure the Node name, IP address, TCP Port of the target in the initiator.

SendTargets

This is recommended for relatively small SAN. Administrator manually configures the address of a target portal, then initiator establishes a discovery session and uses SendTargets command to find out all the target node names available on that portal. Administrator sets up a range of target addresses for discovery.

SLP

Service Location Protocol (SLP) is a client-server protocol that allows services to be discovered by clients using the protocol defined interactions. There are three components necessary to make this work. An iSCSI initiator has a SLP User Agent (UA) that plays a role of a client, and iSCSI targets have an SLP Service Agent (SA) that plays a role of an SLP server. A Directory Agent (DA) may be developed to address multicast service requests from the server. The initiator can discover the targets by sending either a unicast discovery service request to DA or multicast discovery service request to SA or unicast discovery service to SA directly.

iSNS

The internet Storage Name Service (iSNS) discovery protocol provides both naming and resource discovery services for storage devices on the IP network. The iSNS is modeled on both IP and Fibre Channel. iSNS is also a client server protocol that defines the iSNS server, a directory server with some advanced capabilities, and iSCSI initiators and targets would have iSNS client capabilities. The target iSNS client registers the target with iSNS sever. The initiator iSNS client registers the initiator with iSNS server and queries the list of targets.

Security

Because iSCSI must accommodate untrusted IP environments, the iSCSI specification allows for multiple security methods to be implemented. Encryption solutions that reside below iSCSI such as IPSec require no special negotiation between iSCSI end devices and are transparent to the upper layers. For other authentication implementations such as Kerberos or Public/Private Key exchanges, the iSCSI Login Phase provides text fields for negotiating the type of security supported by both end devices. If the negotiation is successful, the PDUs exchanged between iSCSI devices will be formatted for appropriate security validation required by the agreed upon security routine. The iSNS server may also assist this process by, for example, serving as a repository for public keys.

Initiators and targets

The storage router manages access between devices defined as iSCSI targets or iSCSI initiators.

- **iSCSI target** (also logical target) - iSCSI target systems are end node devices that are typically storage systems, storage routers, or bridges. A storage system with iSCSI support is called native iSCSI storage. The iSCSI target listens on TCP port 3260.
- **iSCSI initiators** (IP hosts) - An iSCSI initiator is a system that initiates or starts the transfer of information to or from a iSCSI target. IP hosts see the physical storage devices grouped by iSCSI targets as if directly attached to the hosts.
- **iSCSI software initiators -** There are software iSCSI initiators available from many vendors. The software iSCSI protocol runs on the host and allows a standard Ethernet NIC to handle iSCSI traffic. Software iSCSI offers low cost with a performance penalty and CPU overhead.

iSCSI HBA

In some implementations, high performance iSCSI adapters perform both TCP and iSCSI on the interface card. While this adds cost to the adapter, it provides both high-speed iSCSI transport and minimal CPU overhead. The SCSI commands, status and data encapsulated by iSCSI are served up by the adapter card directly to the host operating system.

Bridging and routing

iSCSI storage routers or bridges are gateway devices that support interconnecting other storage protocols such as Fibre Channel or SCSI to IP networks and enable block level access across different networks.

The bridge or router maps another protocol device to an iSCSI target. In the case of Fibre Channel, the router maps Fibre Channel targets as pseudo iSCSI targets. Similarly the router maps iSCSI hosts as pseudo Fibre Channel initiators. When iSCSI hosts make requests for reads or writes of storage data, those requests are switched or routed through the IP network with the destination IP address of an IP storage switch. The router performs the conversion of the iSCSI request into its Fibre Channel equivalent and passes it to the Fibre Channel target storage device. The complementary conversion is performed as the Fibre Channel target responds to the iSCSI host. Because the IP storage router is mediating both sides of the storage conversation at wire speed, the conversion process is transparent to both hosts and targets.

iSCSI boot

iSCSI supports booting initiators from an iSCSI target. The iSCSI HBA may have built in initiator capabilities enabled in firmware.

B-Series MP Router

The B-Series MP Router offers Fiber Channel routing, FCIP SAN extension functionality, and iSCSI-to-Fibre Channel Bridge capability all in a single chassis.

Hardware support

This section lists the hardware devices supported with the B-Series MP Router for iSCSI.

Storage arrays

Devices supported for iSCSI access on the MP Router must be connected directly to an MP Router or a B-Series Fibre Channel switch on the MP Router backbone fabric. The supported arrays are:

- MSA1000
- MSA1500
- EVA3000/5000
- XP12000, XP48/512, XP128/1024
- VA7410/7110

Fibre Channel switch hardware support

The iSCSI capability with the MP Router supports all B-Series Fibre Channel switches in the backbone fabric.

The MP Router backbone fabric consists of one MP Router at a minimum and includes all MP Routers and B-Series switches connected to the MP-Router through E-ports. iSCSI initiators connected to the MP Router's iSCIS ports can access only storage connected directly to MP Routers or B-Series switches in the backbone fabric (Figure 29 on page 91).

Edge fabrics—Fibre Channel switches and their attached storage connected to the MP Router through EX-ports—are not accessible to iSCSI initiators connected to the MP Router's iSCSI ports (Figure 30 on page 91).

Software support

The B-Series MP Router iSCSI support is for single-path devices only. No multi-path software is supported.

Operating systems and network interface controllers

The following operating systems are supported with the MP Router:

- Microsoft Windows 2000 SP3
- Microsoft Windows 2003 Standard Server and Enterprise Editions

All HP supported NICs for Windows are supported.

Network teaming

HP Network teaming is supported with Windows 2000 and Windows 2003.

iSCSI initiator software

Microsoft iSCSI Initiator v1.06 is supported with the MP Router.

Configuration rules

Table 66 and Table 67 describes the configuration rules for the MP Router.

Table 66: MP Router iSCSI configuration rules

Rule	Maximum
Maximum # of iSCSI ports/MP Router	12
Maximum # of TCP sessions/MP Router	96
Number of TCP sessions/port	8

Table 67: MP Router iSCSI host rules

Rule	Maximum
Maximum # of targets/host	1

iNAS server—HP ProLiant Storage Server iSCSI Feature Pack

HP ProLiant Storage Server iSCSI Feature Pack adds iSCSI target capability to HP Storage Server products, enabling customers to combine file, print, and block storage services on a single platform.

Overview

HP ProLiant Storage Server iSCSI Feature Pack is powerful and easy-to-use software that adds iSCSI target functionality to HP Storage Server (NAS) devices designed for small and medium business environments. The combination creates an iSCSI Storage Server solution that is capable of hosting file, print, and application storage (block) services on a single platform.

Unlike environments with separate file, print, email, and database servers, or environments using proprietary technologies, this approach delivers single-platform manageability, easy scalability, and centralized backup. It provides investment protection by enabling low-cost storage consolidation using industry-standard hardware and software on existing Ethernet infrastructure.

HP ProLiant Application Storage Manager is included with HP iSCSI Feature Pack (T3669A only) and delivers radically simple storage management to customers using HP iSCSI Feature Pack to host Microsoft Exchange 2003 Storage Groups on HP NAS servers. By reducing process steps and knowledge requirements, it reduces time and training required to setup and monitors email stores, and ensures 'best practices' implementation through automation.

The HP iSCSI Storage Server solution is ideal for small or medium businesses that wish to take advantage of the simplification of storage consolidation and are application-focused on Microsoft® Exchange, Microsoft® SQL, or Oracle Database. It is easy-to-use and affordable, yet delivers powerful functionality usually reserved for higher-priced and more complex storage architectures.

iNAS documentation includes the following:

- HP ProLiant Storage Server iSCSI Feature Pack User Guide - Part Number: T3662-90901
- HP ProLiant Storage Server iSCSI Feature Pack Release Notes - Part Number: T3662-90902

These documents are available at:

http://h18006.www1.hp.com/storage/storageservers.html

HP ProLiant Storage Server iSCSI Feature Pack - iSCSI overview

HP StorageWorks iSCSI Feature Pack for HP StorageWorks NAS products provides virtualization, allocation of disk storage, and centralized management for iSCSI host applications.

HP iSCSI Feature Pack is comprised of a Windows-compatible iSCSI target initiator and an iSNS (Internet Storage Name Service) server. All software management has been integrated under a tabular “iSCSI” window in the Windows Server 2003 Web Administration screen.

The following figure illustrates how your NAS system provides file-level storage for your workstations, and how your iSCSI Feature Pack provides block-level storage for your application and file servers.

HP ProLiant Storage Server iSCSI Feature Pack iSCSI hardware and software support

Hardware support

- HP ProLiant ML110 Storage Server
- HP ProLiant DL100 Storage Server
- HP ProLiant ML350 G4 Storage Server
- HP ProLiant ML370 G4 Storage Server
- HP ProLiant DL380 G4 Storage Servers (Base, External SCSI, and External SATA models)
- HP StorageWorks NAS 500s
- HP StorageWorks NAS 1200s
- HP StorageWorks NAS 1500s
- HP StorageWorks NAS 2000s
- HP ProLiant DL380 G4 Storage Server (SAN Storage model)-Gateway Edition only
- HP StorageWorks NAS 4000s-Gateway Edition only

Application support

- Microsoft Exchange 2000 Server
- Microsoft Exchange 2003 Server
- Microsoft SQL 2000 Server
- Microsoft SQL 2003 Server
- Oracle Database 9i and 10g

Management software support

- Management controls integrate into HP Storage Server management GUI

iSCSI initiator support

- Microsoft iSCSI Initiator (32-bit version)

iSCSI initiator rules

- Up to 50 simultaneous initiators
- HP StorageWorks iSCSI Feature Pack has not been qualified with initiators running Microsoft® XP Home Edition or Microsoft XP Professional Edition.
- Many non-Microsoft initiators, such as Linux, HP-UX and Novell, may work but have not been qualified with this solution.
- Hardware initiators may work but are currently unsupported.

HP ProLiant Storage Server iSCSI options (license upgrades)

HP ProLiant Storage Server iSCSI snapshots

- Snapshot functionality upgrade license for HP ProLiant Storage Server iSCSI Feature Pack.

- For use with Microsoft iSCSI initiators only.
- Safeguards against accidental deletions, file corruptions, and virus attacks by creating point-in time images of data.
- Ensures quiescence of application hosts running Microsoft Exchange, SQL Server, or Oracle Database to ensure that delta changes are replicated with 100% transactional integrity.
- Allows users or applications to request fast, space-efficient delta snapshots via Microsoft's Volume Shadow Copy Service (VSS) interface, and performs policy-based, automatic delta snapshots of application hosts to significantly reduce potential data loss and ensure business continuity.
- User may install and license one of the following included initiator-based snapshot agents: Microsoft VSS, Microsoft Exchange, Microsoft SQL, and Oracle Database agents for a single Microsoft iSCSI initiator.

HP ProLiant Storage Server iSCSI clustering

- Clustering services upgrade license for HP ProLiant Storage Server iSCSI Feature Pack (Gateway Edition only).
- Activates two-node iSCSI target capability using Microsoft Cluster Services
- Increases application server availability by eliminating a single point of failure via a dual network connection to the IP network, as well as a dual I/O channel to each storage device.

HP ProLiant Storage Server iSCSI direct backup

- Direct backup functionality upgrade license for HP ProLiant Storage Server iSCSI Feature Pack.
- For use with Microsoft iSCSI initiators only.
- Facilitates centralized, zero impact backup.
- Allows Administrators to leverage their backup software of choice to centralize the backup and recovery of application data stored using HP iSCSI Feature Pack directly from and to the HP Storage Server
- User may install and license one of the following included initiator-based snapshot agents: Microsoft VSS, Microsoft Exchange, Microsoft SQL, and Oracle Database agents for a single Microsoft iSCSI initiator.

Microsoft® Exchange Server 2003 recommended design principles

The HP ProLiant Storage Server iSCSI Feature Pack is a new solution that enables placement of Exchange Server 2003 files on an HP ProLiant Storage Server running Microsoft Windows Storage Server 2003.

The HP ProLiant Storage Server iSCSI Feature Pack is installed on the Windows Storage Server computer (NAS device) to provide iSCSI target functionality. An iSCSI initiator is installed on the Exchange Server 2003 computers to add iSCSI functionality to each.

This section outlines the HP recommended design principles for almost all deployments.

Network design

Part of the appeal that iSCSI brings to the storage market is low cost, standardized network adapters and topology hardware in addition to years of expertise developed in deploying these networks.

However, it should be stressed that existing networks be evaluated for suitability regarding their capacity to support iSCSI storage.

In any deployment, HP recommends using a dedicated Gigabit Ethernet network between the Exchange server and the Windows Storage Server NAS. This ensures adequate performance as well as helps to provide data security against network sniffing of Exchange data. An alternative would be to use IPSec to secure the connection if it is not possible to use private, secured network—but there will be a performance impact.

If you plan to locate your Exchange server any distance away from the Windows Storage Server NAS, check with your networking hardware vendors on the specifics regarding the maximum supported distance. The maximum distance will vary according to the cable type and specifications.

Hardware selection

Check the Windows Server Catalog at:

http://www.microsoft.com/windows/catalog/server/default.aspx?subID=22&xslt=about&pgn=moreinfo

to to ensure than any iSCSI hardware component that you select is qualified under the Designed for Windows Logo program. If an iSCSI hardware device has passed the Designed for Windows Logo program, the hardware is also supported by Exchange Server 2003 and by Exchange 2000 Server with no additional qualification.

Exchange storage design

The most important criteria for design and selection of Exchange storage are:

- Isolation of Exchange transaction logs from databases
- Selection of best RAID protection for performance and fault tolerance
- Hardware RAID controller with sufficient write-back caching for performance

Exchange transaction logs and databases must be stored on separate disk volumes to provide both data protection and efficiency (separation of sequential writes and random read/write access, respectively).

The HP ProLiant Storage Server iSCSI Feature Pack and NAS can be configured to place the Exchange transaction logs on the Exchange server and the databases on the Windows Storage Server NAS.

Little performance difference was measured between placing the transaction logs local to the Exchange server compared with placing them on an iSCSI disk. However, you must consider how to recover the transaction logs if the Exchange server fails and is replaced with new hardware; in that case, the log drives would need to be moved to the new server.

Transaction logs should be placed on a RAID 1 (mirror pair) array (volume), or for additional disk space, four or more disks in a RAID 1+0 (striped mirror).

Performance on the transaction log volume is enhanced by decreasing the response time, which is accomplished by write-back caching. On HP Smart Array controllers with battery-backed write cache, such as the Smart Array 5i Plus and later, the write cache percentage should be set at 100%. (Dedicated read cache memory is built into the controller.) This setting will also benefit performance on the database arrays (volumes).

For database arrays (volumes), the choice of RAID protection on the disk arrays is often a trade-off between maximum storage and performance.

While RAID 5 can provide data protection, it does so at the cost of performance. RAID 1+0 was shown in the testing to provide the absolute best performance given the same number of disk spindles.

For example, six spindles in a RAID 5 array cannot provide sufficient I/O rate to support the same number of Exchange mailboxes as the same number of disks in a RAID 1+0 array. RAID 1+0 is preferable for the database volume—even when using 36-GB drives in an array, if the number of disks required to support the I/O is used, this provides ample storage for 100-MB mailboxes.

Exercise caution when sizing if you are using the newest disk drives (for example, 146 GB or larger), as a few spindles can support the required database storage capacity but will not be able to support the required I/O performance.

In summary, it is most important to place the Exchange logs and database files on a RAID 1+0 array on a hardware RAID controller with sufficient write-back caching. Even if the logs need to be accessed over the Gigabit network (on the NAS), there should be adequate network performance, and the RAID controller performance is paramount.

Exchange sizing for supported load

One of the questions that must be answered in storage sizing, is how much performance does the (average) e-mail user require? This average load is then multiplied by the total number of users to determine what size of system is needed. Or, conversely, the capabilities of the system are examined and the maximum number of users that can be supported is determined.

If a production Exchange environment is in place, the 'perfmon' object for Disk Transfers per second on the database disk can be divided by the number of Active Connections to measure the current I/O per user. However, sizing for averages can lead to poor performance during peak periods, so additional overhead is needed for these peak, stressful periods.

The maximum recommended user load may actually be less than this number depending on the user profile (heavier usage) and additional services running on the Exchange server (connectors, anti-virus scanning, content indexing, etc.).

In addition to the level of activity by the e-mail users, another sizing consideration is the size of the mailboxes. The medium profile creates on average a 60-MB mailbox, and the heavy profile creates on average a 100-MB mailbox. Larger mailboxes do affect not only your storage sizing but also performance criteria, as Exchange performance is affected by managing the larger mailboxes.

Sample Microsoft® Exchange Server 2003 configurations

Low capacity: HP ProLiant DL100 Storage Server (up to 500 mailboxes)

The maximum recommended user load may actually be less than this number depending on the user profile (heavier usage) and additional services running on the Exchange server (connectors, anti-virus scanning, content indexing, etc.).

Also, another HP ProLiant server could easily support the 500 Exchange users. The main sizing concerns are sufficient RAM (1 GB) and enough disks for the Exchange database volume— especially if using RAID 5 (or Advanced Data Guarding [ADG]). As discussed earlier, RAID 1+0 is preferable for the database volume. Note also that the transaction logs can be placed in the Exchange server, which takes advantage of the additional storage capacity of the HP ProLiant DL380 server, but the Smart Array 5i RAID controller option kit is highly recommended as it allows the cache to be set at 100% write cache, which is important for performance.

Medium to high capacity: HP ProLiant DL380 G4 Storage Server (over 5000 mailboxes)

The HP ProLiant DL380 server was chosen as the Exchange Server because it can be configured with up to six internal drives and fault-tolerant options for redundant fans and power supplies. However, other HP ProLiant servers could easily support the 900 to 2,000 Exchange users.

The main sizing concerns are sufficient RAM (1 GB) and enough disks for the Exchange database\ volume—especially if using RAID 5 (or Advanced Data Guarding [ADG]). As discussed earlier, RAID 1+0 is preferable for the database volume. Note also that the transaction logs can be placed in the Exchange server, which takes advantage of the additional storage capacity of the HP ProLiant DL380 server, but the Smart Array 5i RAID controller option kit is highly recommended as it allows the cache to be set at 100% write-back, which is important for performance,

To provide the additional disks, a Smart Array 6402 and a Smart Array 6404 controller were placed in the DL380 G4 Storage Server. Four rack-mountable HP StorageWorks Modular Smart Array 30

(MSA30) disk enclosures (featuring redundant power supplies and up to 14 drives each) were added, for a total of 56 disks. An array of 10 disks on each enclosure was created using RAID 1+0 and presented to a host for the database volume.

The best practice recommendation is to place the transaction log disks on the RAID controller with the most write cache (protected by battery backup). This can be either in the Exchange server or on the iSCSI storage array (especially if a Smart Array has been added for capacity expansion).

HP StorageWorks SR2122-2 IP Storage Router

The HP StorageWorks SR2122-2 IP Storage Router offers FCIP SAN extension functionality and iSCSI-to-Fibre Channel bridge capability in a single chassis.

You can configure the SR2122-2 to run in one of two modes:

- Single-mode (FCIP or iSCSI routing only)
- Multi-mode (one port pair as FCIP and the other port pair as iSCSI).

For FCIP information, see "SAN extension" on page 209.

For more SR2122-2 configuration information, see the following:

- *HP StorageWorks FCIP/iSCSI Storage Router 2122-2 Command Line Interface Reference Guide*
- *HP StorageWorks IP Storage Router 2122-2 User Guide*
- *HP StorageWorks SR2122-2 IP Storage Router Getting Started Guide*

These documents are available at:

http://www.hp.com/support/iprouter

iSCSI-supported hardware

The SR2122-2 is supported in iSCSI configurations with specific hardware components such as:

- Storage arrays, page 249
- Fibre Channel switches, page 249

Note: For support information, see the SR2122-2 Product Support Matrix. This document is at: http://h18000.www1.hp.com/products/storageworks/san/documentation.html

Storage arrays

The SR2122-2 supports the following storage arrays for iSCSI:

- MSA1000
- RA/MA8000
- ESA/EMA12000
- EMA16000
- Enterprise Virtual Array
- VA7100
- VA7400/7410
- XP128/1024

For iSCSI, not all storage arrays or array options like Business Copy and Continuous Access are supported on all supported operating systems. Contact an HP storage representative for specific support information.

Fibre Channel switches

The SR2122-2 is supported with the HP B-Series, C-Series and M-Series switches. For information about these switches, see:

- B-Series switches and fabric rules on page 81
- C-Series switches and fabric rules on page 97
- M-Series switches and fabric rules on page 105

Software supported with iSCSI

The SR2122-2 is supported in iSCSI configurations built with software components such as:

- Operating systems and network interface controllers on page 250
- Network teaming on page 250
- SR2122-2 management software on page 250
- iSCSI initiator software on page 250

Operating systems and network interface controllers

- HP-UX 11i v1, 11i v2 — All NICs are supported by HP-UX.
- Microsoft Windows 2000 SP2 with either Microsoft hotfix Q302895 or Q248720 and Microsoft hotfix Q318271, SP3, Microsoft Windows 2003 Enterprise Edition, Standard Server — All NICs are supported by HP for Windows.
- Microsoft Cluster Services (MSCS) support (HP iSCSI initiator only).
- Red Hat Advanced Server 2.1 — All NICs are supported by HP for Linux.
- Secure Path (Windows 2000, Windows 2003).

Network teaming

- Compaq Network Teaming (Windows 2000, Windows 2003).

SR2122-2 management software

The following HP management software is supported:

- Compaq Insight Manager (CIM) 7
- HP OpenView Storage Area Manager (SAM)

iSCSI initiator software

- HP iSCSI initiator
- Microsoft iSCSI initiator (MSCS and Secure Path are not supported)

SR2122-2 iSCSI configuration rules

SR2122-2 router rules

Table 68: SR2122-2 router rules

Router rule	Maximum
scsirouter instances per SR2122-2 router (and per SR2122-2 router cluster)	12
iSCSI host connections per SR2122-2 SCSI router instance	32
Active logical units (LUNs) per SR2122-2 router	120
Active targets per SR2212-2 router	120

For the SR2122-2:

- The 2nd fibre Channel port (FC2) is not supported as a redundant iSCSI SAN port for FC1. FC2 can, however, be configured as an FCIP port
- Direct connect of the Fibre Channel ports to any HP storage array is not supported.
- The Management port must be in a different subnet than the SCSI Router Instances.
- The Fibre Channel ports appear as host bus adapters to the SAN and all storage arrays.

iSCSI host

Table 69: iSCSI host rules

iSCSI host rule	Operating system	Maximum
Maximum targets accessed per iSCSI host	Windows 2000, Windows 2003, Linux, HP-UX	120
Maximum active LUNs per target	Windows 2000, Windows 2003, Linux, HP-UX	120

Operating system

- Linux Clustering is not supported.
- HP Secure Path for MSA1000, RA/MA8000, EMA/ESA12000 and for Enterprise Virtual Array for Linux are not supported.
- HP Auto Path for VA/XP for Windows 2000 and Linux are not supported.
- HP Secure Manager on XP and VA is not supported.
- Windows MSCS and Windows Secure Path not supported with the Microsoft iSCSI Initiator

Storage array rules

- The HSG80 is supported in SCSI-3 Transparent Failover and Multiple-Bus Failover modes.
- The SR2122-2 accesses only one controller port when used with RA8000/MA8000 or Enterprise Virtual Array storage without HP Secure Path. Controller failover is disabled.
- The MSA1000 is supported with the SR2122-2 accessing only one MSA controller port. Controller failover is disabled.

Fibre Channel switch/fabric rules

- The SR2122-2 is supported on the HP B-series, C-series and M-series switches.

Note: The SR2122-2 is not supported for iSCSI on HP-UX with M-series switches.

- Zone the SR2122-2 only with the storage devices that it will access. Zoning the SR2122-2 with other storage is not supported.

Management software rules

- HP OpenView Storage Area Manager (SAM) provides property support only. It will identify the device, and can launch a device web server interface or telnet. A device specific plug-in for the SR2122 is available at the SAM Website.
- CIM 7 supports the SR2122-2's SNMP management capabilities.
- The SR2122-2 does not provide storage array management functions. Use the recommended application/element manager to configure the storage array.

Example configurations

The SR2122-2's SCSI routing provides IP hosts with access to Fibre Channel storage devices as if the storage devices were directly attached to the hosts, with access to devices being managed primarily in the storage router.

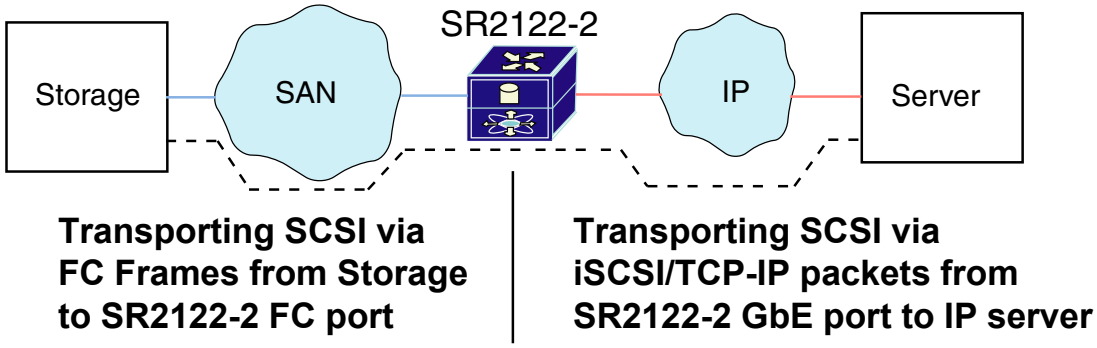

Fibre Channel storage connects to the Fibre Channel interface of the SR2122-2 Storage Router deployed for iSCSI. The SR2122-2 connects to the IP network through one of its Gigabit Ethernet interfaces. The server accesses the storage served from the router through the IP network.

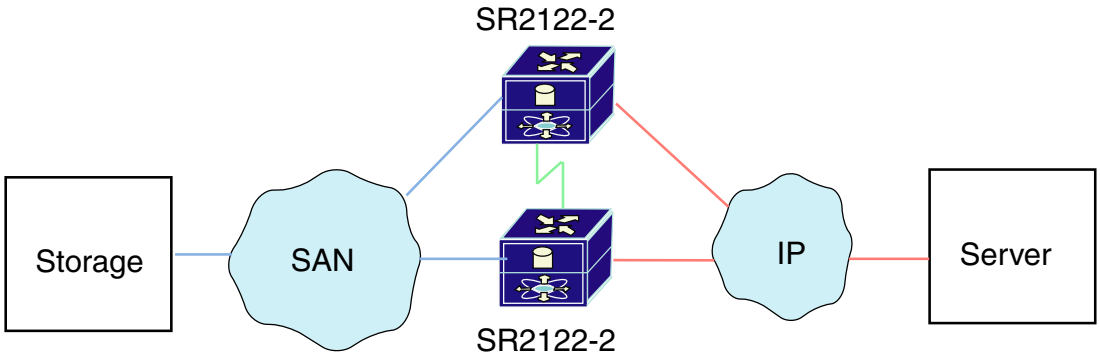

In this example configuration, two storage routers are connected (clustered) together. The storage routers back each other up in case of failure.

Note: A storage router can participate in an SR2122-2 cluster only if it is dedicated to iSCSI routing.

Clustered storage routers continually exchange High Availability (HA) information through a separate network connected to each storage router HA interface. The HA information includes shared configuration data and flags to indicate a failure in the cluster.

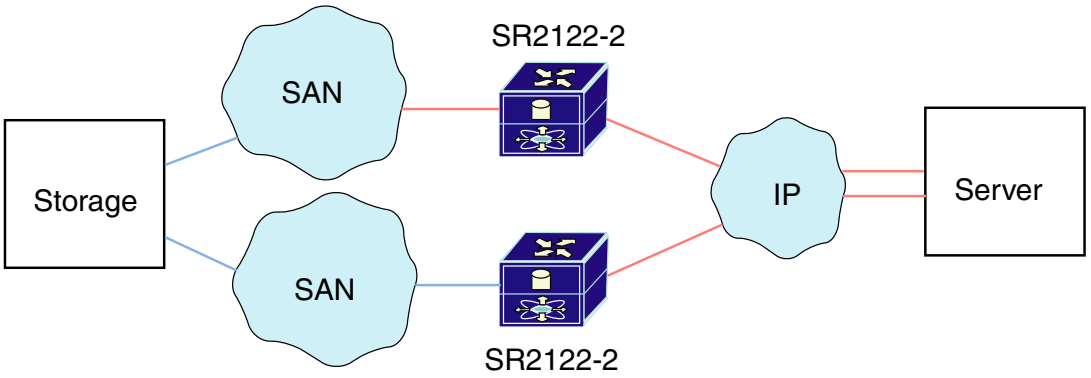

This example shows a more reliable iSCSI configuration in which pairs of SR2122-2 Storage Routers and SANs provide full redundancy. The configuration maintains connectivity between the storage and server if:

- An SR2122-2 fails
- One of the SANs fails

Note: For multiple paths between SANs, the multipath management must be done by an entity outside the SR2122-2 (for example, by management applications on the Fibre Channel host such as HP Secure Path).

More detailed configurations include:

- Example of Multiple OS Systems in a Non-Redundant Path Configuration
- Windows 2000 Servers with NIC Teaming: 2 Node SR2122-2 Cluster
- Secure Path Configuration
- Maximum SR2122-2 Cluster Configuration Using HA Ports

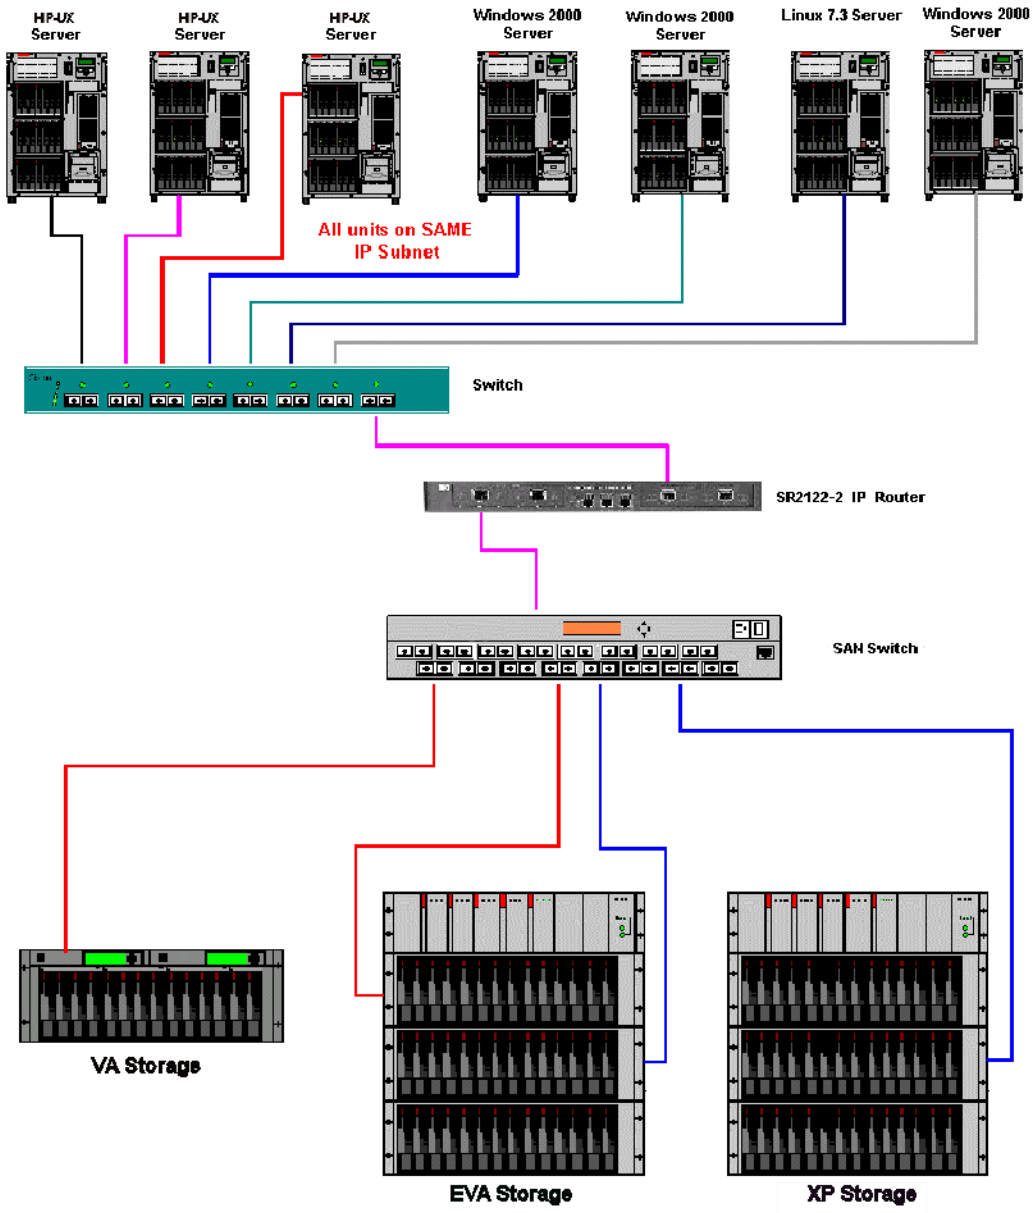

Figure 68: Example of multiple OS systems in a non-redundant path configuration

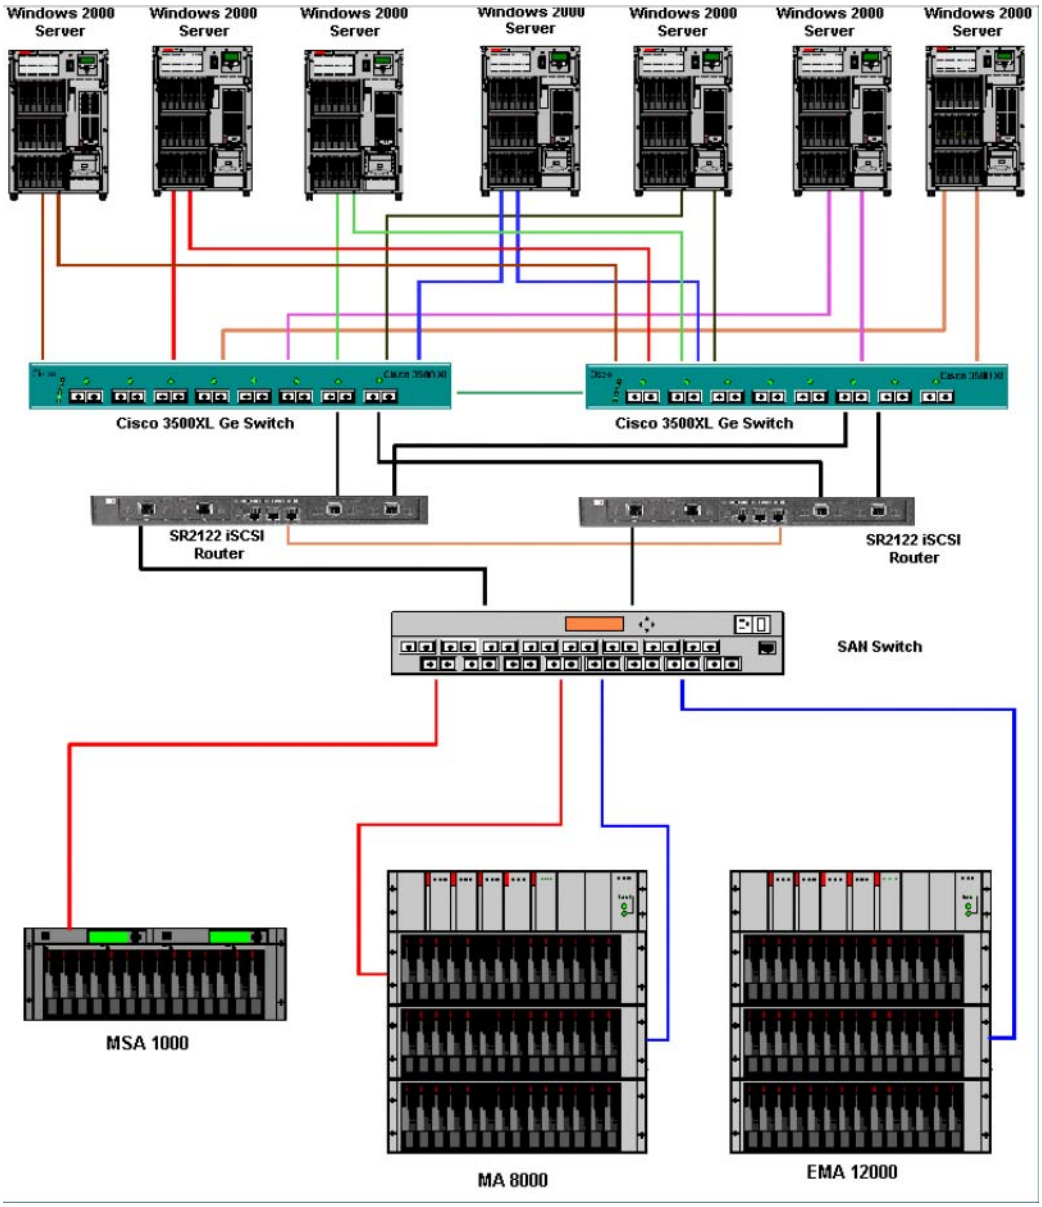

Figure 69: Windows 2000 servers with NIC teaming: 2 node SR2122-2 cluster

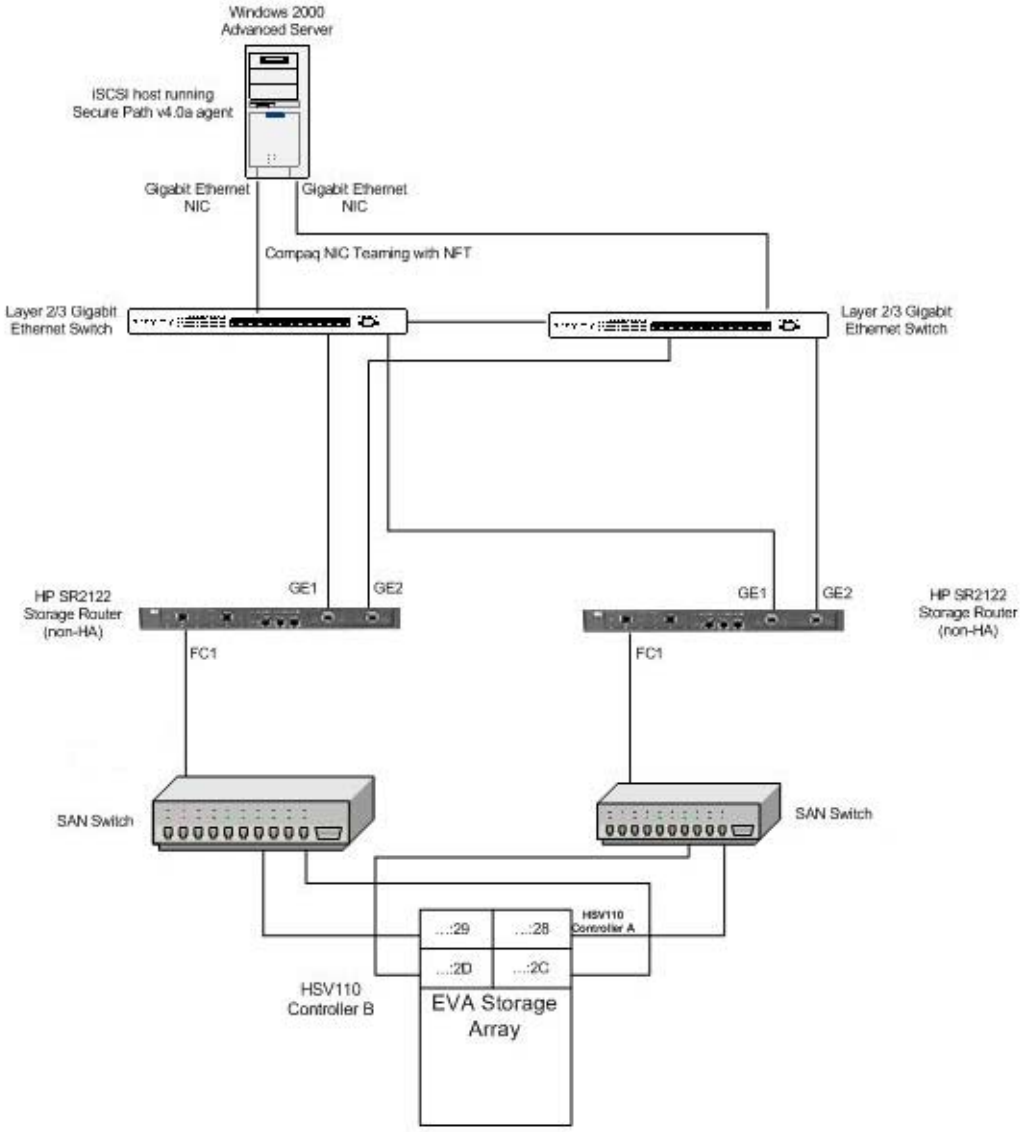

Figure 70: Secure Path configuration

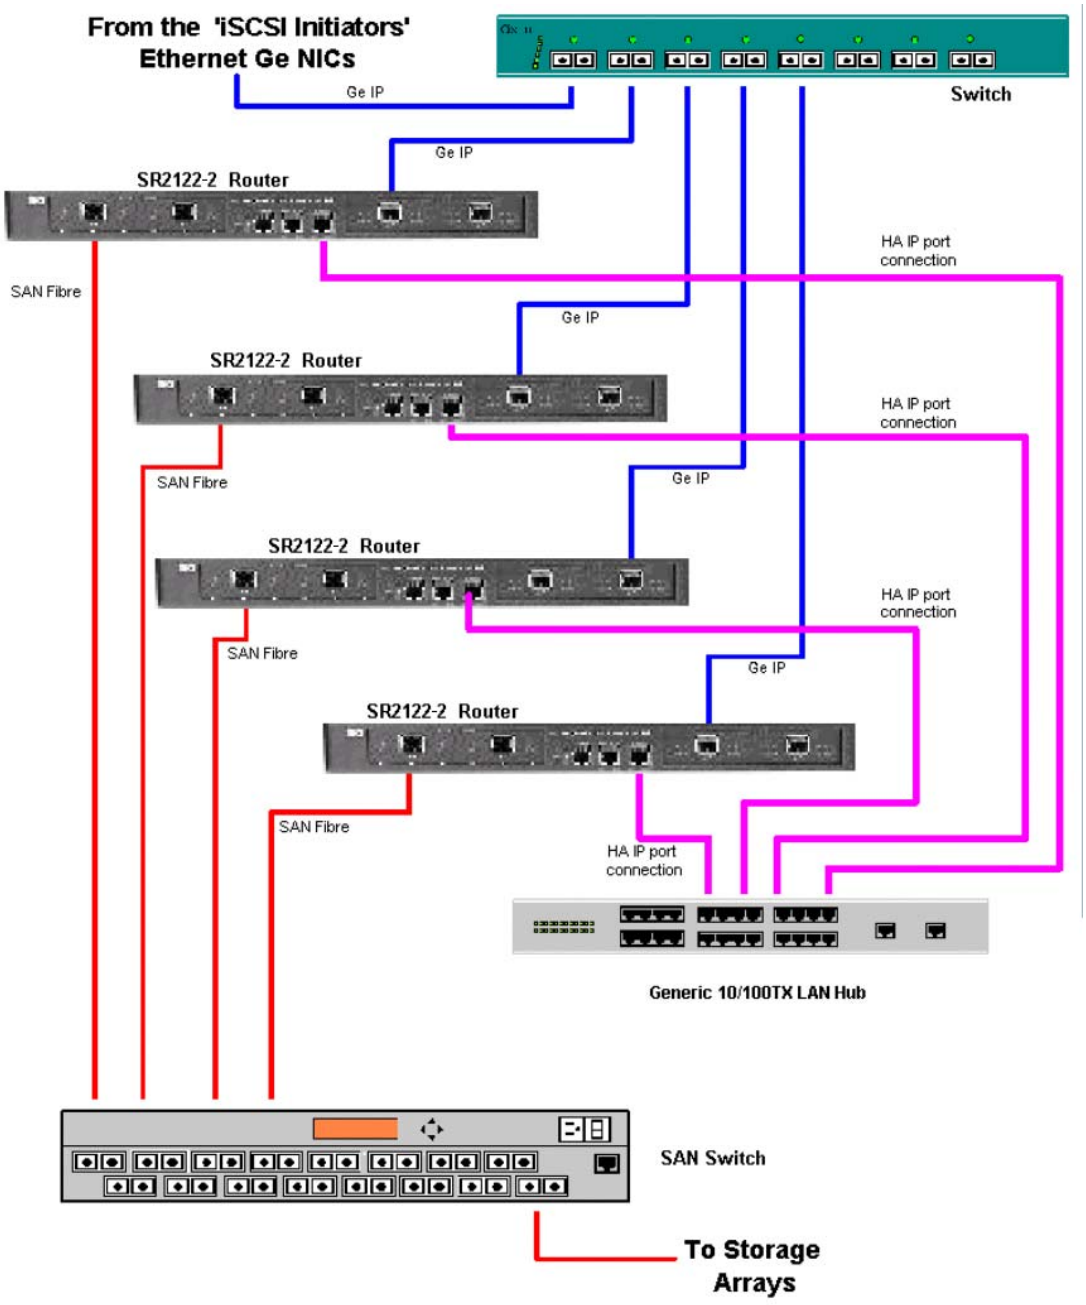

Figure 71: Maximum SR2122-2 cluster configuration using HA ports

C-series switches and modules

The C-Series IP storage (IPS) services extend the reach of Fibre Channel SANs by using open-standard, IP-based technology. The switch allows IP hosts to access Fibre Channel storage using the iSCSI protocol.

The IPS module provides transparent SCSI routing by default. IP hosts using the iSCSI protocol can transparently access targets on the Fibre Channel network.

Available C-Series documentation includes:

- *Cisco MDS 9000 Family Command Reference, Release 2.x*—Part Number: OL-6970-01
- *Cisco MDS 9000 Family Fabric Manager Configuration Guide, Release 2.x*—Part Number: OL-6965-01
- *Cisco MDS 9000 Family Software Configuration Guide, Release 2.x*—Part Number: OL-6973-01
- *HP StorageWorks C-Series IPS, 14/2 module, and MDS 9216i switch getting started guide* Part Number: AA-RW7PA-TE
- HP StorageWorks C-Series IPS, 14/2 module, and MDS 9216i switch support tables Part Number: AA-RW7QA-TE

Additional C-Series documentation is available via the HP website:

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

IPS Service module overview

The IPS Service module provides IP hosts access to Fibre Channel storage devices. The IP host sends SCSI commands encapsulated in iSCSI protocol data units (PDUs) to a C-Series IPS port over a TCP/IP connection. At this point, the commands are routed from an IP network into a Fibre Channel network and forwarded to the intended target.

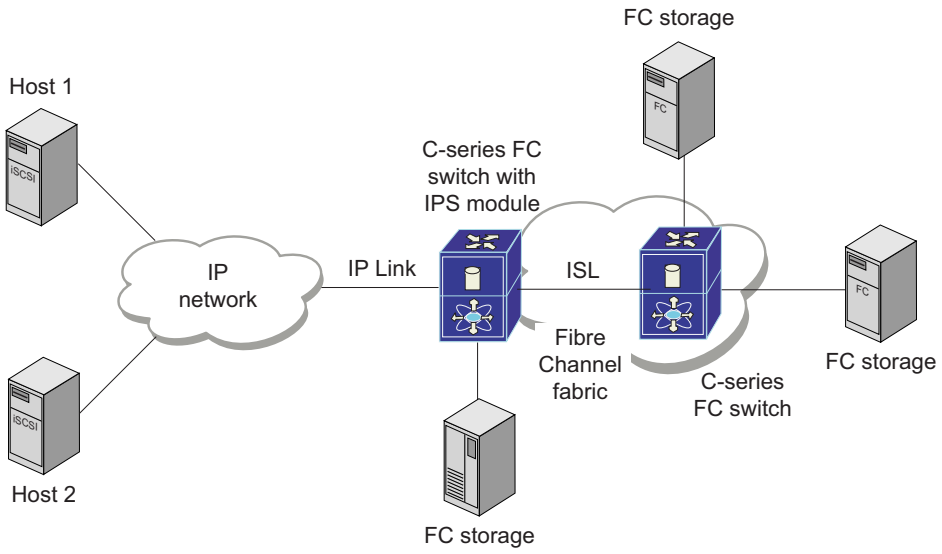

In conjunction with presenting Fibre Channel targets to iSCSI hosts, the IPS module presents each iSCSI host as a Fibre Channel host (in transparent mode), that is, a host bus adapter (HBA) to the Fibre Channel storage device. The storage device responds to each IP host as if it were a Fibre Channel host connected to the Fibre Channel network

Hardware support

This section lists the hardware and devices that are compatible with the C-Series IPS Service Module

Storage arrays

This section lists the storage array support for iSCSI when using the C-Series IPS Service modules. Not all storage arrays or array options like Business Copy and Continuous Access are supported on all supported iSCSI operating systems. Contact an HP storage representative for specific support information. Direct connection of storage arrays to the MDS IP 4/8-Port Storage Services Module and the MDS 14/2 Multiprotocol Services Module is not supported.

- MSA1000, MSA1500
- RA/MA8000
- ESA/EMA12000
- EMA16000
- Enterprise Virtual Array (EVA3000, EVA5000)
- VA7100/7110
- VA7400/7410
- XP48/128/1024/12000

Fibre Channel switch hardware support

The C-Series IPS Module is supported with the HP C-Series product line switches listed in Chapter 5 of this guide.

The following C-Series switches have iSCSI functionality:

Table 70: C-Series switches supporting iSCSI

Switch	Maximum number of Fibre Channel ports	Supported IP modules
MDS 9216	48	IPS-4, IPS-8, or 14/2
MDS 9216A	48	IPS-4, IPS-8, or 14/2
MDS 9216i	46	Embedded 2 ports and IPS-4, IPS-8, or 14/2
MDS 9506	128	IPS-4, IPS-8, and 14/2
MDS 9509	224	IPS-4, IPS-8, and 14/2

- MDS9216, MDS9216A, MDS9216i
- MDS9506, MDS9509

Software support

Operating system and network interface controller

- Windows 2000 Server, Advanced Server - All NICs supported by HP
- Windows 2003 Standard Edition, Enterprise Edition - All NICs supported by HP
- Windows 2003 64-bit - All NICs supported by HP

- Red Hat AS/ES/WS - All NICs supported by HP
- SuSE ES/SS - All NICs supported by HP
- HP-UX 11i v1, 11i v2 (0409) - All NICs supported by HP
- IBM AIX - All NICs supported by HP
- Netware - All NICs supported by HP
- Sun Solaris - All NICs supported by HP

Compaq network teaming software support

Compaq Network Teaming is supported with Windows 2000 and Windows 2003.

C-Series IPS management software support

- Cisco Fabric Manager
- Cisco Device Manager

iSCSI Initiator software support

- HP Windows iSCSI Initiator
- HP-UX iSCSI Initiator
- Cisco Windows iSCSI Initiator
- Microsoft Windows iSCSI Initiator
- SourceForge Red Hat iSCSI Initiator
- SourceForge SuSE iSCSI Initiator
- Cisco IBM iSCSI Initiator
- Novell iSCSI Initiator
- Cisco Solaris iSCSI Initiator

Configuration rules

This section lists the iSCSI limits and rules when using the MDS IP 4/8-Port Storage Services and MDS 14/2 Multiprotocol Services modules in MDS 9506 or MDS 9509 directors, or when using the 14/2 Storage Services module in the MDS 9216i.

Multipath software including HP Secure Path for MSA1000/1500, RA/MA8000, EMA/ESA12000, Enterprise Virtual Array, VA and XP is not supported on any iSCSI initiator.

Without multipath capabilities, the iSCSI initiator can only access one path of the storage controller. This will disable controller failover protection.

The C-Series IPS Module is supported on any supported HP Infrastructure C-Series Switch fabric. Please read “HP StorageWorks SAN Design Reference Guide” for the latest HP Infrastructure switch topologies and fabric rules.

Table 71: C-Series iSCSI limits

C-Series iSCSI limits	Maximum
Maximum number of Initiator/Target (IT) pairs per port	200
Maximum number active logical units (LUNs) per IT pairs	256
Maximum number of IT pairs/LUNs combinations per GigE port	1200

Here are some examples of maximum configurations for initiator/target pairs:

1. 200 iSCSI Initiators each connecting to 1 Target (Storage Controller Port)
2. 100 iSCSI Initiators each connecting to 2 Targets
3. 50 iSCSI Initiators each connecting to 2 Targets and
   100 iSCSI Initiators each connecting to 1 Target

Note: The maximum supported 200 TCP connections (initiator-target pairs) with 256 LUNs per connection yields a theoretical 51,200 possible LUNs. Simultaneous access to 1,200 LUNs via a single 100 Megabit Ethernet port may provide inadequate performance. The speed of the port and the number of configured TCP connections and LUNs configured should reflect the bandwidth requirements for the sum of all simultaneous storage accesses to that port.

Operating system rules for iSCSI with MDS IP 4/8-Port Storage Services and 14/2 Multiprotocol Services modules

- Host based multi-pathing software is not supported.
- Operating system clustering is not supported.

Network Attached Storage

13

This chapter covers the following major topics:

NAS / SAN integration overview

Customers typically base their decisions on SAN or NAS implementations by selecting file or block formats and by determining the hardware and software components they need. But with NAS integration in a heterogeneous SAN environment using the StorageWorks NAS devices, the decision to use either NAS or SAN systems becomes irrelevant.

With NAS/SAN integration, the customer benefits are significant and include the following:

Table 72: NAS/SAN integration features and benefits

Feature	Impact	Benefit
The decision to use NAS or SAN systems is irrelevant because customer should use both.	Fully integrated, converged storage architectures.	Ease of choice. Eliminates technical trade-offs. Easy to implement and grow.
Supports multiple operating systems.	Allows centralized data and increases flexible storage capacity and efficiency. Data can be shared across multiple operating systems regardless of the device or operating system.	Provides strategic power in the marketplace and consolidates data. Saves money by providing better disk utilization.
iSCSI Feature Pack	Provides iSCSI target capabilities for the DL380 Storage Server with selected storage configurations.	Ability to serve block-based storage such as is required by Microsoft Exchange, SQL Server and Oracle. Also provides an IP gateway into the Fibre Channel SAN.
Storage networks integrate with existing hardware	Less interruption to production systems and easier to manage.	Reduces Total Cost of Ownership (TCO).
Enterprise-wide file sharing reduces file duplication.	The right information to the right person at the right time in the right format.	Data sharing, which reduces storage capacity requirements.
Reliable and efficient access to data and application information 24x7.	Improves both local and wide-area data retrieval.	Improved data availability across the enterprise.
Seamless integration of storage devices and architecture.	New applications integrate reliably to ensure new data does not overload your system.	Faster time to market. Lower maintenance costs.
Data storage in centrally managed locations rather than across multiple application servers.	Facilitates administration, backup, and security.	Centralized deployment of specialized resources and skill sets.
Storage resources shared among a much larger number of processing systems and users.	Improved efficiency and simplified management.	Better asset allocation.
Data is highly accessible.	Business continues without interruption.	Provides strategic power in the marketplace.

StorageWorks NAS features

HP ProLiant DL380 G4/G2 Storage Server features

The HP ProLiant DL380 G4/G2 Storage servers fuse NAS and SAN, offering customers the greatest scalability and flexibility, providing cost effective management for their storage resources. This latest innovation from HP provides enhanced performance along with simplified, centralized storage and system management, ultimately saving customers resources, time, and money, lowering their total cost of ownership (TCO). The StorageWorks NAS delivers the fusion of NAS and SAN in a common, networked storage pool that provides customers with the flexibility to choose file (NAS) or block (SAN)-level access to best suit the needs of their applications.

HP ProLiant Storage Server iSCSI Feature Pack

The HP ProLiant Storage Sever iSCSI Feature Pack is powerful software that adds iSCSI target functionality to HP Storage Server (NAS) devices designed for small and medium business environments. This iSCSI Storage Server solution is capable of hosting file, print, and application storage (block) services on a single platform. For more information about this feature, see the iSCSI Feature Pack web page at:

http://h18006.www1.hp.com/products/storage/software/inas/index.html

- Windows Storage Server 2003 Support
- Multi-protocol File Serving (Windows, UNIX/Linux, NetWare, AppleTalk)
- iSCSI Feature Pack
- Users, groups share creation/management
- Quotas
- Manageability
- Backup support
- Anti-virus support
- Storage virtualization
- Snapshot capabilities
- High Availability
- Cluster support
- Redundant hardware components
- Data replication Support
- Integrated Lights Out (iLO) connectivity
- Featuring HP SAN connectivity
- 10/100 Ethernet (TCP Offload Engine - TOE, Optional)
- Gigabit Ethernet (TCP Offload Engine - TOE, Optional)
- Fibre Channel storage connectivity using the XP, VA, EVA, MA/RA/ESA/EMA, and MSA storage arrays
- Services; warranty uplifts
- Installation and configuration service
- CarePAQ Priority Services supplying 24x7 support with a maximum response ranging down to 2 hours for hardware and 30 minutes for software

HP ProLiant DL380 G4/G2 Storage Server hardware

The StorageWorks NAS is qualified in a heterogeneous open SAN as both a stand-alone or clustered server and follows the same Windows 2003 SAN hardware rules found in this guide.

See the *HP ProLiant DL380 G4 Storage Server QuickSpec* document for additional information, including the latest IP network controller hardware support.

StorageWorks NAS b3000v2 features

The StorageWorks NAS b3000v2 is the entry-point into NAS/SAN fusion solutions with enterprise-level availability, scalability and performance in a turnkey package that includes:

- Multi-protocol File Serving (Windows, UNIX/Linux, NetWare, AppleTalk)
- Users, groups share creation/management
- Quotas
- Manageability
- Backup support
- Anti-virus support
- Storage virtualization
- Snapshot capabilities
- Availability
- Cluster support
- Redundant hardware components
- Data replication Support
- Integrated Lights Out (iLO)
- Fibre Channel storage connectivity with MSA1000, EVA3000, EVA5000, and VA storage
- 1Gb and/or 2Gb HP SAN connectivity
- 10/100 Ethernet (TCP Offload Engine - TOE, Optional)
- Gigabit Ethernet (TCP Offload Engine - TOE, Optional)
- Services; warranty uplifts
- Installation and configuration service

StorageWorks NAS b3000v2 hardware

The StorageWorks NAS b3000v2 is qualified in a heterogeneous open SAN as both a stand-alone or clustered server and follows the same Windows 2000 SAN hardware rules found in this guide.

Note: See the StorageWorks NAS B3000v2 QuickSpec document for additional information, including the latest IP network controller hardware support:

http://h18006.www1.hp.com/products/quickspecs/11339_div/11339_div.html

---

StorageWorks NAS e7000v2 features

The StorageWorks Network Attached Storage (NAS) e7000v2 fuses NAS and SAN, offering customers the greatest scalability and flexibility, providing cost effective management for their storage resources. This latest innovation from HP provides enhanced performance along with simplified, centralized storage and system management, ultimately saving customers resources, time and money, lowering their total cost of ownership (TCO). The StorageWorks NAS e7000v2 delivers the fusion of NAS and SAN in a common, networked storage pool that provides customers with the flexibility to choose file (NAS) or block (SAN)-level access to best suit the needs of their applications.

- Multi-protocol File Serving (Windows, UNIX/Linux, NetWare, AppleTalk)
- Users, groups share creation/management
- Quotas
- Manageability
- Backup support
- Anti-virus support
- Storage virtualization
- Snapshot capabilities
- High Availability
- Cluster support
- Redundant hardware components
- Data replication Support
- Integrated Lights Out (iLO) connectivity
- Featuring HP SAN connectivity
- 10/100 Ethernet (TCP Offload Engine - TOE, Optional)
- Gigabit Ethernet (TCP Offload Engine - TOE, Optional)
- Fibre Channel storage connectivity using the XP, VA, EVA, MA/RA/ESA/EMA, and MSA storage arrays
- Services; warranty uplifts
- Installation and configuration service
- CarePAQ Priority Services supplying 24x7 support with a maximum response ranging down to 2 hours for hardware and 30 minutes for software

StorageWorks NAS e7000v2 hardware

The StorageWorks NAS e7000v2 is qualified in a heterogeneous open SAN as both a stand-alone or clustered server and follows the same Windows 2000 SAN hardware rules found in this guide.

See the *StorageWorks NAS Executor E7000v2 QuickSpec* document for additional information, including the latest IP network controller hardware support:

http://h18006.www1.hp.com/products/quickspecs/11004_div/11004_div.html

StorageWorks NAS 8000 features

HP StorageWorks NAS 8000 solutions provide easily managed network-attached storage (NAS) solutions in dedicated storage and SAN configurations for customers that require file-sharing flexibility. With an HP operating system optimized for file serving, NAS 8000 solutions attach directly to Ethernet networks, and deliver low maintenance and high uptime. Cluster technology is available for environments requiring mission-critical access to data.

The NAS 8000 solutions support Windows, UNIX, and Linux. Network administration, user access, and storage configurations are all easily managed through the Command View NAS or command line interfaces. The NAS Data Path Manager software enables management and control of the data paths.

- Multi-protocol File Serving (Windows, UNIX/Linux)
- Users, groups share creation/management
- Quotas
- Manageability
- Backup support
- Anti-virus support
- Snapshot capabilities
- High Availability
- Cluster support
- Redundant hardware components
- Integrated Lights Out (iLO) Connectivity
- 10/100 Ethernet
- Gigabit Ethernet
- Fibre Channel storage connectivity using the XP, VA, EVA 5000, EMA & MA arrays
- Services; warranty uplifts
- Installation and configuration service

StorageWorks NAS 8000 hardware

The StorageWorks NAS 8000 is qualified in a heterogeneous open SAN as both a stand-alone or clustered server and follows the same Linux SAN hardware rules found in this guide.

See the *StorageWorks NAS 8000 QuickSpec* document for additional information, including the latest IP network controller hardware support:

http://www.hp.com/products1/storage/products/nas/8000/specifications.html

Also see the *HP StorageWorks NAS 8000 SAN Storage Configuration Guide* technical white paper available at:

http://welcome.hp.com/country/us/eng/prodserv/storage.html

StorageWorks NAS SAN configuration and zoning rules

NAS Product	Source for Rules
HP ProLiant DL380 G4 Storage Server HP ProLiant DL380 G2 Storage Server StorageWorks NAS b3000 v2 StorageWorks NAS e7000 v2	Windows 2003 SAN Configuration and Zoning information
StorageWorks NAS 8000	Linux SAN Configuration and Zoning information

StorageWorks NAS SAN fabric rules

The HP ProLiant DL380 G4 Storage Server, HP ProLiant DL380 G2 Storage Server, NAS b3000v2, NAS e7000v2, and the NAS 8000 are supported in SAN fabrics consisting exclusively of switch models listed for the B-Series product line or exclusively of switch models listed for the M-Series product line.

StorageWorks NAS SAN storage rules

For additional information on supported storage system firmware versions, contact your HP field representative.

HP ProLiant DL380 G4 Storage Server storage rules

The HP ProLiant DL380 G4 Storage Server supports MSA, EVA, MA/RA/ESA/EMA, VA, and XP arrays. Versions of firmware supported are consistent with those supported under Windows 2003. See the individual storage array documentation or contact your HP representative for firmware support levels.

HP ProLiant DL380 G2 Storage Server storage rules

The StorageWorks NAS e7000v2 supports XP, VA, MSA, EVA, and MA/RA/ESA/EMA arrays. See the individual storage array documentation or contact your HP representative for firmware support levels.

StorageWorks NAS b3000v2 storage rules

The StorageWorks NAS b3000v2 supports MSA, EVA3000, EVA5000, and VA arrays. See the individual storage array documentation or contact your HP representative for firmware support levels.

Note: Full support of the latest array firmware requires the update of Secure Path to 4.0c and the A16 driver for the Emulex HBA.

StorageWorks NAS e7000v2 storage rules

The StorageWorks NAS e7000v2 supports XP, VA, MSA, EVA, and MA/RA/ESA/EMA arrays. See the individual storage array documentation or contact your HP representative for firmware support levels.

Note: Full support of the latest array firmware requires the update of Secure path to 4.0c and the A16 driver for the Emulex HBA.

StorageWorks NAS 8000 storage rules

The StorageWorks NAS 8000 supports SAN storage using the RA8000, MA8000, ESA12000, EMA12000, Enterprise Virtual Array 5000, HP Virtual Array 7xxx series, and XP series disk arrays. Recommended firmware versions are as follows:

- XP FW is 21.05.06 (for Tru64 FW is 21.04.32)
- VA 7x00 FW is HP18
- VA 7x10 FW is A100
- EVA5000 v2 FW is v2.002
- EVA5000 v3 FW is v3.000
- MA/RA/ESA/EMA FW is V87F-0

The StorageWorks NAS 8000 supports SAN storage using the Enterprise Virtual Array 5000.

The StorageWorks NAS 8000 can access storage on a variety of devices within a SAN device, including the HP Virtual Array 71x0 and 74x0 series storage, and HP XP series disk arrays using SecureManager VA or SecureManager XP.

Volume 5

Management and best practices

SAN management and best practices are presented in these chapters:

SAN management

14

HP is rapidly transitioning from the traditional server, storage, and component level-based management to a SAN-level application architecture and implementation using the Storage Management Appliance (SMA) or a general purpose Windows server running the SAN management software.

Just as important as the quality and feature set of the SAN's hardware is the effectiveness of the SAN management applications in tying these devices together and simplifying the complexity of the storage network. Whether using an HP standard topology, or a custom design using the StorageWorks SAN design rules, IT managers need to configure, monitor, and maintain the SAN, as well as plan for, and accommodate, growth.

The HP Open SAN management strategy is to:

- Simplify storage management using standardized web-based graphical user interfaces (GUIs) residing on easy-to-use, easy-to-implement storage management appliances.
- Centralize the management of multi-vendor Heterogeneous Open SANs in distributed and consolidated environments.
- Automate policy-based management.
- Optimize functionality by exploiting all available management levels such as appliances, SAN fabrics, and servers/storage.

Key to the StorageWorks SAN management strategy is the use of the Storage Management Appliance. HP SANs can be designed for local, centralized, or distributed data access. Regardless of the arrangement or location of the storage components and preferred data access method, storage environment management can be centralized using a Storage Management Appliance.

The Storage Management Appliance connects directly to the storage network through a Fibre Channel switch providing full access to all supported devices in the storage environment. Strategically located out of the SAN data path, the appliance allows data transfers to proceed independently between computers and storage devices. The appliance optimizes SAN availability and performance while streamlining manageability.

Note: For more information about using an Storage Management Appliance SAN, see "Storage Management Appliance rules and recommendations"on page 176.

Storage Management Appliance features/functionality

- Simple, unintrusive management of SAN elements
- High SAN performance because the appliance is located out of the data path
- High SAN availability, because data transfers occur independent of the appliance
- Support for multiple management and monitoring applications
- A web-based, centralized user interface
- No console operations for increased SAN management security
- Support for heterogeneous platforms attached to the SAN
- Higher utilization for processing applications on host servers
- Rack mountable, ease of installation and administration

OpenView Storage Management Appliance software

HP OpenView Storage Management Appliance software is included with and resides on the Storage Management Appliance, giving you access to the storage management appliance functions. Logging into Storage Management Appliance software anywhere over the web provides a single aggregation point to launch a variety of HP's SAN management applications to monitor and manage your storage network.

Zoning the HP Storage Management Appliance in a heterogeneous server environment

Whenever a storage management appliance is placed in a fabric with heterogeneous servers, it is recommended that a dedicated storage management zone be created. This zone is specifically for the storage management appliance and the elements it is to monitor and manage.

For example, create a zone called SANAPP_1_ZONE that would contain the appliance host bus adapter port WWN and the port WWNs of all the HSG or HSV controllers managed by this Storage Management Appliance. Because fabric devices can be in multiple zones, this will have no effect on other zones containing the same HSG and HSV controller port WWNs.

The storage management appliance communicates with HSG or HSV controllers in-band, within the Fibre Channel fabric itself. It is not necessary or recommended to include either the switch WWNs or server HBA port WWNs in this zone. Communications to these devices are done out-of-band; outside the fabric via TCP/IP.

Note: In the Continuous Access EVA environment, the storage management zone must include all storage management servers and any arrays involved in the Continuous Access replication relationship.

HP OpenView Storage Area Manager overview

HP OpenView Storage Area Manager is comprised of a comprehensive software portfolio that simplifies and automates the management of storage resources and infrastructure. It manages tape and disk, and direct and network-attached storage, across multivendor devices and distributed environments. From its central management console, IT storage administrators effectively monitor storage and storage service availability, performance, usage, cost and growth, while optimizing resources and cost.

HP OpenView Storage Area Manager also enables users to define, monitor and measure storage service levels, helping to guarantee quality of service and increasing the value of storage investments. IT management can determine and set enterprise-wide device, capacity and performance management, usage metering, storage allocation and access control.

The Storage Area Manager product suite includes the following applications:

- Storage Node Manager, for device management
- Storage Accountant, for usage metering and billing
- Storage Allocater, for storage provisioning and access control
- Storage Builder, for capacity management
- Storage Optimizer, for performance management

You can install all applications from the Storage Area Manager CD-ROM; however, you must purchase and enter operational licenses for each application in order to use them after the initial evaluation period.

The software suite's building block architecture allows each of its five software tools to operate and be available separately, so users can add functionality when needed and budgeted. Each tool focuses on a particular aspect of storage management, yet is designed to provide a seamlessly integrated view of the storage environment when used in conjunction with the other tools.

Key benefits:

- Simple, automated operations
- Intuitive, easy-to-use tools and automated wizards drive staff efficiency and shorten learning curve
- Central console and common reporting structure to manage and monitor storage service availability, performance, usage, cost and growth
- Automated identification of wasted or stale or secondary storage frees capacity
- Usage metering and billing recovers cost and enables charge-back
- Logical, online storage provisioning for storage allocation without impacting operations
- Performance monitoring identifies bottlenecks before they impact business operations
- Automated reporting saves time and increases service quality
- Automated host access notification ensures data is safe from network intruders
- Virtualized storage access control provides highest levels of data integrity
- Continuous health status and event monitoring quickly isolates and solves problems
- Multi-vendor device and system support maximizes storage investments, and provides open choice for future storage acquisitions

Storage Area Manager architecture

The HP OpenView Storage Area Manager hardware and software architecture is comprised of the bridge, management server, managed hosts, management clients, and the Manager of Managers (optional).

Bridge

The bridge is a Web server application that allows other applications access to Storage Area Manager's functionality, and enables Storage Area Manager's integration with other OpenView enterprise applications. The bridge consolidates information from multiple management servers for use by the application integrating with the bridge.

It is automatically installed on the management server when OpenView Storage Area Manager software is installed from the CD.

Management server

The management server is a server application that hosts the majority of Storage Area Manager's storage management functionality. Its framework includes the Storage Area Manager database, discovery subsystem, event-handling subsystem, configuration files, and server components for each of the five software tools that comprise the product suite.

The management server software is installed from the Storage Area Manager CD on a dedicated Windows 2000 server or workstation. A single management server manages a single storage domain, which consists of storage resources that are visible to the SAN hosts associated with the management server. Storage Area Manager can manage direct-attached, SAN-attached, or network-attached storage resources.

Managed host

The managed host contains the host agent software, which includes all components that require access to storage resources visible to the hosts. These components include discovery, status and event inquiry, and performance and capacity data collection.

The host agent software can be installed remotely from the management server or locally from the Storage Area Manager CD onto a Windows, HP-UX, Solaris, Linux, AIX, Tru64 UNIX OpenVMS, or NetWare host. Upon successful installation, the host becomes associated with and dedicated to the management server. The host agent runs as a service on Windows hosts and as a daemon on UNIX hosts.

Management client

The management client is a graphical user interface (GUI) application that uses a common navigation and presentation framework to display the storage information stored by the management server.

The management client software is automatically installed on the management server. It also may be downloaded from the management server to remote client systems running Windows, HP-UX, Solaris, or Linux.

Manager of Managers

The Manager of Managers (MoM) is a graphical user interface (GUI) application that consolidates storage information from multiple storage domains. This allows administrators to view, from a single location, the high-level status and filtered event information of a large or geographically dispersed storage network. They also can launch the management client for one specific storage domain to view the detailed information displayed by the client.

The MoM software can be downloaded from the management server to remote Windows, HP-UX, Solaris, or Linux hosts. It is an optional piece of the Storage Area Manager architecture.

OpenView enterprise applications

The Storage Area Manager software suite integrates with various OpenView enterprise applications. Through the bridge, the Storage Area Manager Smart Plug-in (SPI), and the integration packages contained on the Storage Area Manager CD, its information and control can be integrated with:

- HP OpenView Reporter
- HP OpenView Operations for Windows
- HP OpenView Operations for UNIX
- HP OpenView Internet Usage Manager
- HP OpenView Service Navigator
- HP OpenView Service Desk

For more information on the integration of Storage Area Manager with these enterprise applications, see the Storage Area Manager documentation.

Hierarchical multi-domain architecture

The Storage Area Manager software suite is designed to support hundreds of managed devices and thousands of LUNs spread across both logical and physical (direct, NAS, SAN) domains in the storage infrastructure.

Its optional Manager-of Managers (MoM) capabilities further ease the management of large distributed environments. It enables administrators to see a complete view of the distributed storage infrastructure, and allows individual administrators access to those infrastructure components for which they are responsible.

SAN management categories

SAN management is wide ranging, covering many aspects of the day-to-day activities used for monitoring and managing, as well as simplifying, the complexity of the storage network.

This section classifies SAN management into four major categories:

- Fabric management
- Storage management
- Data management
- SAN usage and monitoring

SAN fabric management

SAN fabric management can be thought of as the control of the SAN infrastructure or "traffic flow" within the SAN. This pertains to control and management of device communication or access within the SAN, such as switch zoning, or LUN level Selective Storage Presentation (SSP). This also includes managing SAN interconnect components, individually and collectively, throughout the fabric.

SAN storage management

Storage management allows control of the specific storage system configuration such as redundant paths, creation and management of storagesets (LUNs), setting of RAID levels, and the setting of platform specific SAN interface characteristics and parameters.

SAN data management

SAN data management applications help ensure that data is available and accessible. The data being stored on the SAN is part of a company's assets. It is imperative to keep this data available to system applications with minimal, if any, downtime. Techniques such as cloning, snapshots, data replication, and backups protect the data from disasters.

SAN/storage usage and monitoring

SAN and storage usage and monitoring applications are necessary to provide SAN event notification and fault/failure information for service before SAN anomalies can adversely impact the enterprise. They may also provide reporting and billing information for determining the amount of storage and quality of service delivered.

SAN management application deployment

Within the different categories of management tools, individual tools are implemented either on the storage management appliance, within fabric interconnect components, or within servers/storage systems. Table 73 lists the management tools by category, and identifies where the specific tools reside.

Note: Some applications may reside in more than one category.

Table 73: SAN Management tools & location

SAN Management Application	Appliance Based	Fabric Based	Server Based	Storage Based
SAN Fabric Management				
HP SANworks Network View	Yes	No	No	No
HP OpenView Storage Node Manager	No[1]	No	Yes	No
StorageWorks Fabric Watch	No	Yes	No	No
SAN/Fibre Channel Switch Management	No	Yes	No	No
HP StorageWorks Fabric Manager	No	Yes	Yes	No
HP StorageWorks HA-Fabric Manager	No	Yes	Yes	No
SAN Storage Management				
Storage Management Appliance Element Manager for HSG	Yes	No	No	No
Command View EVA	Yes	No	No	No
HP SANworks Network View	Yes	No	No	No
OpenView Storage Node Manager	No[1]	No	Yes	No
OpenView Storage Allocater	No[1]	No	Yes	No
StorageWorks Command Console	No	No	Yes	No
Storage System Array Controller Software (ACS) Command Line Interface (CLI)	No	No	No	Yes
Storage System Scripting Utility (SSSU)	No	No	Yes	No[2]
RA4000/4100 Array Configuration Utility (ACU)	No	No	Yes	Yes
MSA1000 (ACU, ACU-XE, ACU-XE(Offline)	No	No	Yes	Yes
Secure Path Manager	Yes	No	Yes	No
Storage Provisioner	Yes	No	No	No
SAN Data Management				
StorageWorks Business Copy (BC)	Yes[3]	No	Yes	No
OpenView Storage Virtual Replicator	No	No	Yes	No
StorageWorks Data Replication Manager (DRM)	No	No	No	Yes[4]
StorageWorks Command Scripter	No	No	Yes	No[5]

Table 73: SAN Management tools & location (Continued)

SAN Management Application	Appliance Based	Fabric Based	Server Based	Storage Based
Continuous Access EVA user interface[6]	Yes	No	No	No
SAN/Storage Usage & Monitoring				
OpenView Automation Manager	Yes	No	No	No
SANworks Network View	Yes	No	No	No
OpenView Storage Node Manager	No[1]	No	Yes	No
OpenView Storage Builder	No[1]	No	Yes	No
OpenView Storage Accountant	No[1]	No	Yes	No
OpenView Storage Optimizer	No[1]	No	Yes	No

1.HP OpenView Storage Node Manager, Allocater, Builder, Accountant and Optimizer are supported as options in the SAN Appliance.
2.This product is a character cell interface to configure and control an Enterprise Virtual Array.
3.BC Version 2 and later.
4.DRM requires ACS Version 8.xP Software. It is best managed via a character cell command line interface.
5.This product is a front-end to the Storage System CLI.
6.Requires Command View EVA, Continuous Access EVA License, VCS V3.

SAN fabric management tools

Storage Management Appliance Network View

HP SANworks Network View is a Storage Management Appliance application providing at-a-glance views of SAN configuration and availability. SAN devices, their Fiber Channel interconnects, and associated status are automatically discovered and represented in an intuitive topographical display. SAN device management is made easy by double clicking on a device icon. A Java based design allows remote SAN management from any web-enabled console having Internet Explorer or Netscape browsing capability. Network View serves as a SAN management consolidation point.

Software features/functionality

- Simplified SAN management from one application
- At a glance SAN visualization
- SAN administration from remote locations
- Automated SAN availability monitoring and notification when faults arise
- A consolidation point and launch pad for device specific device and storage management tools
- Fibre link connection mapping for established SANs or for future planning
- Scalable for future SAN growth
- A host independent solution

Storage Management Appliance Network View spans three SAN management categories by providing:

- SAN Fabric Management - Network View can view, monitor and manage Fibre Channel Switches, Tape Routers as well as interswitch links (ISL) right from the topology map. By either clicking on the device icon or the device folder Network View will automatically launch the device's web GUI.
- SAN Storage Management - HSG elements can also be viewed, monitored and managed from the Network View topology map. By clicking on the element icon or device entry, Network View will call up the respective Element Manager application.
- SAN monitoring - Network View can monitor the condition of the fabric hardware by displaying and reporting the condition of server HBAs, Fabric Interconnects, and HSG elements via E-mail, pager or SNMP traps. Performance can be monitored either in real time or a SAN history may be maintained to play back at any time.

Network View setup in a large SAN

Network View discovers, monitors and manages various Fibre Channel devices either through in-band or out-of-band communications with that device.

HSG Elements

Network View uses in-band and out-of-band communication to discover the HSG controllers or elements. Network View will automatically populate its database and topology map based on the HSG elements discovered by the appliance. HSG elements are displayed by the controller serial numbers discovered.

To clarify display and error reporting, change the properties of the HSG Element to a more intuitive name. Right-click on the element icon or device list to change properties. An example pattern for a name is: *location-controller type-failover mode- ACS version.*

For example: RACK05TOP-G80-T-V8.6F would indicate the element resides in the top of Rack 5, and are HSG80 controllers running in Transparent Failover with ACS code V8.6F.

Fibre Channel switches/Fibre Channel routers

Network View uses out-of-band communication (TCP/IP) to discover Fibre Channel Switches and Fibre Channel Routers. When Network View is launched for the first time a Configuration pop-up window will appear indicating the database is empty. You may then add a range of Fibre Channel Switch/Router IP addresses for Network View to map and monitor.

You may also at any time add additional IP addresses by clicking on the Configure button on the topology map, adding the new IP addresses, then Start discovery.

By default, Network View will display the DNS name of the IP address, if any, or the IP address itself. Right-clicking on the device icon or name will allow you to edit the name within the properties box that is displayed in the topology map.

A suggestion would be: *FCtopology-device-xx.*

For example: RING-SWITCH-01 would indicate the first Fibre Channel switch in the ring topology.

Note: Use at least 2 characters for numbers to keep the display sorted properly.

Server Host Bus Adapters

Network View does not initially discover server HBAs until a Device Manager Agent is installed on the server. During the agent install it will prompt you to input the appliance name running Network View. It is this device manager service that “pushes” the HBA information over TCP/IP to the appliance. Server device managers are available on Window NT, Windows 2000, Sun Solaris, Tru64 UNIX, HP-UX, IBM AIX, OpenVMS, and NetWare.

Network View will display the DNS name of the server in the topology map and it cannot be renamed. However it is suggested to append the HBA name located under the Host name for monitoring and error reporting purposes.

A suggestion would be to prefix the existing entry with: *servername-topology.*

For example: SERVER04-RING-Emulex-LP8000-Port0 would indicate this is SERVER04’s ring topology adapter.

For further information, including agent O/S versions, please read the Network View QuickSpec and release notes.

HP OpenView Storage Node Manager

HP OpenView Storage Node Manager provides a central management console from which multi-vendor storage and infrastructure resources can be centrally and automatically discovered, mapped, monitored, configured and maintained. In addition, Storage Node Manager serves as a common launch point for device applications. Storage Node Manager is part of the Storage Area Manager suite and is available as an individual product.

Key features and benefits include:

Table 74: Storage Node Manager features and benefits

Feature	Benefit
Centralized Management Console	Operators may troubleshoot from a single console. Adding, deleting or changing storage configurations and tracking data center environment changes are handled through a single interface.
Status and Event Monitoring	Maximize storage availability by proactively and quickly isolating and resolving events with alerts and alarms (event logs maintained for review at any time).
Auto-Discovery of Devices	Maximize availability—storage network changes are continuously and automatically identified and mapped.
Device Applications	Reduced configuration and troubleshooting time. Device applications launch from the management station.
Multi-Vendor Host Support	Choice among market leaders.
Graphical Device Maps	Visualize all aspects of the storage network, including redundant connections and device zones
Device Icons	Manage more efficiently through a standard set of icons.
Customizable Location Fields	Improved asset management through the clearly identified physical location of all storage devices in large, distributed (e.g., campus) environments.
Fibre Channel Zone Presentation	Efficient management. Easily identify zone members in the maps.

Fabric Watch

Fabric Watch allows the SAN manager to monitor key fabric and switch elements, making it easy to quickly identify and escalate potential problems. It monitors each element for out-of-boundary values or counters and provides notification when any exceed the defined boundaries. The SAN manager can configure which elements, such as error, status, and performance counters within an HP SAN Switch, are monitored.

Fabric Watch can be accessed through a web GUI, a Telnet interface, an SNMP-based enterprise manager, or by modifying and uploading the Fabric Watch configuration file to the switch.

Fabric Watch monitors the following elements:

- Fabric events (such as topology reconfigurations, zone changes)
- Switch environment (fans, power supplies, and temperature)
- Ports (state changes, errors, and performance)
- GBICs

With Fabric Watch, each switch continuously monitors error and performance counters against a set of defined ranges. This and other information specific to each monitored element is made available by Fabric Watch for viewing and, in some cases, modification. This set of information about each element is called a *threshold*, and the upper and lower limits of the defined ranges are called *boundaries*. If conditions break out of acceptable ranges, an *event* is considered to have occurred, and one or more alarms (reporting mechanisms) are generated if configured for the relevant threshold.

Please see Table 10 on page 83 for hardware support.

HP StorageWorks HA-Fabric Manager

As SANs expand and become more complicated, IT administrators need an efficient tool for managing the enterprise. HP StorageWorks HA-Fabric Manager (HAFM) is a comprehensive storage resource management application used to configure and manage HP's M-Series switch product line. HAFM simplifies SAN management, optimizes storage resources, and minimizes storage networking risks.

HAFM features include:

- complete management of the SAN from a single console
- integration with leading multi-vendor applications
- high levels of access and security
- scaling from department-level SANs to enterprise networks
- savings in time, money, and personnel resources
- detailed logging, diagnostics, and proactive alerts that monitor and ensure fabric health
- streamlined troubleshooting processes
- ease to use

HP StorageWorks HA-Fabric Manager - new features:

- Persistent Fabrics
- Improved Zoning
- User Interface Enhancements
- 2Gb/s Management Support
- New Product Managers

HAFM can be run locally on the HAFM Server platform or remotely on any network-attached user workstation in the enterprise. The Java-based deployment support gives IT administrators the flexibility to run HAFM from virtually any type or size of client device including Sun Solaris, AIX, HP-UX, Linux, Windows NT, Windows 95, Windows 98 and Windows 2000.

Please see Table 28 on page 107 for hardware support.

HP StorageWorks Fabric Manager

HP Fabric Manager is an application that manages multiple StorageWorks SAN switches and fabrics in real time. Fabric Manager provides the essential functions for efficiently configuring, monitoring, dynamically provisioning, and managing StorageWorks SAN fabrics on a daily basis.

Fabric Manager is tightly integrated with other HP StorageWorks SAN management products, such as Web Tools and Fabric Watch. Organizations can use Fabric Manager in conjunction with other leading SAN and storage resource management applications as the drill-down element manager for single or multiple fabrics.

Highlights

Fabric Manager version 4.x enables the user to:

- Provision, monitor, and administer large numbers of switches and multiple StorageWorks SAN fabrics with greater efficiency.
- Perform management tasks across multiple devices and fabrics as a single management operation.

- Intelligently group multiple HP B-series Fabric switches or ports to facilitate aggregated management.
- Visualize and track changes to SAN configuration and state information through multiple views at multiple levels of detail.
- Launch Fabric Manager from other enterprise management applications as the element manager for the fabric or multiple fabrics.
- Track SAN assets by using detailed table views that can be exported to a spreadsheet.
- Discover details about devices logged into the fabric, including Host Bus Adapter (HBA) asset information.
- View the SAN layout through a topology map that specifies Intersite Link (ISL), switch, and device details.
- Identify, isolate, and manage SAN events across large numbers of switches and fabrics.

SAN management: C-Series product line switches

The C-Series switches can be managed in several different ways:

- via the serial port/manager console
- via telnet over IP
- with the Cisco Fabric Manager over IP or IP over Fibre Channel

The Cisco Fabric Manager is a fabric-based, web-loaded application that provides fabric level (Fabric View) and switch (Device View) level management functions from most web-enabled clients. Fabric Manager also provides a Summary View mode that reports on port statistics.

SAN/Fibre Channel switch management

The HP Fibre Channel SAN Switches are high performance, scalable switch fabrics designed for creating large SANs. The management functions let you control and monitor fabric topology, frame throughput, error statistics, fans, cooling, media type, port status, and a variety of other information to aid in system debugging and performance analysis.

The administrative and diagnostic functions of the SAN switch are accessible from IP over the RJ-45 10/100BaseT Ethernet port or any Fibre Channel port. You can use any Simple Network Management Protocol (SNMP)-based management product to access the SNMP agent. You can also use any supported web browser to use the Java Web Management Tools.

Supported management methods include:

- SNMP
- Telnet
- Web-based Management Tools launched via Network View
- Telnet command subset via switch front panel display (FC SAN Switch/16 only)

OVSAM

There is an HP OpenView Device Plug-In's (DPIs) for B-Series, C-Series, and M-Series Switches. These DPIs extend Storage Node Manager and Storage Optimizer support to discover, map, and monitor the health and performance of the HP-supported switches. These DPIs can be obtained through this URL:

http://www.openview.hp.com/products/dpi/

SAN storage management tools

Command View EVA

Command View EVA is a SAN management application to configure and monitor HSV controllers. For each controller pair, Command View EVA enables you to:

- Initialize an Enterprise Virtual Array and create a pool of disk drives
- View, configure, and upload code to the controllers and disk drives
- View and configure virtual disks, and host properties
- Dynamically expand volumes for operating systems that support dynamic volume expansion
- Make temporary snapshots of volumes for backup purposes (with supported firmware, requires a license)
- Make snapclones to create an exact copy of another Virtual Disk at a particular point in time (with supported firmware, uses the same license as snapshots)
- View Enterprise Virtual Array event logs

VCS features and functionality

These features do not reflect the more restrictive requirements of solutions like Business Copy and Continuous Access EVA.

- Support for up to 240 disk drives per storage system
- Management of up to 512 virtual disks per disk pool ranging in size from 1GB to 2TB per virtual disk
- Dynamic capacity expansion and virtual disk data load leveling
- Distributed sparing of disk capacity
- Virtually Capacity-Free Snapshot (Vsnap)
- Virtually Instantaneous Snapclone
- Dual redundant controller operation for increased fault tolerance
- Multiple Bus Failover Support
- Battery Back-up
- Asynchronous Disk Swap (Hot Swap)
- Clustered Server Support
- Mirrored Write-Back Cache Support
- Read-Ahead and Adaptive Read Caching Support
- Virtual RAID Arrays (Vraid0, Vraid1, Vraid5)
- Non-disruptive software upgrade capability
- Supports connection of up to 256 hosts
- Multi-Vendor Platform Support
- Controller Password Protection for Configuration Control
- Selective Storage Presentation and SAN-based data zoning
- GUI Interface for management and monitoring

Supported management methods include:

- SSSU
- Storage Management Appliance

VCS works in a heterogeneous environment that includes Tru64 UNIX, OpenVMS, Microsoft Windows NT and Windows 2000, and Sun Solaris. This application is at the storage system level.

Command View EVA restrictions

Current restrictions of the Command View EVA must be enforced:

- A maximum of 16 Enterprise Storage Systems can be managed by one Storage Management Appliance.
  - In a Continuous Access environment, this maximum reduces as the distance between the SMA and the EVA storage increases.
- An Enterprise Storage System can be actively managed by only one Storage Management Appliance.

General HSV storage system configuration process

The following steps highlight the configuration process for storage systems. See your storage system user guide for more information.

1. Set up the storage system according to the product user guide.
2. Ensure that you have connected the Fibre Channel from HSV controller ports to optical interfaces found on the Fabric Switch.
3. Ensure that appliance and the HSV controller port WWNs are in a Storage Management Appliance zone.

Initially, Enterprise Storage Systems are display as “UNINITIALIZED” on the Command View EVA browser window. It is recommended that, when the storage system is initialized, a intuitive name is used for display and monitoring convenience.

1. In the navigation pane, click on a controller icon.
2. In the Content Pane, click the INITIALIZE button.

   A pop-up message confirms that you are initializing the storage system. It also states that any data associated with the selected system will be lost, and then asks if you wish to proceed with the initialization procedure.

   If you have not previously entered the license key for the storage system, you will be prompted to do so.
3. Enter a name for the storage system.

   A suggestion would be:

   location-controller type- VCS version -

   For example: RACK05-V110-V2002 would indicate the element resides in Rack 5, and contains HSV110 controllers running VCS code V2002.
4. Specify the number of disks in the default group. Enter from 8 up to the total number of drives in the subsystem.

5. Click on the **Advanced Options** button and set the date and time option. It is recommended that you synchronize the time with the Storage Management Appliance time. If this practice is used with all controllers, then you will have synchronized times on all event log entries. Always use the same Storage Management Appliance to initialize all EVAs or synchronize the time of all Storage Management Appliances to the same source.

   Leave the Console LUN ID set to "**0**".

Element Manager for HSG

Element Manager for HSG is a Storage Management Appliance application to configure and monitor HSG80/60 controllers. For each controller pair, Element Manager for HSG enables you to:

- View existing virtual disk, controller, physical disk, and host properties
- Make changes to these properties for different configurations
- Configure Remote Copy sets and add associations (with supported firmware)
- Dynamically expand volumes for operating systems that support dynamic volume expansion
- Make temporary snapshots of volumes for backup purposes (with supported firmware)

HSG Element Manager restrictions

Current restrictions of the HSG Element Manager must be enforced:

- A maximum of 25 HSG Storage Systems can be managed by one Storage Management Appliance.
- An HSG Storage System must not be visible to more than one Storage Management Appliance.

Note: The HSG Element Manager is not a supported management tool for DRM environments. Contact HP consulting services for information about managing DRM.

Storage Management Appliance and HSG storage system communication

When configuring an HSG controller in a Storage Management Appliance storage environment, you will need to enable CCL (i.e. setting the controllers to SCSI-3 or SCSI-2 with CCL Enabled) or provide a dedicated LUN (i.e. setting the controllers to SCSI-2 CCL Disabled) for the Storage Management Appliance. If you are using a dedicated LUN instead of CCL, verify that the LUN is presented to the Storage Management Appliance through each controller host port connection for the Storage Management Appliance (there are two host ports per controller). The LUN may be a partition. For more information, see your HSG controller user guide.

General HSG Storage System configuration process

The following steps highlight the configuration process for storage systems. See your storage system user guide for more information.

1. Set up the storage system according to the product user guide.
2. Ensure that you have connected the Fibre Channel from HSG controller ports to optical interfaces found on the Fabric Switch.

3. Ensure that the appliance and the HSG controllers port WWNs are in a Storage Management Appliance zone.

Note: The following steps are only necessary if the HSG controller is configured for SCSI-2 CCL Disabled.

---

   a. Connect to the HSG controller via the serial interface. See your HSG controller manual for further information on the serial connection.
   b. Start a terminal session. See your HSG controller manual.
   c. At the prompt in the terminal session, enter the Command Line Interface (CLI) command SHOW CONNECTION.
   d. Verify, via HSG connection table, that the Fibre Channel HBA of the Storage Management Appliance is online. Use CLI command SHOW CONNECTION.
   e. Create Logical Unit Number (LUN) and enable access from the LUN to the Storage Management Appliance.

4. Verify that the HSG controllers are discovered. To verify HSG status, you will need to configure and launch Element Manager for HSG and click on OPTIONS. Enable controllers that are discovered.

HSG elements are displayed by the controller serial numbers discovered. It is recommended to change the properties of the HSG Element to a more intuitive name for display and error reporting reasons.

1. In the navigation pane, click on a controller serial number displayed.
2. In the Content Pane, edit the ALIAS field and save changes.

A suggestion would be: *location-controller type-failover mode- ACS version.*

For example: RACK05TOP-G80-T-V8.6F would indicate the element resides in the top of Rack 5, and are HSG80 controllers running in Transparent Failover with ACS code V8.6F.

HSG Storage System Array Controller Software/Command Line Interpreter

HSG Array Controller Software (ACS) for Fibre Channel Arbitrated Loop and Switched Fabrics provides storage controller software capability for the StorageWorks HSG60 and HSG80 Array Controllers in Fibre Channel arbitrated loop and switched fabric environments. HSG Array Controller Software is designed to be common across multiple operating system platforms. However, there may be operational differences between platforms, and there may also be features that are not supported on every platform.

Management of storage systems based on the HSG60 or HSG80 is provided directly through the controller serial port using a terminal or a terminal emulator (such as Microsoft Windows NT HyperTerminal) using the CLI interface. The CLI provides all the commands necessary to configure controller failover modes and parameter settings, controller and host connections to the SAN, storageset creation, SAN LUN access (SSP), RAID levels, and cache settings. The CLI also provides access to the array controller utilities. The utilities are used to monitor controller functions and statistics, and to allow storage system component replacement procedures to be conducted while the storage system is active.

Selective Storage Presentation

Selective Storage Presentation (SSP) provides a way to control SAN access at the storageset or LUN level. SSP allows each server or HBA's storagesets (LUNs) to be presented exclusively to those that are allowed access. Additionally, SSP allows the setting of host modes and LUN offsets for each HBA connected to the storage system. The host mode is specially tailored to the storage communication techniques of the operating system type. The LUN offset feature of SSP allows higher numbered LUNs in a storage array to be presented in a range required by specific operating systems. The SSP feature also provides a way to track the numerous Fibre Channel HBAs within servers attached to a SAN by identifying each by name and WWN.

ACS features/functionality

Solutions such as Data Replication Manager may impose stricter limits than those shown here.

- Host Interconnect and Protocol Services
- Microsoft Cluster Server (MSCS) Support
- Dual Redundant Controller Operation
- Testing and diagnosis of the HSG array controller
- SCSI device control
- Transparent Controller Failover Support
- Multiple-Bus Failover Support
- Asynchronous Disk Swap (Hot Swap)
- ACS system management services
- Local program support
- Mirrored Write-Back Cache support
- Read Ahead Cache support
- Disk Mirroring capability (RAID 1)
- Disk Striping capability (RAID 0, 0+1)

- RAID capability (RAID 3/5)
- Storageset Expansion
- Disk Partitioning capability

Supported management methods include:

- Terminal emulation through the HSG's serial port using the CLI
- Command Console
- Command Scripter

ACS works in a heterogeneous environment that includes Tru64 UNIX, OpenVMS, Microsoft Windows NT and Windows 2000, Novell NetWare, Sun Solaris, HP-UX, SGI IRIX, IBM AIX, Linux x86, and Linux Alpha. This application is at the storage system level.

HP OpenView Storage Allocater

The HP OpenView Storage Allocater software delivers a central, unified method for virtualized storage access control and LUN-level storage assignment, enabling you to build and manage complex storage area networks (SANs) with heterogeneous servers and storage devices.

Storage Allocater is part of the Storage Area Manager suite and is also available as an individual product.

Key features and benefits include:

Table 75: HP OpenView Storage Allocater features and benefits

Feature	Benefit
LUN level storage assignment	Optimized storage utilization
Storage security controls	Prevent data loss and unauthorized access
Highly scalable	Cost effective, simplified management: one software solution may be used in a range of configurations
Add/remove/assign storage without host reboots	Increased system availability
Mirrored SAN configuration	High system availability thanks to no single point of failure
Automated storage network with host and storage discovery device capabilities	Dramatically reduce configuration time and eliminate a major source of errors
Intuitive Graphical User Interface (GUI)	Maximize productivity and minimize training time with familiar techniques (drag-and-drop) and controls (view filters)
Share groups for cluster configurations	Visualized storage prevents errors when setting up cluster server and shared tape device environments
Fibre Channel topology independence	Enhanced SAN configuration flexibility
Supports open system, heterogeneous environments	Easy-to-use, cost-effective, single tool provides storage security
Native file system and raw disk support	Simplifies moving existing storage to SANs
Automated rogue host detection and notification	Enhanced data/information security, eliminate component/driver tampering

StorageWorks Command Console

StorageWorks Command Console (SWCC) is a feature-rich, graphical user interface providing local and remote management of StorageWorks HSG60 and HSG80 array controllers. It is a user-friendly tool for monitoring, configuring, and troubleshooting HP HSG60 and HSG80 storage arrays and controllers.

SWCC can be connected to your StorageWorks controller in several ways. Once connected, the program issues commands and interprets the responses sent by the controller. The user interface displays the logical and physical layout and status of a selected subsystem in graphical form. Command Console consists of two major components: the Client and the Agent. The Client, which includes the user interface and some additional services, provides a window into your storage subsystems. The Agent is a host-resident program that is an interface between the Client and the host's storage controller to interpret and transfer information.

The Agent acts as the Client's assistant in controlling your storage subsystem. The Agent continuously monitors the subsystem and notifies the Client of changes. Commands sent from the Client are received by the Agent and are routed to the storage subsystem via the subsystem's Fibre Channel bus. Subsystem status is transmitted back to the Client from the Agent via the network connection.

Software features/functionality

- Easy, graphical configuration of the storage subsystem using the graphical user interface.
- Graphical view of the controller and its physical and logical storage elements.
- Status monitoring of the storage subsystem using intuitive icons.
- Fault notification by pager, electronic mail, and event log entries.
- Management of multiple host systems through a TCP/IP network connection.
- Direct serial port connection.
- Direct SCSI port connection (Windows NT and Windows 2000 Only).
- Robust security that prevents unauthorized access to configuration capabilities.
- The Client supports Microsoft Windows NT 4.0 and Windows 2000.
- The Agent supports HP-UX, Tru64 UNIX, OpenVMS, IBM AIX, SUN Solaris, Linux, Novell Netware, and Windows (NT4.0 and Windows 2000).

All G80 DRM should be placed in separate zones from SWCC.

Array Configuration Utility for RA4000/4100/MSA1000

The HP Array Configuration Utility (ACU) software (for Smart Array products, StorageWorks RAID Array 4100/4000 systems, and MSA1000 systems) makes it easy to configure and expand your disk drive arrays. This graphical tool is very intuitive: by using its Configuration Wizards, you have the ability to configure your array controller, add additional disk drives to an existing configuration, or completely reconfigure your disk drive array.

Software features/functionality

- Selective Storage Presentation: allows RA4100 and MSA1000 array sets to be partitioned to multiple servers for SAN access.
- Online RAID Level Migration: allows for online post-configuration change to RAID level without destroying data or volume information.
- Online Capacity Expansion: lets you add storage to an operational RA4100 or MSA1000, reducing expensive server downtime.
- Online Volume Extension: allows for the capacity growth of existing logical volumes.
- Global Online Spare: reduces the risk of data loss by facilitating automatic rebuilds after a drive failure.
- Logical Drive Capacity Extension: allows the user to increase the size of existing logical drives online under Windows NT and offline for other operating systems.
- Pre-Failure Warranty: Drives installed in an RA410 or an MSA1000 and monitored under HP Insight Manager are supported by a Pre-Failure (replacement) Warranty.

Note: Pre-Failure Warranty allows for the replacement of designated drives in an RA4100 before they actually fail when using HP Insight Manger on HP servers.

Note: Some operating systems may not support all of these features.

Secure Path multi-path software

Depending on the platform or operating system, high availability functionality may or may not be embedded in the operating system I/O drivers. Tru64 UNIX and OpenVMS operating systems have the ability to create and maintain multiple paths over the SAN to the same LUN, with support for these functions embedded. For those operating systems that do not support multi-pathing, HP provides this capability using HP Secure Path.

The HP Secure Path product provides continuous data access for HP RAID storage systems accessed by operating systems that are both HP-based and not HP-based. When combined with the inherent fault-tolerant features of the RAID Array, this configuration effectively eliminates single points of failure in the storage system.

When a host bus adapter, cable, or controller in a path fails, the failure is detected and I/O is automatically re-routed to the functioning, alternate path. This process, called failover, requires no resource downtime and ensures high availability of data. Storage units that have experienced failover may be configured to failback automatically after a path is restored. Failback can also be done manually through the use of the Secure Path Manager (Windows) or via spmgr (UNIX).

Software features / functionality

- Switched fabric and loop support
- Automatic path failover
- I/O load distribution
- User-selectable failback
- Supported on the Storage Management Appliance

Secure Path works in a heterogeneous environment. See "Heterogeneous server rules" on page 127 and "SAN storage system rules" on page 171.

Secure Path Element Manager on the Storage Management Appliance

Managing Windows NT, Windows 2000, and Windows 2003 Secure Path servers throughout the enterprise is now available using the Storage Management Appliance. Secure Path Element Manager uses the easy-to-use Storage Management Appliance Web GUI interface to manage and monitor hosts and HSG60/80 and EVA3000/EVA5000 storage subsystems and integrates with the Storage Management Appliance's notification utility.

The notification utility provides centralized functionality. The web-based Notification console is used to provide a single, modular, networked software unit that has the ability to handle Event Logging, SMTP, SNMP and command line launching operations.

Secure Path Element Manager uses TCP/IP to communicate with Secure Path servers. Adding the Storage Management Appliance server name to the Client list on the Secure Path Server will allow Secure Path Element Manager to discover the Secure Path server and add it to a profile.

SAN data management tools

Business Copy

Business Copy (BC) is web-based application software that manages controller-based clone and snapshot operations. Cloning is a mirroring copy function that allows you to create an exact copy of a LUN; snapshot provides an instantaneous point-in-time copy function. BC can be used to meet business continuance requirements by minimizing application downtime required for system backups and data migration activities. BC automates the creation of command files that control the cloning or snapshot operation. BC also allows you to mount the clone or snapshot on a second host on the same controller.

BC is a host-based tool that can be accessed by the user directly via a GUI or remotely via a browser. The tool then interacts with the requested application(s) to stop new I/O and flush pending I/O to disk. Once the I/O is stopped, BC instructs the storage via the in-band CLI to perform a snap or clone operation. Finally, on completion of that operation, BC restarts the application.

BC automates the creation of command files that control the cloning or snapshot operations. BC also allows users to mount the clone or snapshot to a new host. The new host can then act as a dedicated backup server or data warehouse server. All operations are performed on the clone or snapshot, minimizing performance impact on the production system.

Software Features/Functionality

- Web-based application
- Supported on the SMA or a general purpose server with Command View EVA installed
- Easy management of complex cloning and snapshot operations
- Supports LAN-less backup
- Simplified, centralized storage management

Business Copy works in a heterogeneous host environment that includes Tru64 UNIX, Microsoft Windows NT, Windows 2000, and Sun Solaris. This application is at the server level.

Business Copy on the Storage Management Appliance

Managing Windows NT and Windows 2000 Business Copy (BC) servers throughout the enterprise is available using the Storage Management Appliance.

Business Copy uses TCP/IP to communicate with BC servers. Adding the Storage Management Appliance server name during the BC server agent setup will allow the BC application on the Storage Management Appliance to discover the BC server.

Virtual Replicator

The HP OpenView Storage Virtual Replicator (VR) combines a rich set of innovative capabilities that enhances and simplifies storage management for Microsoft Windows NT and Windows 2000 environments. Through virtualization, online volume growth, snapshot and management features, the software complements the standard capabilities within the operating system.

Virtual Replicator utilizes industry-standard server, storage, and network-interconnect components, protecting an organization's current and future storage investments.

Storage Virtual Replicator provides the ability to create instant, virtual snapshots of production data without having to physically copy it. A snapshot, which looks exactly like the original disk from which it was copied, takes seconds to create and allows customers to back up and restore data with minimal impact to users and applications. Customers can schedule automated snapshot backups using the integrated policy-based scheduling and scripting features.

Software features/functionality

- Virtualization:

  Allows companies to respond quickly to rapidly changing storage capacity requirements. With storage virtualization, multiple storage arrays can be grouped into a pool of disk space for individual or clustered systems to use. Multiple high-capacity virtual disks, up to 2 terabyte in size, can be created from a pool for users and their applications. System administrators can tailor disk space to specific requirements for maximum utilization of storage resources.

- Online volume growth:

  Enables easy, nondisruptive growth for Windows 2000 with zero downtime. Online Volume Growth allows a system administrator to grow an existing volume on a Virtual Replicator virtual disk and also on a Windows 2000 basic disk. The system will remain online, and the data on the volume will remain intact.

- Snapshots:

  Enable the instant creation of multipurpose virtual replicas of production data without the requirement of a physical copy. Snapshots function identically to ordinary physical disks with both read and write capability. Whenever a quick copy of production data is needed, snapshots can be used with minimal disruption to running applications. For example, the snapshot can be the source for backup using standard backup tools. Snapshots can remain online for restore operations, testing, and data mining.

- Management:

  Simplification through easy-to-use interfaces using Microsoft Management Console or a command line. Interactive wizards are available to guide the administrator through all management tasks and create automatic schedules of operations.

Virtual Replicator provides server-based virtualization and is supported on Microsoft Windows NT and Windows 2000 (Server and Advanced Server.) VR is cluster-aware to ensure business continuance.

Continuous Access EVA

Continuous Access EVA is a Fibre Channel storage controller-based data replication (remote mirroring) solution to support disaster tolerance requirements. Continuous Access EVA works with the HP StorageWorks Enterprise Virtual Array storage system, which contains the HSV virtualized RAID controller. The HSV controller and the Virtual Controller Software (VCS) Version 3.0 and 3.01 enhance the virtualization with remote replication technology.

Continuous Access EVA copies data online and in real time via synchronous (and in Version 3.01 or later asynchronous) replication to a remote EVA through a local or extended storage area network (SAN). Additionally, data replication can be bidirectional, meaning that a storage array can be both a source and a destination. A single EVA may have a replication relationship with up to two other arrays, using different DR groups for each relationship. A particular LUN can be replicated in only one direction between the two storage arrays. Write I/O is sent to the

source and then replicated by Continuous Access to the destination. Properly configured, Continuous Access EVA can be a complete disaster-tolerant storage solution that guarantees data integrity in the event of a storage system or site failure.

A basic two site Continuous Access EVA configuration is shown in Figure 72.

While Continuous Access EVA may be used to satisfy data distribution or data migration requirements, this version of the Design Guide does not provide design recommendations for those solutions.

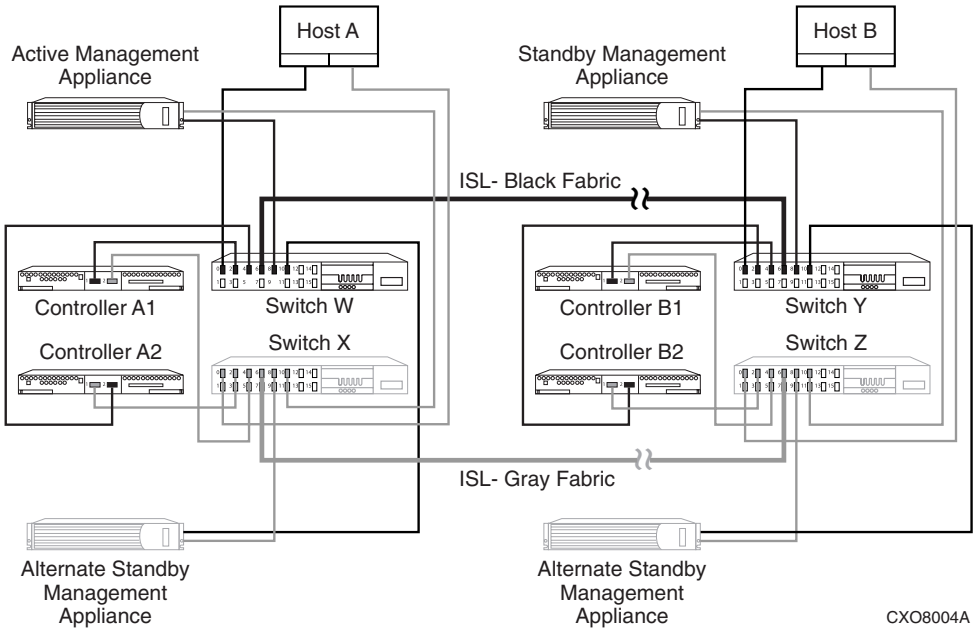

Figure 72: Continuous Access EVA basic configuration

Features

The following are prominent features of the Continuous Access EVA V1.1 solution:

- In-order synchronous or asynchronous remote replication (remote mirroring)
- Automated failover support
- Normal and fail-safe data protection modes of operation
- Dual redundant controller operation for increased fault tolerance
  - No single point of failure
  - Pairs of arrays share replication relationship and any one array may share a relationship with up to two arrays
  - Replicated write-back cache support
  - Read-ahead and adaptive read caching support
  - I/O continuation during normalization and merging
- Intersite link suspend and resume operations
- Multi-vendor platform support
- Dynamic capacity expansion, if supported by OS
- Merge of write history log in write order
- Failover scripting

- Multiple bus failover support
- Continuous Access user interface
- Asynchronous disk swap (hot swap)
- Controller password protection for configuration control
- GUI interface for management and monitoring
- Selective storage presentation for SAN-based data zoning
- B-Series, C-series, and M-Series switch support
- Virtual RAID arrays (Vraid0, Vraid1, and Vraid5)
- Virtual RAID (Vraid) techniques that:
  - improve performance
  - increase disk utilization efficiency
  - dynamically expand storage capacity
- Virtual disk data load leveling
- Clustered server support
- Distributed sparing of disk capacity
- Non-disruptive software upgrade capability
- LiteTouch Management
- Battery back-up
- Bi-directional replication
- Copy set size of 1 GB to 2 TB in 1 GB increments
- Up to 128 remote copy sets
- Up to 128 DR groups
- Up to eight copy sets per DR group
- Management of up to 512 virtual disks per EVA ranging in size from 1 GB to 2 TB per virtual disk
- Maximum of 240 FC-AL drives (with expansion cabinet) per storage system
- Maximum of 8 storage systems on the SAN management zone
- Maximum of 128 Fibre Channel adapter (FCA) connections per pair of arrays, and if a dual fabric, then 64 per fabric
- Maximum of 256 LUN presentations from an EVA to a single FCA
- Switch and hop count limits as specified in Chapter 4, Chapter 5, and Chapter 6
- Dark Fibre to 35 km at 2 Gbps or 100 km at 1 Gbps
- 2 Gbps end-to-end Fibre Channel solution
- 100 ms latency with wide area network gateways
- Virtually capacity-free snapshot (Vsnap)
- Virtually instantaneous snapclone
- Snapclone across physical disk groups
- Multiple snaps of the same data (both source and destination)
- Maximum of 8 snapshots or snapclones per DR group at the local or remote site
- The option of selectable World Wide Names (WWNs) for Vsnaps, snapshots, and snapclones

Data Replication Manager

The HSG80 Data Replication Manager (DRM) Software is the software component of the HSG80 array controller used in switched fabric environments for remote data replication. Data Replication Manager Software is a storage-based disaster tolerance and workload migration solution that provides the ability to copy data, in real time, to a remote location, up to 100 km away using direct Fibre Channel or further using either Fibre Channel over ATM or Fibre Channel over IP links. This is done without any host involvement. The HSG80's dual host port design, when used in DRM configurations, allows for long distance mirroring in a switched No Single Point of Failure (NSPOF) Fibre Channel topology.

The DRM software executes in the HSG80 array controller and processes I/O requests from the hosts, performing the local and remote device-level operations required to satisfy the requests. This is done through the use of a pair of initiator and target controllers sharing a switched NSPOF Fibre Channel fabric. Host generated Reads are performed on the local copy of the data. Host generated Writes to the local storage are copied by DRM from the local controller directly to the remote controller automatically. This capability provides the ability to maintain the same data at both locations, providing disaster tolerance protection.

Software features/functionality

- Online, real-time data replication to a local or remote site
- Data replication over a Fibre Channel SAN
- Cloning at Initiator and Target sites
- Snapshot support at Target site
- Cascaded switches support
- Full Fibre Channel-to-ATM connectivity with line speeds of T1 through OC3
- Full Fibre Channel-to-IP (FCIP) connectivity with line speeds of T1 through 1 GbE
- Replicate up to 100 km(~63 miles) with Very Long Distance GBIC
- Asynchronous and synchronous transfer modes
- Write History Logging and "Mini-Merge" reconstruction
- Stretched Clusters capabilities for Microsoft Windows NT and OpenVMS
- Association sets
- Non-RCS LUN support
- Switch Zoning support
- Wavelength Division Multiplexing

Data Replication Manager works in a heterogeneous host environment that includes HP OpenVMS, HP Tru64 UNIX, HP-UX, IBM AIX, Microsoft Windows NT, Windows 2000, Novell NetWare, and Sun Solaris. The application is at the storage system level.

See the DRM Design Guide, which is available online, for additional details.

Command Scripter

Command Scripter is application software that provides command-level control of HP StorageWorks systems equipped with HSG60, HSG80, HSZ70, and HSZ80 Array Controllers. With Command Scripter, you can create, edit, and run script files that contain StorageWorks Command Line Interpreter (CLI) commands. This allows automation of frequently performed StorageWorks operations.

Two interfaces are included in Command Scripter: a command line interface for local, direct connection to StorageWorks controllers and a web-based interface, which requires StorageWorks Command Console (SWCC) for centralized, remote connection via browser.

Software features/functionality

- Web-based interface for centralized, remote connection to StorageWorks array controllers
- Command line interface for local, direct connection to array controllers
- Select agent host and StorageWorks subsystem
- Create and edit CLI script files
- Run saved CLI script files
- Execute a single CLI command
- Display CLI command history

Command Scripter works in a heterogeneous host environment that includes Tru64 UNIX, OpenVMS, Microsoft Windows NT, Windows 2000, Sun Solaris, HP-UX and AIX (command line interface only), and AIX (command line interface only). This application is at the server level.

Storage System Scripting Utility

The Storage System Scripting Utility (SSSU) is the character cell interface for a user. The host based application should use the EMClientAPI to access the Command View EVA. The API transports SOAP/XML requests over the wire to the element manager, handling security and communication. The EMClientAPI provides an efficient machine interface to the Command View EVA, specifically designed for host-based applications.

SAN storage usage and monitoring tools

Automation Manager

Automation Manager provides a tool with which a storage administrator can automate the management of a storage area network. Automation Manager runs, controls, and manages predefined policies that storage administrators can configure for their environment. Predefined policies are provided with the product as Perl scripts. In addition, you can create and import your own management scripts.

Automation Manager also provides the following utilities to assist in managing storage operations:

Reports – View and print status reports about storage operations.

Agents – View and download an agent to hosts on which scripts reside. Agents enable Automation Manager to communicate with and run batch jobs on host systems.

Notification – Set up different notification types for Automation Manager events. The notification utility provides centralized functionality. The web-based Notification console is used to provide a single, modular, networked software unit that has the ability to handle Event Logging, SMTP, SNMP and command line launching operations.

HP OpenView Storage Builder

The HP OpenView Storage Builder is a storage inventory and resource planning tool for direct attached, SAN attached and network attached storage, and enables you to monitor, manage storage capacity and plan for future storage demands.

Storage Builder is part of the Storage Area Manager suite and is available as a separate product.

Key features and benefits include:

Table 76: HP OpenView Storage Builder features and benefits

Feature	Benefit
Centralized view of allocated vs. unallocated storage, and used vs. unused storage-by application, host, storage device, LUN, partition, volume, directory and user	Understand how much storage is assigned, being used, available for deployment, and how to balance capacity across systems/users. Make better use of storage resources. Lower total cost of ownership.
Automated thresholds warning system: set thresholds on hosts, partitions, volumes, directories and user	Receive early warning of capacity shortfalls that could cause system outage or user inconvenience. Capacity quotas on a per-user basis ensure storage growth is in line with company goals.
Group hosts, interconnects, bridges, NAS devices and storage components to reflect departments or physical locations, then create screens, reports and thresholds for these storage groups	Better identification of major users and heavily used devices. Establish norms and quotas to ensure storage growth is in line with company goals.
Identification of junk or stale files selectable by extension, such as MP3 or games, that waste valuable storage space	Free up primary storage capacity for use in meeting business goals.
Screens and reports that rank hosts by the amount of storage accessed each day	Achieve better asset utilization, higher availability and centralized management.
Historical trending of storage capacity data through screens and reports. Future extrapolation of historical data	Anticipate storage capacity shortfalls. Plan for and justify just-in-time purchases of additional storage capacity. Improve storage resource management efficiency and utilization rates, which lowers total cost of ownership.
Tabular and graphic reports showing allocated vs. unallocated storage, as well as consumed vs. free storage	Clearly communicate facts, trends and analysis concerning storage resources to staff and upper management.
Volume management	Provide better visibility into the host to LUN utilization mapping.
Capacity information can seamlessly be integrated into HP OpenView Internet Usage Manager (IUM)	Allow for centralized usage analysis, billing and charge-back.
Applications are automatically discovered and elements of the applications are reported and mapped to the storage devices	Provide better visibility of how the applications are utilizing the storage devices in order reduce costs associated with application capacity.

HP OpenView Storage Accountant

The HP OpenView Storage Accountant provides a toolset to measure, or meter, storage assigned to users (customers/organizations) for financial analysis, budgeting and charge-back.

Storage Accountant is part of the Storage Area Manager suite and is available as an individual product.

Key features and benefits include:

Table 77: HP OpenView Storage Accountant features and benefits

Feature	Benefit
Create and manage customer accounts and organizations	Better customer service. Analyze storage service usage on customer/organization basis. Define greater levels of granularity within one customer, organization.
Define and apply service levels	Reduce costs by providing the required type of storage based on usage analysis.
Assignment of storage to accounts	Measure assigned storage for tracking consumption, budgeting and financial analysis.
Automated billing Detailed usage and billing views and reporting	Recover costs of providing storage services. Manage relationships.
CSV, HTML and XML output	Export charge-back information to third-party applications for billing and financial analysis.
Audit log maintenance	Track events related to customers, service levels and storage consumption.
Seamless integration with HP OpenView Internet Usage Manager (IUM)	Centralize usage analysis and billing/charge-back.

HP OpenView Storage Optimizer

The HP OpenView Storage Optimizer provides performance monitoring of all components of the storage area network (SAN), including hosts, storage devices, and infrastructure.

Storage Optimizer is part of the Storage Area Manager suite and is available as an individual product.

Key features and benefits include:

Table 78: HP OpenView Storage Optimizer features and benefits

Feature	Benefit
Monitors key metrics of SAN performance, drilling down to individual node level (host, switch or storage array)	Ensures service levels and availability of business processes are met.
Management of storage resources via a single centralized station	Centralize storage management on a SAN. Company-wide cost savings in a tight job market.
Multiple data presentation formats-GUIs, CLUIs and interval summation reports	Receive information in a user-preferred format.
Automated baselining and over-baseline notification for performance metrics	High availability and reliability. Proactively ensure that SAN SLAs are met.
Graph historical performance metrics and extrapolate historical data	Proactively identify SAN infrastructure trends/anomalies. Evaluate the impact of upgrades. Identify future performance demands. Improve system efficiency.

SAN security

15

Information security is a fundamental issue that must be dealt with while managing any data center. HP understands the importance and complexity of establishing and maintaining a secure information storage environment. HP storage products are designed to make it easy to protect the availability, integrity, and confidentiality of the customer data that they hold.

HP is working with other storage vendors in the Storage Networking Industry Association to develop enhanced SAN security technology. See: http://www.snia.org/tech_activities/storage_security for additional information.

HP is also working with the Fibre Channel standards community to develop storage network security protocols. See http://www.t11.org for information on the Fibre Channel Security Protocols (FC-SP) project.

This chapter describes storage aspects of information security in a StorageWorks SAN environment. Major topics covered in this chapter include:

Basic security model

The ideal mass storage system provides fast storage and retrieval of information for a number of servers.

This one line summary leaves unspoken a number of additional expectations: It is expected that data written to the storage system today will be available tomorrow. It is expected that the data will be the same when it's read as it was when it was written. And it's expected that the data is not available to any server or any person not specifically authorized to have access. These three possibilities are covered under the general headings of availability, integrity, and confidentiality.

These additional expectations form the basis for defining the availability and security of the data in the mass storage system. For example, the data should be available even if a hardware or software component in the storage system fails: RAID and remote mirroring technology are methods used to maximize data availability.

Three types of attacks, corresponding to the three aspects of information security, can be made on a computer system. Data can be made unavailable for access. Data can be deleted or modified without permission. Data can be examined without permission. Any computer security system must deal with these types of attacks.[1]

The security of a computer system is the responsibility of a security manager. This person defines the operational rules and procedures that are required to maintain the desired security level. To achieve the desired security level in an HP SAN system, the operational rules and procedures should incorporate the guidelines discussed in this chapter.

The basic approach to making a system secure is to define one or more security domains. A security domain is a logical grouping of related components in the storage system, along with a set of rules that specify the amount of communication that is allowed between the components. Devices such as servers and storage systems that are within a given security domain are allowed to communicate with each other. The security manager defines the communication–if any–that is allowed between domains. The security system works by controlling every possible communication path between the security domains, so that data cannot be moved between domains without authorization.

The boundaries of the security domains are barriers that control access to the components. The boundaries also control communication between domains through the network or storage bus connections. Any potential path between security domains must be reviewed to make sure that only approved access is permitted. This can be an extremely complex undertaking.

1. An excellent introduction to computer security may be found in *Computer Security Basics*, by Deborah Russell and G.T. Gangemi Sr, published by O'Reilly. A more detailed discussion of network security methods and protocols my be found in *Network Security Essentials*, by William Stallings, published by Prentice-Hall.

Summary of SAN security practices

HP StorageWorks SAN hardware and software components incorporate features that can be used to implement a secure data storage system. The following table shows the appropriate use of these security features in various environments. The Enterprise Storage System environment is a typical mid-sized to large IT installation used in a business. The Service Provider Storage System environment is a large installation where several customers share a single IT infrastructure. These environments are discussed in more detail in later sections of this chapter.

Table 79: How to use SAN security features

SAN Storage Security Feature	Enterprise Storage System	Service Provider Storage System
Physical security of SAN environment.	Suggested. All personnel are employees, but it is always better to keep sensitive systems away from informal access.	Essential. Personnel are competitors, so the systems must be kept in a secure environment.
Use of zones.	Optional. Use port or WWN zoning as required to manage Operating System conflicts.	Optional. Use port zoning as required to manage Operating System conflicts.
Use of Selective Storage Presentation (SSP).	Essential. Use as required to manage access to data.	Essential. Use as required to manage access to data.
Controlled access to storage system management using serial line interface.	Suggested. Limit physical access to machine room.	Optional. Storage systems are physically secure in this environment.
Controlled access to storage system management using in-band interface.	Optional.	Optional.
Restricted use of multiple switches.	Optional. No additional risk is added.	Optional. No additional risk is added.
Restricted use of multiple storage systems.	Optional.	Essential. Each customer must be located on a different storage controller pair.
Restricted use of Storage Management Appliance.	Optional. Appliance applications are password protected.	Recommended. Appliance applications are password protected, but a shared infrastructure is sensitive to competing interests.
Use of logical unit visibility control on Modular Data Router tape controller.	Essential. Use as required to manage access to data.	Essential. Use as required to manage access to data.
Event logging enabled.	Essential. Needed to track possible intrusion attempts.	Essential. Needed to track possible intrusion attempts.

Data path and management path security

HP divides the responsibility for SAN security into two parts. Data Path Security refers to the protection of the communication path used to move user data through the SAN. Management Path Security refers to the protection of the communication path used to move management information through the SAN.

This is a functional distinction, because in some cases the same physical connection is used for both user data and for management information. For example, the Storage Management Appliance communicates with an HSV storage controller using the Fibre Channel connection that is also used to send user I/O traffic.

Table 80 shows the Data Path Security and the Management Path Security features available in HP SAN products.

Table 80: HP SAN products data path and management path security features

Data Path Security	Management Path Security
Selective Storage Presentation	Passwords on user interfaces
Zoning by port and WWN	Security of sign-on to Element Manager
Port binding	Secure communication between storage management appliance and storage array
Fabric binding	Control of IP access to device management ports
Switch binding	
Communication packet encryption (future)	
Data encryption on storage media (future)	

The HP storage security model is implemented as three distinct areas. The overall security of the storage system is an integral part of the total solution security, and is deployed within the context of a comprehensive understanding of the system, developed and delivered by HP Professional Services. The software components of the storage system provide Management Path Security by controlling operator access rights and by securing the SAN management communication paths. The hardware components of the SAN provide Data Path Security by controlling storage array access and by governing the SAN fabric configuration control mechanisms.

Personnel and operating practices

The most important security feature in any environment is the attitude and operating practises of the personnel. The system managers and operators must have a positive view of security, and must be able to balance the need for data security with the need for reasonable user access.

Responsibility for maintaining SAN security should be assigned to a security manager, and this person should have the authority to enforce reasonable security guidelines. The security manager is responsible for making the trade-off between required user access capability and access restrictions required to maintain the required level of security.

HP professional services can assist in developing a suitable operating protocol for your SAN environment. The HP Security Services Portfolio includes a comprehensive end-to-end lifecycle range of services for designing, building, integrating, managing, and evolving sound solutions. The Security Healthcheck Services provide quick, comprehensive security vulnerability and risk assessments of your installation, including the storage systems and related storage network infrastructure. See http://www.hp.com/hps/security for additional information on HP Security Services.

Professional services for SAN security

The establishment of a comprehensive security environment for a large computer system is a complex task. In addition, a failure or breach of the security system may result in the loss of important business information. For these reasons, HP requires that licensed security options must be installed as part of a professional services contract.

The HP SAN Security Services product includes a security review and planning feature and an ongoing security auditing process. The initial security review and planning steps are done before the security products are installed, and result in a report summarizing the security environment and requirements of the proposed installation. The security products are installed when required. The ongoing audit process includes a periodic review of the environment, management and operational practises, and any security logs or other data that is recorded by the system.

Risk of security problems is minimized by using the HP professional services for SAN security.

Security features of HP StorageWorks SAN components

The components of an HP StorageWorks SAN are shown in Figure 73.

Hardware components include the Host Bus Adapter (HBA) residing in each Application Server, the Fibre Channel Switches that make up the SAN fabric (or fabrics in a multiple fabric SAN), the Disk Storage Systems (including their RAID controllers, cache memory, disks, and related management components), the Tape Storage Systems (including their Network Storage Router gateways), the SAN Management Server, the Storage Management Appliance, and various communication cables.

Software components include the server operating systems, the StorageWorks Command Console software, the SAN Management Software, Web Browser and Terminal Emulator interfaces to the Fibre Channel Switch and Storage System management tools, and the MDR management interface.

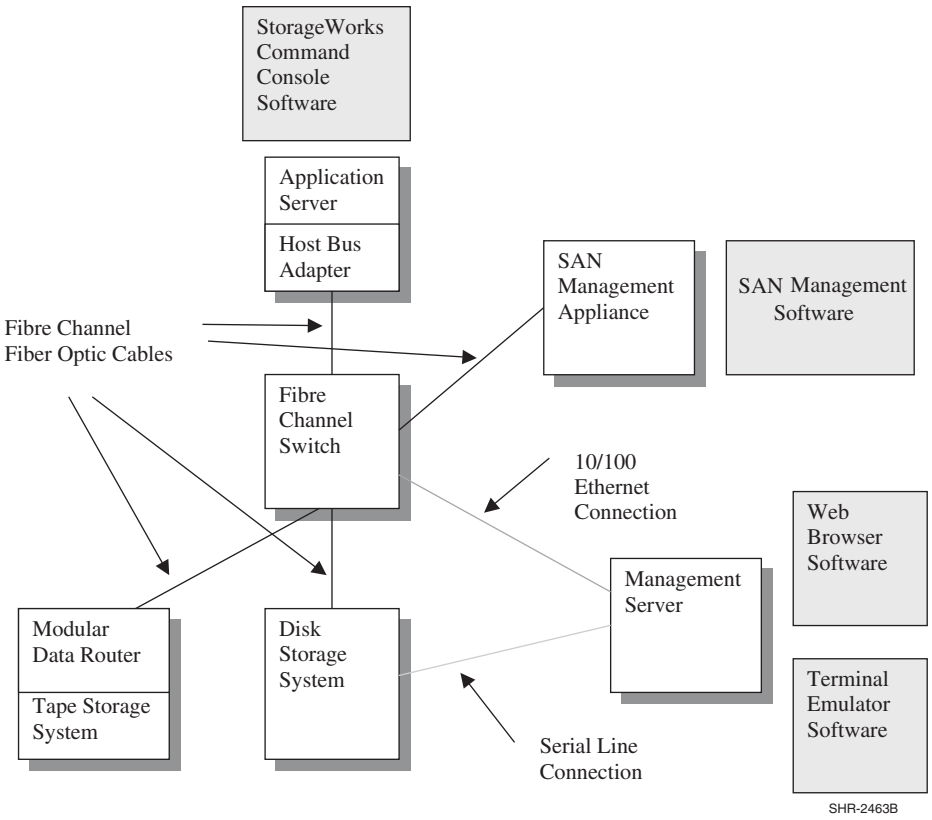

Figure 73: SAN components

The security features of each SAN component are listed below. SAN security is a rapidly developing technology, and the information in this chapter reflects the status of the technology as of the date of publication.

Fibre Channel fiber optic cables

Fiber optic cables used for Fibre Channel communication do not emit electromagnetic radiation. This reduces the risk of security intrusion by means of remote sensing. However, it is not particularly difficult to make a physical tap into an active fiber optic cable. To maximize communication security, the cables should be kept within a secure area.

If a Fibre Channel cable is disconnected, the loss of signal is detected by the connecting devices and is logged in the devices' event logs. Verify that event logging is enabled on all connecting devices that have this feature.

10/100 Ethernet

Because of the difficulty of securing a distributed system, many IP LAN installations suffer from a low overall level of security. The storage security manager should verify that good passwords are in use on all the SAN components that are connected to a LAN, including the application servers, the Storage Management Appliance, the management server, and the Fibre Channel switches.

Serial line

Serial line interfaces are used to connect a terminal (with its associated keyboard and display) to a server or other SAN component. Serial line connections are made using an RS-232 physical interface. The EIA-423 protocol is used, and the connection runs at a low speed (typically 9600 baud). The serial line protocol itself does not have any provision for access security.

The security manager should verify that good passwords are in use on all the SAN components that have serial line connections, or that these connection points are in a secure area.

Host Bus Adapter

The host bus adapter (HBA) is the basic interface between the SAN and each server. The microcode in an HBA can be changed by using a utility program. In the case of Windows NT, a microcode load can be done on an active system, and the server does not need to be re-booted to resume normal I/O activity. A new host bus adapter may be installed in an operational server”

In Fibre Channel, there is no equivalent functionality to the “promiscuous” mode of operation that historically could be used on 10 Mbps CSMA/CD Ethernet networks. The security risk associated with HBAs in a Fibre Channel environment is low because the switches filter all traffic. Only traffic intended for a given server is communicated between the switch and that server's HBAs.

If the operating system driver is changed, then the system must be rebooted. This minimizes the likelihood of undetected changes to driver software.

Fibre Channel switch

Fibre Channel switches are connected together to form a SAN fabric. The switches are the foundation of the SAN system. HP offers three families of switch products, the B-Series Fabric Line, the M-Series Line, and the C-Series Line, each with a unique set of features and capabilities. The security features differ between the three families, as described in the sections below. See the product documentation for additional information that is specific to these products.

Standard security features of M-Series product line switches

The following security features are included with all members of the M-Series product line family of Fibre Channel switches.

Switch zones

The switches in a fabric cooperate to enforce data access zones. Servers are identified either by the switch port to which they are connected, or by their unique World Wide Names (WWN). These two methods are called “port zoning” and “WWN zoning”, respectively.

The advantage of port zoning is that it is easy to configure, while the disadvantage is that if a server is moved from one port to another, the zone configuration must be changed to reflect the new connection topology. The advantage of WWN zoning is that it is independent of port, so servers may be moved from one port to another without changing the zone settings. The disadvantage is that an HBA in a server could, at least theoretically, take on the WWN of another HBA and thus gain unauthorized access to the wrong zone.

The purpose of zones is to manage the interaction of servers in a SAN, preventing interference between the operating system drivers. In heterogeneous configurations the drivers may interfere with each other, and in homogeneous operating system environments the capacity of certain driver data tables may be exceeded. Zoning is used to manage these operational factors. The security manager should verify that event logging is enabled to record unintended and unauthorized changes to the SAN configuration.

Passwords

All user interfaces to switches in the M-Series Line are protected by passwords. HP strongly recommends customers change the passwords on all switches.

Management system communication

The Ethernet connection between the switch and the management station is protected by a secure protocol.

Optional security features of M-Series product line switches

The following security features may be activated on all members of the M-Series product line family of Fibre Channel switches by the use of a license key. This key is supplied in the HP SANtegrity Binding product. See the HP SANtegrity Binding product description for additional information on these features.

Fabric binding

When Fabric Binding is activated, only switches and directors that are identified in the Fabric Membership List are authorized to join the fabric.

Switch binding

When Switch Binding is enabled, only devices that are identified in the Switch Membership List are allowed to connect to the fabric.

Enterprise fabric mode

When Enterprise Fabric Mode is active, the security system automatically activates the following capabilities on all switches in the fabric and does not allow any to be deactivated:

- Fabric Binding (includes Insistent Domain ID)
- Switch Binding
- Domain RSCNs
- Rerouting Delay

Standard security features of B-Series line switches

The following security features are included with all members of the B-Series Line family of Fibre Channel switches.

Switch zones

The switches in a fabric cooperate to enforce data access zones. Servers are identified either by the switch port to which they are connected, or by their WWN. These two methods are called “port zoning” and “WWN zoning”, respectively.

The advantage of port zoning is that it is easy to configure, while the disadvantage is that if a server is moved from one port to another, the zone configuration must be changed to reflect the new connection topology. The advantage of WWN zoning is that it is independent of port, so servers may be moved from one port to another without changing the zone settings. The disadvantage is that an HBA in a server could, at least theoretically, take on the WWN of another HBA and thus gain unauthorized access to the wrong zone.

The purpose of zones is to manage the interaction of servers in a SAN, preventing interference between the operating system drivers. In heterogeneous configurations the drivers may interfere with each other, and in homogeneous operating system environments the capacity of certain driver data tables may be exceeded. Zoning is used to manage these operational factors. The security manager should verify that event logging is enabled to record unintended and unauthorized changes to the SAN configuration.

If more than 64 zones are defined in a single SAN, there may be cases where the port zoning table overflows. In this case the switches revert to WWN zoning. See the user guide for the switch you're using for additional information on this topic.

Passwords

All user interfaces to switches in the B-Series product line are protected by passwords. The default passwords are available to the public, so it is extremely important to change them when the switches are installed.

Optional security features of B-Series product line switches

Enhanced Brocade Fabric Manager 4.0

Brocade Fabric Manager provides a comprehensive SAN configuration control utility. Fabric Manager enables customers to configure and manage multiple B-Series product line switches from a single console. Features available with Fabric Manager 4.0 include SAN-at-a-Glance overviews with a topology map, call home support to send automatic notifications of system failure, enable remote support and isolate faults, and enhanced port management support, including port grouping.

See www.brocade.com for additional information on Brocade Fabric Manager 4.0.

Secure Fabric OS

HP StorageWorks Secure Fabric OS protects your SAN by using the strongest, enterprise-class security methods available, including digital certificates and digital signatures, multiple levels of password protection, strong password encryption, and Public Key Infrastructure (PKI)-based authentication, and 128-bit encryption of the switch's private key used for digital signatures.

Features include Fabric Configuration Servers ("trusted" switches), Management Access Controls, Device Connection Controls (Access Control Lists), Switch Connection Controls, and Secure Management Communications. The trusted switches provide a central location for controlling SAN security. Device ACLs and Switch Connection Controls prevent unauthorized devices and switches from connecting to the secure fabric. All inter-switch management communication as well as communication to the management console is secured using encrypted passwords.

For additional information on configuring your HP StorageWorks SAN using the HP StorageWorks Secure Fabric OS, see:

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

To access the technical documentation at this site:

- Locate the **Networked Storage** section of the web page.
- Under Networked Storage, go to the **by type** subsection.
- Click **SAN Infrastructure**. The SAN Infrastructure page appears.
- Locate the **fibre channel switches** section.
- Go to the **infrastructure** subsection.

Storage system

Products in the HP HSG80-based and Enterprise Virtual Array series of storage systems incorporate security controls on all the interfaces to the storage system.

Each storage system consists of a pair of HSG80 or HSV storage controllers, along with assorted supporting hardware.[2] The storage system is connected to one or more servers, and presents logical disks to those servers. Each logical disk has a logical unit number (LUN).

The Selective Storage Presentation (SSP) feature allows visibility of logical units to be restricted to a subset of the servers connected to the storage system.

Physical access control

The storage system is typically housed in a standard HP rack with locking front and rear doors. The locks for these cabinets all use the same key, so the security aspect of the locks is only sufficient to deter the most casual intrusion. The locks can be changed to provide additional physical security if desired.

Controller management

Basic control of the storage system is performed using various buttons and lights on the front and rear panels of the RAID controller shelf. These controls allow the controllers to be halted or restarted. The HSG80 controller microcode is stored on PCMCIA cards that are inserted into these panels. Physical access controls to the controller shelf must be maintained to prevent unauthorized manipulation of these controls and to prevent unauthorized replacement of the controller microcode.

One option for initial setup of the storage system as well as for ongoing operation is to use a serial line connection to each HSG80 RAID controller. This connection is typically made between a controller and a terminal emulator program running on a nearby computer. All storage system management operations can be done using this interface. Physical access to the controller shelf must be maintained to avoid unauthorized use of this interface.

2. See the HSG80 and HSV controller documentation for a complete description of the features of the HP family of storage systems.

Another option for the initial setup and ongoing operation of the storage system is to use the in-band Fibre Channel management system. This system sends SCSI commands to logical units on the storage system to control the logical unit definitions and the SSP settings. A server may send these commands to any logical unit to which SSP allows it access.

Data access control

The Selective Storage Presentation feature of the storage system is the method used to control access to user data. Access is allowed to each logical unit by one or more servers.

The SSP settings may be controlled by any server having access to any logical unit on the storage system. This includes the SWCC agent, the SSSU tool, and the Storage Management Appliance, and could include a purpose-built intrusion application running on a server connected to the SAN. If a computing environment has multiple security domains, then the domains must not coexist on a single storage system.

For example, consider the configuration shown in the following figure. Server A and Server B have access to logical unit D and logical unit E respectively. Server A and logical unit D are in one security domain, and Server B and logical unit E are in a separate security domain. Because both have access to Storage System C, then Server A may change the SSP settings to prevent Server B from accessing any logical units on the storage system.

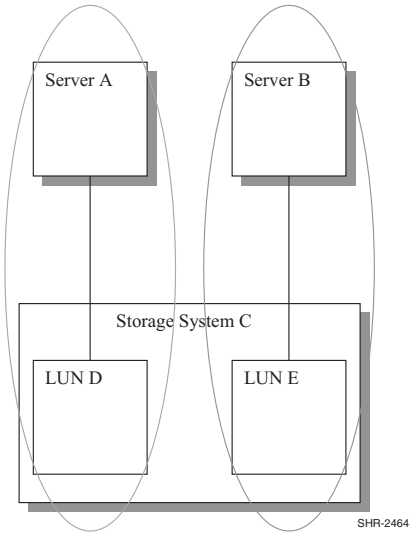

Figure 74: Multiple security domains on one storage system

LUN security in the XP based disk storage systems

Secure Manager XP provides security at LUN level, which is not available through switch zoning. LUN security can be enabled on a per port basis and allows permitted WWNs of hosts to be added to host group or groups on the selected ports.

LUN security in the VA-based disk storage systems

Secure Manager VA is an optional software for the VA arrays that provides extra security at LUN level. This is accomplished by mapping LUNs against pre-configured host HBA WWNs, thus creating a secure host table internally. The array will not permit access to the LUNs if the WWN of the host HBA is not present in the table. The total number of hosts or HBAs allowed per VA controller varies depending on the model. See the respective user manuals for details.

EVA management access control

A management agent can control many storage systems, and many management agents can control a storage system. Without password protection, any management agent on the fabric can access any storage system on the fabric. A password is used to increase the security within your storage subsystem. Specifically, password protection:

- Allows a management agent to control only certain storage systems.
- Allows only certain management agents to control a storage system.

All management functions for Enterprise Virtual Array storage subsystems are done via the Storage Management Appliance. Two levels of security are implemented for the Enterprise Virtual Array to control unauthorized access to the storage subsystem.

The first level controls access to the MA itself. User access to the MA is controlled by a username and password method that uses the WEBM security model. Without the correct username and password, an unauthorized user cannot access the MA.

Secondly, the storage subsystem has an optional password protection to control which MA can manage which storage subsystems. The password is established by entering a password into the operator control panel (OCP) of one of the controllers. Use Command View for HSV Management Agent options to enter the password used by that MA to access particular Storage Subsystems.

In addition to the optional storage subsystem zoning on the fabric, this should prevent someone from putting an unauthorized MA on the fabric and attempting to manage a EVA storage subsystem.

StorageWorks Command Console management software

StorageWorks Command Console (SWCC) is a client-server storage management software product that supports in-band management of HP EVA and HSG80-based storage systems. An agent program runs on a server and communicates with any storage system attached to that server. The SWCC client program runs on a second, remote server to provide the GUI. The two servers communicate by using a TCP/IP connection between the two servers.

The Command Scripter tool also uses the SWCC agent to communicate with storage systems.

User access to the SWCC agent is controlled by a username and password. Any SWCC client accessing the agent to perform management tasks will be asked for this password. The communications between the management station and the host servers connected to the storage controllers is protected by single-use key encryption. In addition, remote configuration can be optionally disabled.

Communication between the agent and the controller is done by using SCSI commands on the Fibre Channel connection between the server and the controller. The agent communicates with a logical unit on the controller.

Storage System Scripting Utility

Storage System Scripting Utility (SSSU) is a character cell interface that allows a user to configure and control Storage Controllers generically on a storage area network (SAN). Simple or initial configuration requests can be handled easily and expediently through this simple character cell interface, such as the initial creation of LUNs presented to the host. SSSU meets this requirement with an interface that allows the user to issue simple, terse commands.

SSSU uses the Storage Management Appliance to communicate with EVA storage systems. User access to the MA is controlled by the user name and password method that uses the WEBM security model.

Storage Management Appliance

HP offers an optional integrated SAN management system that uses an appliance connected to the Fibre Channel fabric. The Storage Management Appliance hosts web-based Open SAN Storage Management software. This software provides a wide variety of management tools.

Access to the Open SAN Management applications is controlled by a user name and password method that uses the WEBM security model.

Storage security in an enterprise environment

In a business enterprise, computer systems may be shared between two or more departments. The systems are managed and operated by an Information Systems organization, which has enterprise-wide responsibility for the computing environment. All the people in the enterprise work towards a common business goal, but the day-to-day interests of the departments may vary widely depending on the business climate, time of year, or product development issues. Each department has specific computing requirements that must be met by the IS organization.

There may be wide differences in the need for data security. For example, a typical accounting department has strict security guidelines, while the marketing department may be willing to tolerate more risk.

The IS organization may try to achieve efficiency by placing the computer equipment in a single central location. A considerable amount of computer and storage hardware is required for an enterprise of moderate size. This discussion assumes that the storage for all the departments is located in a single SAN storage system. Servers are distributed throughout the facility.

The IS organization must implement a computing system that meets the security and capacity requirements of all the departments to which it provides service, and the IS security manager must implement a security plan that is suitable for the needs of the enterprise.

To meet the security requirements, many security managers specify a centralized machine room located in a secure area. This substantially reduces the security risk for the storage system, because the ordinary users of the system do not have physical access to the machines.

Security expectations

This is an environment with a requirement for a high level of storage system security. Protection is needed against unauthorized, accidental, and malicious data access attempts. The required security level is set by the department with the most strict security needs.

SAN component security attributes

The following features are used to provide security in this environment.

Traditional user account security is in effect in the servers. This protects each user account against accidental access by an unauthorized user. Disk quotas are enabled for each account. This prevents a user from consuming all of the storage capacity allocated to the server.

The HBAs pass user I/O requests to a Fibre Channel switch. Communication is done using Fibre Channel fiber optic cables. These cables pass from the servers into the secure area that holds the storage systems.

The SAN switches are shared by all the users and servers in all departments in the system, and are located in the secure area. Configuration management of the switches is done by the system manager using the web management interface. The interface is protected by password to prevent unauthorized changes to the switch configuration.

Data is stored in several StorageWorks storage systems. Access to each logical unit is controlled by the Selective Storage Presentation feature of the array controller.

Response to attacks

Two attack scenarios are possible in this situation. Accidental inappropriate data access requests might be made by any user, and malicious attempts to make an inappropriate data access requests might be made by a user.[3]

Inappropriate read and write requests by system users are routinely handled by the operating system. Disk mounting requires a privileged account, and directories are protected by access control lists. The benign server environment puts little stress on the security capabilities of the storage system.

Because the storage systems are located in a secure area, the risk of inappropriate access to the array controllers is limited. There is some risk that the fiber optic cables might be tapped, but this requires a technical approach that is unlikely in this scenario.

Denial of service attacks initiated by system users, whether accidental or purposeful, are not fully protected against. A user could write a program that issues a useless but very high I/O load and consume most of the I/O operation capability of the storage system.

Checklist

For a SAN storage system that requires a moderate level of security, and where the storage systems and Fibre Channel switches are located in a secure area, the following steps are required.

- Good employment practices to minimize malicious attacks.
- Computer system security awareness training for all personnel.
- Routine user account management at the server.
- Disk quotas enabled for all users.
- Locate storage systems and Fibre Channel switches in a secure area.
- Passwords enabled on all switch configuration ports.
- Selective Storage Presentation for all logical units.
- Disable SES management interface to Fibre Channel switches.
- Routine periodic security audits.

The HP StorageWorks Secure Fabric OS is recommended for SANs based on B-Series product line switches.

3. We've ruled out serious attempts to break into the storage system, but unsophisticated attempts to read someone else's data are possible in any computer system environment.

Storage security in a service provider environment

Some organizations provide computing services to their customers on a lease or contract basis. The services may include general-purpose office applications such as Microsoft Exchange or file and print services, or they may be specialized. One example of the latter is the Storage Service Provider, which provides storage capacity to some other organization. In all service provider situations, the service provider is the HP customer, and the service provider has second level customers of its own who purchase the service.

These second level customers are the users of the systems.

These users may be competitors of each other, and it is essential that they be protected against security breaches—accidental or intentional—by other users in the computer system. The security plan must take into account the possibility of aggressive attacks.[4] This is probably the most difficult environment for a storage system security manager.

Physical access to the storage system is controlled by placing it in a secure area. The servers are in separate secure areas, segregated by user so that each user has a unique secure server area.

Security expectations

The requirement is for high security. Each user wants a separate security domain because there is no trust between competitors. Protection against accidental or intentional unauthorized access to data must be provided, and protection against unauthorized changes to the configuration of the storage system is also required. Sophisticated attacks are not expected, but intentional attacks may occur.

At the same time, services providers are very sensitive to cost. There is a desire to share equipment between users to minimize hardware and management cost. This must be balanced against the security requirements.

SAN component security attributes

The following features are used to provide security in this environment.

Traditional account security is in effect in the servers. This protects each user from accidental, unauthorized access. Disk quotas are enabled for each account. This prevents I/O from one account from consuming all of the storage capacity allocated to the server.

If one user attempts an intentional attack on another user's data, it may be expected that this would be done from a privileged account on a server. Account security does not protect against this sort of attack, but the exposure from a privileged account is to the data of other accounts on that system, not other users—because they are on their own servers.

The HBAs pass I/O requests to a Fibre Channel SAN switch. Communication is done using fiber optic cables. These cables pass from the servers into the secure area that holds the storage systems.

The Fibre Channel switches are shared by all of the service provider's customers, and are located in the secure area. Configuration management of the switches is done by the service provider's system manager using the serial line interface.

Data is stored in several StorageWorks storage systems. Access to each logical unit is controlled by the Selective Storage Presentation feature of the array controller.

4. Denial of Service attacks are not considered to be a problem, because the comparatively small number of users on a SAN makes it easy to identify and eliminate this sort of aggressor.

A user may attempt to access a competitor's data. To protect against this possibility, it is important to provide a separate storage system for each of the service provider's customers. While the risk associated with sharing a single storage controller between customers is small,[5] distributing them onto private storage controllers eliminates the risk.

Response to attacks

Two attack scenarios are possible in this situation. Accidental inappropriate data access requests might be made by any user, and malicious attempts to make an inappropriate data access requests might be made by a user.

Inappropriate read and write requests by system users are routinely handled by the operating system. Disk mounting requires a privileged account, and directories are protected by access control lists.

The storage systems are located in a secure area, but there is some risk of intentional attacks. This is prevented by providing a separate storage controller for each user.

There is some risk that the fiber optic cables might be tapped. While this risk is minimal, the Fibre Channel switch logs must be examined regularly and the configuration change alarms on the switches enabled. These will notify the security manager if this sort of activity occurs.

Depending on the service provider environment, it may be possible for a sophisticated attack on the SAN to take place. This could involve equipment such as frame grabbers or phantom switches. It is extremely difficult to protect against a sophisticated attack against any network system, and Fibre Channel is inherently exposed because the data is sent as clear text. If this level of risk is expected, see the following section.

Denial of service attacks initiated by system users, whether accidental or purposeful, are not fully protected against. A user could write a program that issues a useless but very high I/O load and consume most of the I/O operation capability of the storage system.

Checklist

For a SAN storage system that requires a high level of security, and where the storage systems and Fibre Channel switches are located in a secure area, the following steps are required.

- Good employment practices to minimize malicious attacks.
- Computer system security awareness training for all personnel.
- Routine user account management at the server.
- Disk quotas enabled for all users.
- Locate storage systems and Fibre Channel switches in a secure area.
- Passwords enabled on all switch configuration ports.
- Selective Storage Presentation for all logical units.
- Disable SES management interface to Fibre Channel switches.
- Disable SNMP management interface to Fibre Channel switches.
- Disable web browser management interface to Fibre Channel switches.
- Each user (that is, each customer of the service provider) must have a separate array controller.
- Routine periodic security audits.

5. It requires special knowledge and equipment to successfully complete an unauthorized access to data on an array controller.

The HP StorageWorks Secure Fabric OS is strongly recommended in B-Series product line SANs used in service provider environments. The enhanced security provided by this product eliminates the risk associated with having ports from a single SAN exposed to multiple second level customers.

Storage security in a secure environment

Some system environments require extremely high levels of security. These are cases of national security or where the data is so sensitive that the owner is willing to make substantial functionality trade-offs to maintain the desired security level. These systems are safe in the face of the worst cases of overt attempts to break into the system by any means possible.

Security expectations

It is expected that the system will have no exposure to security intrusions. This corresponds to the highest levels of information security.[6] Network systems generally are not able to be audited for compliance with the highest levels of security, because network software is too complex for a comprehensive evaluation. To obtain the highest possible levels of information security, the entire system must be enclosed in a secure environment.

SAN component security attributes

To provide security in this environment, the system is enclosed in a secure area.

Checklist

For a SAN storage system that requires the highest level of security, enclose the entire system in a secure area.

- Perform routine periodic security audits.
- Follow other appropriate actions based on the required system security level.
- Place machines in a secure area.

6. Computer system security ratings are set by NIST/NSA in the US and ITSEC in Europe, and “Common Criteria” is a newly-adopted US standard. Within these classifying bodies, a product can be evaluated at various security levels. Most operating systems are classified at ITSEC's E3 rating or Common Criteria's CAPP protection profile EAL4. The OpenVMS SEVMS product has a E3/B1 rating. Tru64 UNIX has an E2/C2 rating. See http://csrc.nist.gov/cc/.

Continuous Access Storage Appliance

16

The HP OpenView Continuous Access Storage Appliance (CASA) solves a wide range of problems that may be encountered in enterprise storage environments. This chapter provides an overview of CASA as well as information about how to integrate CASA solutions into general HP StorageWorks Fibre Channel SAN installations. The following topics are discussed in this chapter:

Note: The information in this chapter is specific to version 5.6.1 of CASA.

Overview of CASA

CASA provides data replication on SANs consisting of heterogeneous mixes of servers and RAID array storage devices. By making all available storage capacity accessible by all servers—with appropriate access controls as required—CASA helps you optimize the use of the server and storage systems in your installation. Data may be placed on the storage device that makes the most sense, regardless of server driver, host bus adapter, or operating system.

CASA supports data mirrors and point-in-time snapshots. These may be done between heterogeneous storage devices. By providing the opportunity for flexible data placement, CASA allows you to distribute redundant copies of data on the storage device most appropriate for each specific type of copy.

CASA can be used to migrate data between heterogeneous storage devices. By removing limits to where data is stored, CASA facilitates the retirement of legacy equipment and the addition of new equipment. If your data availability specifications change, CASA helps you adapt existing data to different availability configurations. As your data center requirements change to meet new business conditions, CASA provides the adaptable data placement and migration tools to optimize your new SAN configuration.

CASA provides an incremental approach to storage virtualization, because it works in conjunction with traditional SAN solutions. If only a subset of your data requires replication or migration, then adding CASA to your existing SAN won't disturb the rest of the data.

The features and supported configurations described here reflect CASA SANOS software version 5.6.1.

Figure 75 is an overview of a typical CASA configuration.

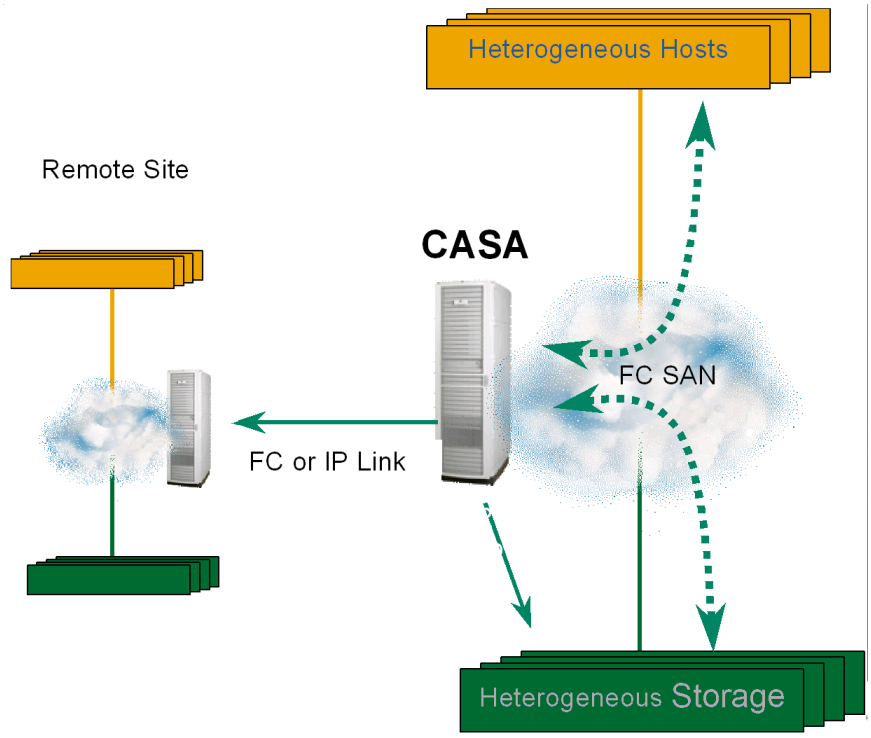

Figure 75: Typical CASA deployment

How CASA works

The Continuous Access Storage Appliance consists of the following components:

- Physical application servers and physical storage arrays
- A CASA appliance with two internal nodes and shared metadata storage
- The CASA utility software
- Management tools used to manage the operating characteristics of the CASA
- Optional racks to hold the appliance, servers, storage, and related equipment

Application servers are connected to CASA using traditional Fibre Channel (FC) Host Bus Adapters (HBAs), cables, and switches. RAID storage arrays are connected using FC cables and switches as required for the specific configuration. The detailed requirements for these components are discussed in this chapter.

Figure 76 shows a schematic of the internal architecture of the CASA. The target ports present virtual disks to the physical servers, while the initiator ports present virtual servers to the physical storage arrays. Shared storage is used within the appliance to store metadata information about the virtual devices. Gigabit Ethernet ports are used to connect the nodes in the appliance.

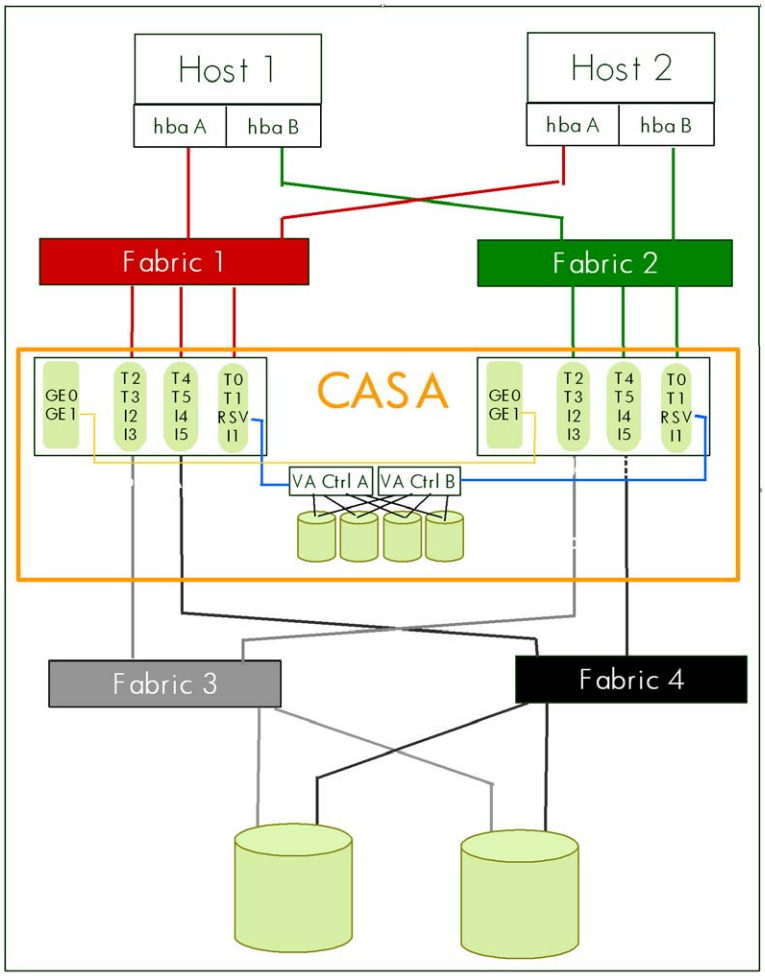

Figure 76: CASA internal architecture

The CASA software implements two kinds of virtual devices:

- The appliance presents virtual target logical units (LUNs) to the application servers.
- The appliance presents virtual initiators to the storage arrays.

The servers "see" the virtual storage devices provided by the appliance, but not the physical storage arrays in your SAN. Similarly, the storage arrays "see" the virtual servers provided by the appliance, but not the physical servers in your SAN.

This isolation of the physical servers from the physical arrays provides the opportunity for tremendous flexibility in the deployment of the servers and storage in the SAN, and is a core element of the storage infrastructure that supports the HP Adaptive Enterprise environment.

Changes to the array configuration can be made without the knowledge of the servers, and changes in the server configuration can be made without the knowledge of the arrays. For example, failures in disk arrays can be made completely transparent to the servers. The failure recovery mechanism in the array works with the virtual server in the appliance to manage the failure, but the appliance masks this activity from the servers on the SAN. Certain appliance failures are visible to the servers—in the same way that array failures are visible to the servers in a traditional SAN—but because all of the storage provided to the server is from the appliance, there is only one storage failure model that the server needs to handle. This means that a server may connect to storage capacity provided by a heterogeneous mix of storage array types, while only implementing a single storage failure handling model. This failure handling model is the one provided by the appliance.

Appliance ports and paths

Each CASA appliance is fully redundant, and includes two peer nodes, each with:

- 6 target ports (host-connect)
- 5 initiator ports (storage-connect)
- 1 dedicated initiator for shared metadata storage
- 3 Gigabit Ethernet ports

The CASA appliance has a total of 12 target ports and 10 initiator ports.

Shared metadata (data about the data) is stored on fully redundant dual controller local storage.

The CASA appliance has redundant paths:

- From hosts to CASA
- From CASA to storage

CASA features

Storage pooling

Under the control of CASA, physical SAN storage is collected into a virtual capacity pool. Unused storage capacity (“stranded capacity”) can be allocated from the virtual capacity pool and then assigned to where it is needed. Mirroring and related replication technology can be applied to the pool in a flexible fashion, without disrupting the servers’ view of the virtual devices. This capability optimizes the utilization of existing storage capacity.

Local data replication

1. Data replication. Data in a given LUN may be mirrored to up to nine other LUNs. This adds to the reliability of the data stored in the physical arrays by protecting against array failure. Mirroring can be used to add flexibility to your backup process by allowing multiple copies of the data to be available at one time.
2. Data snapshot. The Vsnap feature creates a space efficient point-in-time image of a LUN. This capability can be used to create additional static copies (up to nine) of databases or other information. Typically, Vsnap uses only a fraction of the space that would be required for a full mirror of the LUN.
3. For CASA Fibre Channel replication (FCP mirrors), CASAs must be within a campus configuration, typically within 30 km.
4. A cascaded configuration supports 2 CASAs only.

Remote data replication

1. For CASA Synchronous IP replication, CASAs must be within a campus configuration, typically within 40 km.
2. Remote data replication. Mirror copies may be made (up to 3 copies) between multiple CASAs, providing disaster tolerance and the ability to recover quickly from a site failure.
3. Remote cross mirroring between two appliances. The virtual storage capacity pool is distributed across the two sites, and servers at each site may have local and remote mirrors. This provides a fully disaster tolerant configuration that can withstand the failure of either site. For more information on this CASA application, see http://www.hp.com/products1/storage/products/virtualization_appliances/network/sv3000/infolibrary/CASA_CAMs.pdf
4. Synchronous and asynchronous mirroring. Both types of remote mirroring are supported. In the synchronous case, I/O operations issued by a server are not reported as complete until the remotely replicated operation has completed. In the asynchronous case, I/O operations are reported as complete when the local operation completes, which improves performance in cases where distance-related latency is undesirable. Both cases provide guaranteed write ordering technology, so that a disaster recovery operation will have a coherent data image with which to continue operation.
5. “N to 1” replication (for IP replication only.) Up to three sites can be mirrored back to a single central site. This may be used to support centralized backup of multiple sites. This feature allows the use of asynchronous replication to all of the the cascaded sites: No snapshot is required to handle multiple sites.

Figure 77 shows a cascaded configuration.

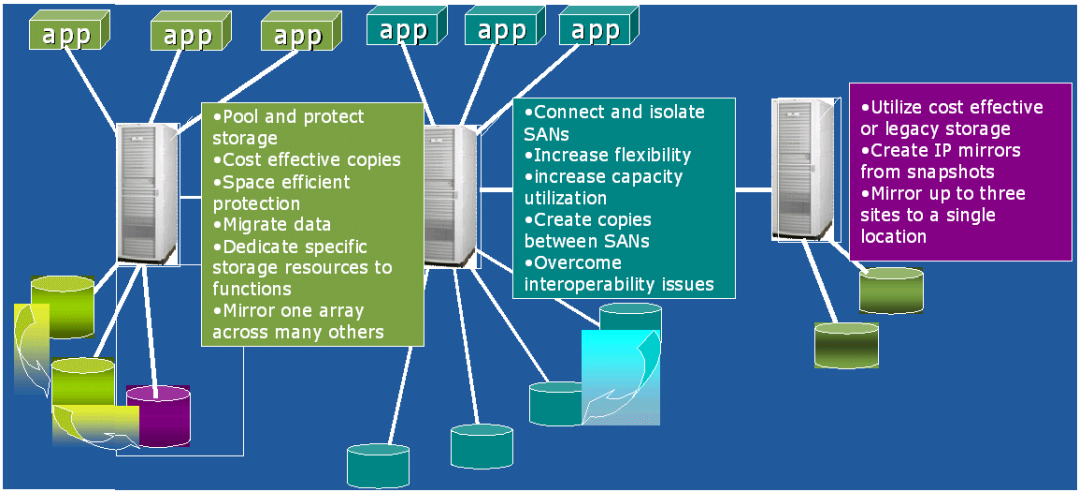

Figure 77: Cascaded CASA configuration with three sites

IP/FCP mirroring

The CASA appliance contains a pair of redundant nodes. These nodes work in tandem as peers to provide a high level of availability to the CASA system. “IP/FCP mirroring” is used to maintain coherence between the two nodes. A journal is maintained on a shared metadata disk array built into the CASA, making the appliance fully redundant. If one node fails, the other node has full access to the shared journals and can fully recover the data.

Mirroring availability continues under a variety of failure conditions:

- Local storage failure
- Remote storage failure
- Either IP link failure
- Either CASA node failure
- In the case of either IP link failure, SCSI requests are routed to the peer node
- If one node loses access to storage, SCSI requests are routed to the peer node

Heterogeneous storage

In all of these cases, heterogeneous mixes of storage arrays are supported. CASA contains HBAs, HBA firmware, driver software, and path failover software appropriate for use with all supported combinations of storage array devices. Because the device characteristics are hidden from the application servers, heterogeneous mixtures of storage array devices may be used without requiring any changes or special configuration options in the application servers.

This powerful feature adds considerable flexibility to the HP StorageWorks SAN. Benefits include:

- Existing arrays may be mixed with new arrays to support data migration.
- Low cost arrays may be mixed with enterprise-class arrays to optimize cost.
- Migration from one type of array to another may be done without impact to the servers.
- Configuration changes required to support new storage requirements may be done without interfering with production work.

Many other important applications for this feature may be imagined without difficulty.

CASA management

CASA environments are managed using the Appliance Management Service (AMS), a centralized web-based user interface. The management service is used to configure and control all aspects of the CASA system. All CASA features are presented in a common fashion to make it easy to control both local and remote devices and the distributed virtual capacity pool. AMS implements a secure management interface for all CASA-related functions.

CASA graphical user interface

The CASA graphical user interface (GUI) supports remote management of CASA appliances. It provides navigation between multiple CASA nodes and appliances without having to login, and provides optional access to the command line interface (CLI) if needed. The GUI incorporates settable user privileges to provide selective access to management operations.

CASA Command Line Interface

The CASA command line interface (CLI) provides remote console based management of CASA appliances. It uses a UNIX shell-like interface that has scripting capability, the ability to process multiple requests from a single file, and flexible navigation between CASA nodes and appliances. The scripting capability allows a CASA to be managed by third party clients.

AMS server

The AMS server software runs on the CASA, and is responsible for handling user management requests from management clients, either the GUI or the CLI.

The GUI or CLI sends an XML request to the management server, which in turn performs the required validation and translates the request to a command that is understood by the appliance. The XML handler is capable of processing management requests for the appliance backend engine and for B-Series switches.

The appliance processes the request, and then sends and appropriate response to the management server, which in turn creates an XML response message and sends it back to the requesting client.

Prior to forwarding any request to the appliance backend engine, the management server first authorizes the request with the security service. See “Security Features” on page 334.

Integration of AMS with OpenView SAM

OpenView Storage Area Manager (OpenView SAM) is used to handle SNMP traps generated by CASA. OpenView SAM 3.0 Suite DPIs are available for integration with Storage Node Manager, Storage Builder, Storage Optimizer, and Storage Accountant. These provide centralized discovery, mapping, performance planning and management, and billing capabilities.

Additional information about CASA management

See the *HP OpenView Continuous Access Storage Appliance System Administrator's Guide* for additional information about managing the CASA system.

Security implications of CASA

Traditional networked storage systems deliver a high level of security. In many cases this security is built into the SAN, because typical SANs are constrained to fairly small physical areas (such as a single machine room, single building, or single campus) and because SAN infrastructure components (such as Fibre Channel switches) incorporate various security control methods. In those cases where a SAN is extended beyond these limits, additional techniques (like encryption of data passed on extended links) must be used to maintain a suitable level of security.

Security features

CASA systems achieve a level of security similar to that of traditional SAN systems by the use of strong access controls and redundancy.

- Passwords protect against intrusions through the management interface.
- Every CASA system is designed using a no-single-point-of-failure topology with redundant components and redundant meta-data storage.
- LUNs are mapped to hosts by unique worldwide name (WWN) to protect data from access by unauthorized servers. New hosts on the network have no access until LUNs are explicitly mapped.
- Hosts can have exclusive storage for independent applications or shared storage to enable failover for clustered applications
- Mapping is network-based, so no host software is required.

The security component provided by the Appliance Management Service (AMS) provides ticket-based authentication and authorization for services on a CASA appliance. This is used by AMS to control access to the CASA appliance, B-Series switches, and its own administrative interface. It also provides an audit trail of authentication and authorization operations, as well as of its own administrative operations.

Security Services provided:

- Identification and Authentication:
  - Challenge-Response Mechanism. Password is never transmitted over the wire.
  - Encryption Based on Shared Knowledge of the Password and User ID.
  - 128 bit Encryption utilizing Blowfish
  - Result: Ticket is Granted
  - Tickets have a tunable timeout—Default is 8 hours.
  - Originator IP Address is contained within the encrypted portion of the Ticket.
  - Ticket must be passed with every request.
- Authorization:
  - Authorization request contains: Ticket, Originating Host, Comma separated list of requested operations
  - The requestor must have privileges required for ALL operations to allow any to be performed.
- User and Role Administration:
  - Add/Mod/Delete/Query Users
  - Add/Mod/Delete/Query Roles

- List Roles for a User
- List Privileges for a User

Architectural Advantages of this approach include:

- XML is a standard language that allows an open, human-readable protocol.
- The security service is usable by clients written in any language that can output XML on a socket connection.
- Implementation in Java enables platform independence.
- XML-based socket level protocol

The strong security features of AMS provide a high level of protection against intrusion through the management interfaces. In addition, if someone were to obtain unauthorized access to the appliance itself, a valid account and password are required to use the GUI or CLI even from the local machine.

See "SAN security" on page 307 for additional information.

Supported systems and software

CASA 5.6.1 supports the following Fibre Channel SAN switches, storage arrays, and server operating systems. Contact your Hewlett-Packard representative for information on specific supported models and version numbers.

Supported Fibre Channel SAN switches

CASA supports the full line of HP StorageWorks SAN switches, as shown in the following tables.

Table 81: HP StorageWorks B-Series product line switches

HP StorageWorks Switch Name		Number of Ports	Firmware Version
HP StorageWorks MSA SAN switch 2/8		8	3.1.1c
HP StorageWorks SAN Switch 2/8 EL, 2/8 power pack		8	
HP StorageWorks SAN Switch 2/16, 2/16 EL, 2/16 power pack		16	
HP StorageWorks SAN Switch 2/32, 2/32 power pack		32	4.1.2b
HP StorageWorks Core Switch 2/64, 2/64 power pack		64 (2 switches per chassis, for a total of 128 ports per chassis)	
HP Switch Name	**Compaq StorageWorks Switch Name**	**Number of Ports**	
HP Brocade 2400 (HP reseller)	Compaq StorageWorks SAN Switch 8	8	2.6.1c
N/A	Compaq StorageWorks SAN Switch 8-EL	8	
HP Brocade 2800 (HP reseller)	Compaq StorageWorks SAN Switch 16	16	
N/A	Compaq StorageWorks SAN Switch 16-EL	16	
HP Surestore FC Switch 6164 (64 ISL Ports)	Compaq StorageWorks SAN Switch Integrated/32 (64 ISL Ports)	32 (counts as 6 switches and 2 hops when applying configuration rules)	
HP Surestore FC Switch 6164 (32 ISL Ports)	Compaq StorageWorks SAN Switch Integrated/64 (32 ISL Ports)	64 (counts as 6 switches and 2 hops when applying configuration rules)	
HP Surestore FC 1Gb/2Gb Entry Switch 8B	N/A	8	3.1.1c
N/A	Compaq StorageWorks SAN Switch 2/8-EL	8	
N/A	Compaq StorageWorks SAN Switch 2/16-EL	16	
HP Surestore FC 1Gb/2Gb Switch 8B	N/A	8	
HP Surestore FC 1Gb/2Gb Switch 16B	Compaq StorageWorks SAN Switch 2/16	16	

Table 82: HP StorageWorks M-Series product line switches

HP StorageWorks Switch Name		Number of Ports	Firmware Version
HP StorageWorks edge switch 2/12		4 to 12	05.05.00-12
HP StorageWorks edge switch 2/16		16	05.01.00-24
HP StorageWorks edge switch 2/24		8 to 24	
HP StorageWorks edge switch 2/32		16 to 32	
HP StorageWorks director 2/64		32 to 64	
HP StorageWorks director 2/140		64 to 140	
HP Switch Name	**Compaq Switch Name**	**Number of Ports**	
N/A	McDATA ES-3016 (Compaq reseller)	16	05.01.00-24
N/A	McDATA ES-3032 (Compaq reseller)	32	
McDATA ED-5000 (McDATA reseller)		32	04.00.00-16
HP Director FC-64	Compaq StorageWorks SAN Director 64	64	05.01.00-24

In addition to the switches listed in Table 81 and Table 82, CASA is also supported with the following Fibre Channel switch models (vendor branded):

Table 83: Brocade and McData Fibre Channel switch support for CASA-only SAN

Switch Brand	Switch Model	Firmware	Hub
Brocade	2400	2.6.1c	No
	2800	2.6.1c	No
	3200	3.1.1c	No
	3800	3.1.1c	No
	3900	4.1.2b	No
	12000	4.1.2b	No
McData	6064 (1 Gb directors)	04.01-02-4	No
	6140 (2 Gb directors)	05.01.00-24	No
	3216	05.01.00-24	No
	3232	05.01.00-24	No

Note: Table 83 lists switch vendor branded switch models supported by CASA only. For general non-CASA SAN configurations, see B-Series switches and fabric rules on page 81, C-Series switches and fabric rules on page 97, and M-Series switches and fabric rules on page 105 for a list of supported HP-branded switch models.

Supported RAID storage arrays

- HP StorageWorks XP48, XP512, XP128, XP1024, XP256
- HP StorageWorks EVA v2, EMA/ESA12000, RA/MA8000
- HP StorageWorks MSA1000
- HP StorageWorks va7400, va7410, va7100, va7110
- EMC Symmetrix 4 and 5
- EMC CLARiiON 4700 and 5700
- Hitachi 9200 and 9900
- Dell Powervault 650F
- IBM Shark ESS 2105 F20

Supported host operating systems

- Windows 2003, 32-bit and 64-bit, requires Secure Path version 4.0c for failover.
- Windows 2000, requires AutoPath version 2.01 or Secure Path version 4.0c for failover.
- Windows NT 4.0, requires AutoPath version 2.01 or Secure Path version 4.0c for failover.
- Solaris 2.6, 7, 8, 9 requires VERITAS DMP for failover.
- HP-UX K, L, R, and V class servers running HP-UX 10.20, 11, 11i v1, 11i v2 requires Secure Path version 3.0c or PVlinks for failover.
- IBM AIX 4.3.3, 5.1 requires AutoPath version 4.02 for failover.
- Red Hat 7.1/Linux Kernel 2.4, Red Hat 8.0/Kernel 2.4.18, Red Hat Advanced Server 2.1/Kernel 2.4.9-e.3 (QLogic) and 2.4.9-e.16 (Emulex), United Linux 1.0 Kernel 2.4.19-64GB-SMP, SuSE Linux 8.0/Kernel 2.4.19-64GB-SMP 32-bit requires native Red Hat failover.
- Novell NetWare 5.1 requires Native NetWare failover.

Configuration rules

CASA supports the full range of HP StorageWorks Fibre Channel SAN configurations as documented in this Guide. The following additional rules apply to all HP StorageWorks SAN installations that include CASA appliances.

Ask your HP representative for additional guidance on configuration rules.

Note: It is required that all host HBA ports are individually zoned to CASA target ports and that all CASA initiator ports are individually zoned to storage target ports.

Number of SAN fabrics

For the purpose of availability, CASA installations normally use four separate Fibre Channel fabrics. Two fabrics are used to provide redundancy for the connection between the application servers and the appliances. Two additional fabrics are used for the connections between the appliances and the storage arrays. For installations where all the storage capacity is to be managed by CASA, this is the preferred configuration because it maximizes the availability of the entire system.

CASA may be used in installations where some of the storage capacity is managed by the CASA and some is directly connected[1] to application servers. In this case two fabrics are required. The failover functionality in the application servers, CASAs, and storage arrays makes this a no-single-point-of-failure configuration; however, there may be additional failover delay associated with the failure of one of the fabrics.

Number of CASAs

CASA is deployed with pairs of nodes in order to provide failover capability. A minimum CASA deployment has two nodes and is described as “one CASA.”

Multiple CASAs may be included in a SAN. Storage capacity is not shared between CASAs, except in those cases where replication is used. There is no specified limit to the number of CASAs that may be deployed in a single SAN, but in practice the connectivity limits of the SAN will restrict the number of CASAs.

Recommended SAN topology

The recommended SAN topology for CASA deployments is core-edge (or director-edge) interconnection. Other topologies may not provide adequate port-to-port bandwidth.

CASA is supported in all HP StorageWorks SAN topologies.

Connection rules

CASA requires a high-performance connection to the SAN for all of its ports. For this reason, the CASA should be connected directly to the core.

Application servers may be connected directly to the core or to edge switches, depending on the application workload.

1. The connection may be a direct physical connection between the application and the storage array, if this is supported for the required server/array combination, or may be through an intermediate SAN. In this discussion “direct connection” includes both possibilities.

Storage arrays may be connected to edge switches or directly to the core, depending on the workload requirements. In many cases the storage array will see a heavy workload and will need to be connected to the core.

Failover software rules

If all of the storage capacity in the SAN is under the management of CASA, the application servers must have failover software appropriate for the CASA. The physical storage devices are consolidated by CASA, so the failover software depends only on the CASA.

The following failover software must be installed on the application servers. Note that this failover management software is used when connecting to the CASA regardless of the storage arrays that are present in the configuration.

- AutoPath VA for Microsoft Windows and IBM AIX
- Veritas DMP for Sun Solaris
- PVLinks for HP-UX
- Native Linux
- Secure Path for Microsoft Windows

In order to handle the event of a path failure between the CASA and a storage array, the following failover software is used in the CASA:

- Native Active-Active (XP, VA, EMC)
- Secure Path (HSG80, HSV110,and MSA)
- ATF (for CLARiiON)

If some of the storage capacity is managed by the CASA and some is directly connected to application servers, then the application servers must have the appropriate failover software for the storage arrays to which they are connected. In these configurations the following issues should be considered.

- May require multiple flavors of failover software on the host (one for CASA storage, one or more for physical storage).
- Requires LUN mapping/masking on storage to allocate CASA LUNs and host accessible LUNs. Ensure that CASA LUNs can only be accessed by CASA by using an appropriate combination of LUN mapping, LUN masking, and zoning.

Example configurations

Three example configurations are shown below. They cover the following cases:

- Single CASA Manages all the Storage Arrays
- Single CASA Manages a Subset of the Available Storage Arrays
- Multiple CASAs Manage the Storage Arrays

Similar configurations may be suitable for customer installations, depending on the specific requirements at hand.

Single CASA manages all the storage arrays

Figure 78 shows a simple CASA configuration with four single-switch SAN fabrics. CASA Node 1 and CASA Node 0 are the redundant pair of nodes and the shared disk that make up "the CASA" in this illustration.

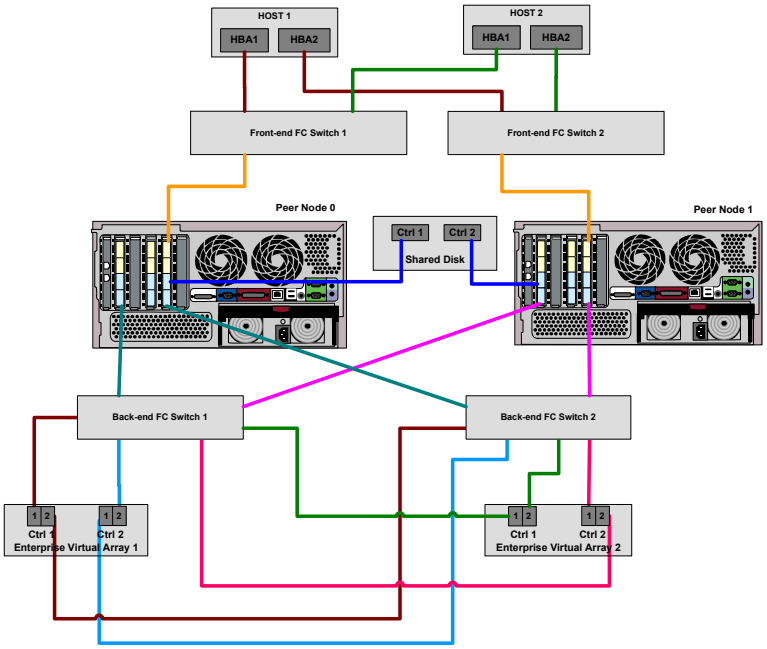

Figure 78: Single CASA configuration

Single CASA manages a subset of the available storage arrays

Figure 79 shows a more complex configuration where one CASA (pair of nodes and a shared disk) manages some of the storage capacity, while other storage capacity is connected directly to the application servers. Additionally, one of the storage devices, Storage Array 1, is configured with LUN masking so that some of its capacity is connected to the servers and some to the CASA. Storage Array 2 is managed by the CASA while Storage Array 3 is connected directly to the application server.

In this case two fabrics are used because a connection between the application server and the storage arrays is needed. Each fabric has a core-edge topology, and the servers are connected to the edge switches as was discussed above. All of the CASA ports are connected directly to core switches.

Note that some of the connections to the storage arrays are not included on this drawing, for the purpose of simplification.

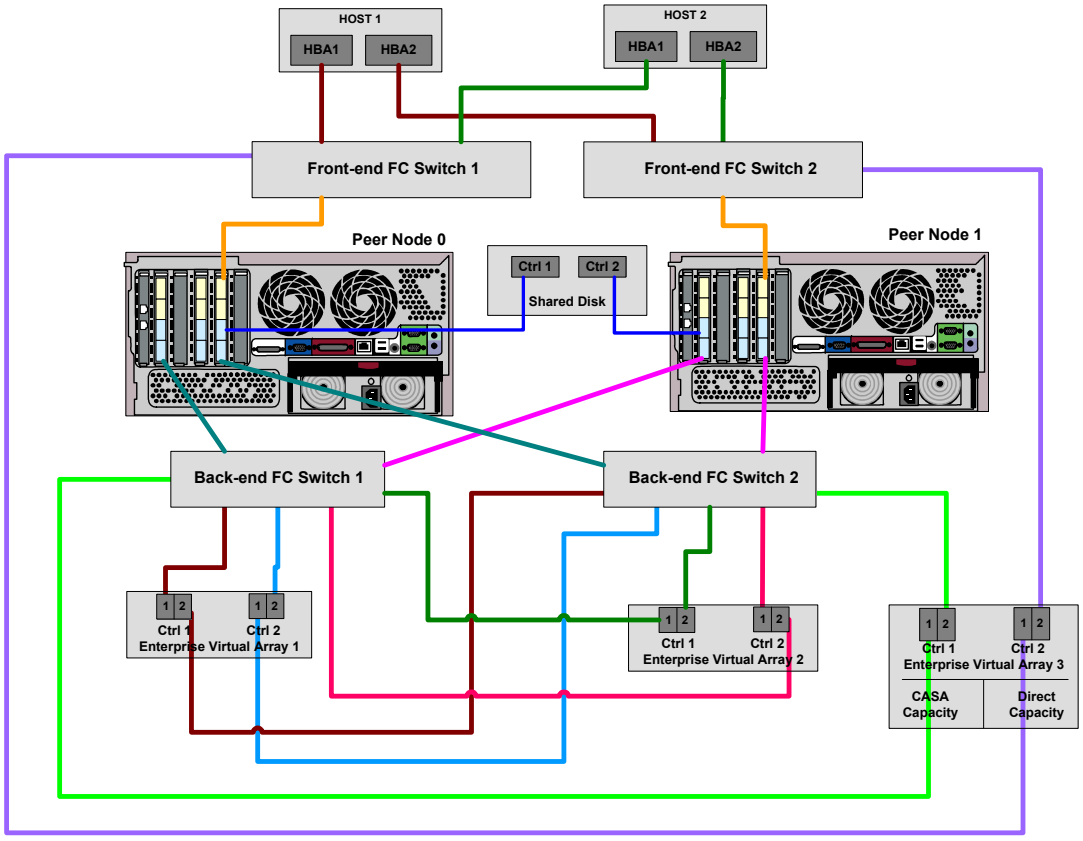

Figure 79: Single CASA mixed with non-CASA storage

Multiple CASAs manage the storage arrays

Figure 80 shows a larger configuration with multiple CASAs and multiple storage arrays.

Each CASA (node pair and a shared disk) controls a specific set of physical LUNs. The LUNs may be located on a single storage array, but in that case LUN masking (for example, Selective Storage Presentation on EVA arrays) must be used to isolate the LUNs.

CASA Nodes C and D are a pair, and have storage capacity assigned to them on storage arrays G and H. CASA Nodes J and K are a pair, and have storage capacity assigned to them on EVA 1. Some of the capacity on EVA 3, and all of the capacity on EVA 2 is assigned for direct access (through the SAN) by the application servers.

If one imagines an even larger configuration, and keeps in mind the requirement that all the CASA ports be connected directly to the fabric core, it is easy to see how the number of CASAs in a single installation is limited only by the number of ports on the fabric core.

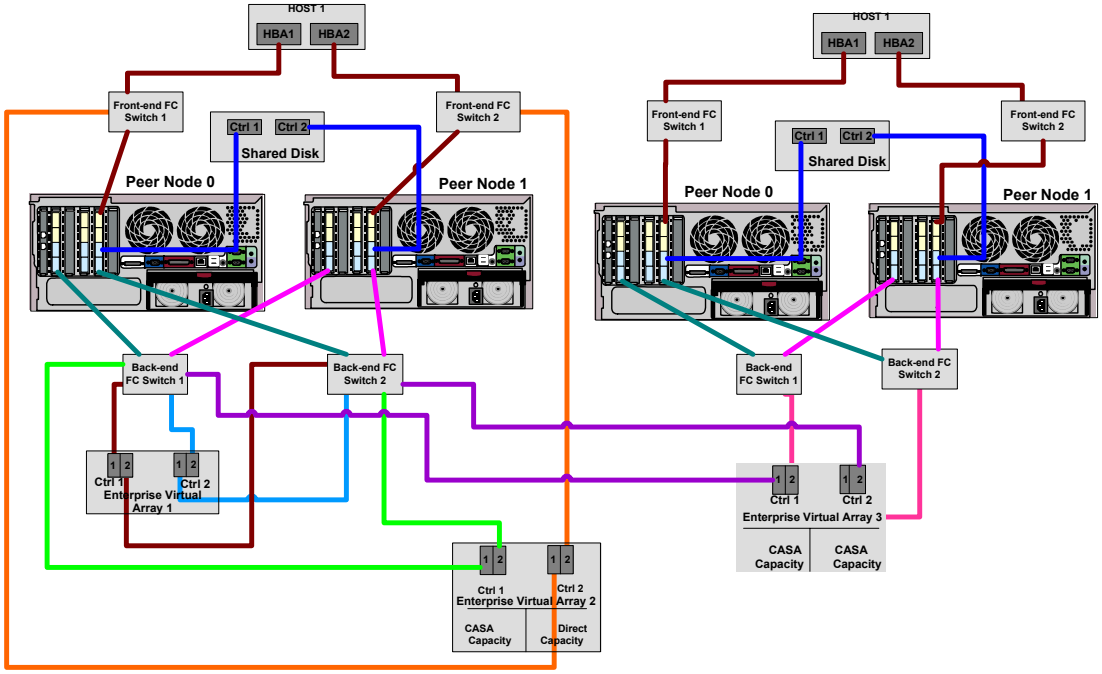

Figure 80: Multiples CASA supporting mix of arrays

CASA services

HP offers three levels of CASA Implementation Services. Simple (one operating system) and moderately complex (three operating systems, 12 hosts, and two Fibre Channel switches) installation services are available at fixed prices. Installation services for complex implementations are quoted after an assessment is performed within the standard HP custom quoting process.

CASA Implementation Services include the following:

- Implementation planning and consulting for CASA.
- Hardware installation and configuration of a CASA solution.
- Basic SANOS configuration including: password management, configuration of shared disk parameters, setting connections between nodes and shared disk, setting management and LAN network properties, and creating peer relationships between the nodes.
- Verify functionality of environment by performing: storage discovery, verification of host registration, creation of sample partitions and expansions, sample LUN mappings, and sample Fibre Channel Mirror and Fibre Channel Mirror configuration.
- Configure virtual LUNs and mappings per customer provided requirements and verify visibility of LUNs to the applicable nodes. The number of nodes verified is bounded by the complexity of the environment (simple, medium, or complex).
- Configuration of WAN IP addresses for each CASA and enablement of WAN interface.
- Configuration of IP Mirror relationship between source and target CASA units.
- Enablement of IP Mirror and creation of sample mirrored volume to verify functionality.
- Preparation of sample source and target LUN for Vsnap snapshot.
- Enablement of Vsnap snapshot, configuration of sample source and target LUNs to verify functionality.

The use of these services is highly recommended, particularly in complex environments.

Additional information sources

See the CASA documentation at:
http://www.hp.com/products1/storage/products/virtualization_appliances/network/sv3000/index.html
for detailed information on supported RAID array storage devices.

Best practices

17

This chapter describes “best practices” for implementing heterogeneous storage area networks (SANs). The information contained in this chapter should be used as a guide for constructing your SAN. Although every attempt has been made to provide a best practice recommendation, some aspects of SAN implementation are a matter of preference. Also, the physical location of servers, storage, computer labs, or building layout and location may dictate particular aspects of your SAN implementation. In part, this is an expected reality and is often easily accommodated, given the inherent flexibility in implementing SANs and Fibre Channel technology.

Rather than just present a list of best practices, the information has been organized into these sections:

Much of what is presented here is the result of the actual experiences of building large SANs within the internal HP engineering environment and at customer sites.

Although this chapter does describe portions of the design process in the planning phase below, it is not meant to convey the entire SAN design process. Contact an HP Enterprise Storage Consultant or the Professional Services organizations for assistance and consultation on designing SANs. HP Storage Services may be contacted through this link:

http://h18005.www1.hp.com/services/storage/index.html

Note: Much of the information in this chapter applies equally to SANs with the B-Series, M-Series, or C-Series Fabric product lines of switches. Any reference to specific switch features pertains only to the B-Series product line.

Planning a SAN

Proper planning considers both present and future requirements. This can be accomplished by over-planning your initial SAN capacity and connectivity requirements to accommodate expected future needs. Whether using an HP standard topology or designing your own topology, select a design that not only offers the best implementation for present usage, but also allows you to expand your SAN over time.

It is important that you allocate an adequate amount of time to plan your SAN. In general, the more detail you can define in the planning phase, the greater the benefit you will realize during the configuration phase.

Consider each of these items during the planning phase:

- **Deployment Strategy**—You can choose to deploy separate smaller SANs or SAN Islands with the idea of increasing capacity by growing the SANs independently or by interconnecting the independent SANs in the future. Smaller SANs are easier to construct, larger SANs offer economies of scale from an operational standpoint, but take longer and are more complex to build.
- **Topology Design**—Consider the topology design compared to the ease of migrating to another, higher capacity design. In most cases this can be accommodated; however, it is always preferable to choose an initial design that can grow, without the need to transition to a different topology.
- **Experience Level**—If you are just beginning deployment of SAN technology, consider starting with a smaller implementation. As you gain experience, deploy larger SANs.
- **SAN Management Strategy**—See *"SAN fabric management tools"* on page 281 and "SAN storage management tools" on page 286 for information about SAN management tools. Define the management strategy and the specific tools that you will utilize to manage your SAN.
- **Technology Advances**—The ideal design considers expected future technological advances, and can easily accommodate the resultant changes. Plan for flexibility in your initial design. Higher port count Fibre Channel switches and faster interconnect speeds are an inevitable evolution of Fibre Channel technology. Ensure that your initial plan addresses and can accommodate expected changes such as these.
- **Document the Design**—This is one of the most important aspects of the planning process. This allows you to fully review and evaluate the design beforehand, evaluate trade-offs, make changes, and effectively communicate specific plans to all groups affected. The other important benefit of documenting your design is that during the later phases of implementation, the documentation serves as the roadmap for the actual implementation.

HP recommends, at a minimum, that you document the following before beginning the actual implementation:

- **Topology Map**—Shows the logical SAN topology and fabric interconnect scheme; conveys the overall design from a strategic standpoint, and can also serve to convey how future growth and technological advances will be accommodated.
- **Configuration Layout**—Shows the physical layout of the entire implementation. More detailed then the topology map, the layout is used during implementation to verify the correct connectivity. This is also extremely helpful if troubleshooting is required in later phases.
- **Storage Map**—Defines the storage system arrangement and configuration in the SAN, and storageset settings such as SSP and RAID levels. This map effectively defines how all of the storage is configured in the SAN.

- **Zoning Map**—Defines the inter-node communication access within the SAN. This map defines which nodes or user ports are allowed to communicate with each other in the SAN.

General planning considerations

It is difficult to make general recommendations about the choice of a SAN topology. There are many variables in large installations that each new configuration requires substantial customized design work. This section provides background information for designs that meet typical large SAN requirements and that are compatible with the future direction of StorageWorks SAN technology.

Advantages of dual fabric SANs

Most large SANs should have two independent fabrics. Each fabric operates independently, and the failure of one fabric does not cause a complete loss of SAN communication.

The reliability of modern electronic hardware is so high that it is difficult to make meaningful predictions of failure rates. Software is used in all components, but it is difficult to estimate the likelihood of software failures. Operator errors are the most likely cause of problems, and the frequency of operator errors depends strongly on operational discipline and employee morale, both of which are very difficult to quantify. All of these potential failure points are minimized by the use of multiple fabrics.

The advantage of dual fabric designs is that they support path failover technology. Path failover is available in most operating systems that are supported in HP SANs. Two HBAs are used in each server, and if the communication path from one HBA to the storage system fails, then the I/O traffic is rerouted through the other HBA.[1]

The two fabrics should be similar in size and topology. This minimizes the risk of asymmetrical performance under certain workloads, and minimizes the total cost of the SAN. Failover software does not support the concept of primary and secondary fabrics.

There is not an automatic increase in cost caused by the use of two separate fabrics. For example, two switches in a single fabric provide about 24 usable ports (depending on the topology). Two separate fabrics, each with a single switch, provide 32 ports at the same cost.

Many of the SAN illustrations in this guide show only a single fabric. This is because most of the design and compatibility requirements apply to each fabric as a complete unit. However, practical SAN designs should have two or more fabrics, each satisfying the configuration rules described in this guide.

Data access patterns

Several HP SAN topologies, are available for a wide range of applications, from small to very large systems. For small installations, the topology may be chosen to maximize connectivity or to minimize cost. SAN performance is not likely to be an issue for a small installation because of the high I/O throughput provided by basic Fibre Channel SAN components.

Large installations must be designed to maximize performance and minimize cost, to support current and future connectivity requirements, and to enable eventual migration to new technologies. Several factors must be taken into consideration to meet these requirements. The factors are categorized into three different data access patterns, one-to-one, many-to-one, and any-to-any.

- **One-to-one**

1. Failover can also be useful in SANs with only one fabric. This protects against HBA failures and certain unlikely problems in array controllers. In general, failover technology should be used in SAN configurations that have two fabrics.

The communication paths within the fabric are used in different ways, depending on the relationship between the servers and the storage systems. In some cases, each specific server stores data on only one or two storage systems. In this case, only a few specific storage systems service all I/O requests from a server, and there is little or no communication between the servers or between the storage systems. A given fabric port sends requests to one (or two) specific fabric ports. This is the traditional server-storage relationship. Many systems still operate this way today.

From the viewpoint of the fabric, the I/O traffic has a "one-to-one" pattern, and the traffic pattern is stable. Each server sends I/Os to a small, specific set of storage systems, and each storage system is associated with only a handful of servers. Only significant changes to the configuration by the system manager will change the connection pattern.

- **Many-to-one**

  Multiple servers accessing data stored in a single centralized pool is another data access pattern. This is a common situation when high performance storage systems have enough capacity to handle a number of servers. In this environment, there is a "many-to-one" I/O traffic pattern on the SAN fabric, and the traffic pattern is stable. Each server sends I/O requests to a small set of storage systems, but each storage system may service a large number of servers. The connection pattern changes only when significant changes to the configuration are made by the system manager.

- **Any-to-any (or many-to-many)**

  In a third case, application servers access data that is distributed across many storage systems. This case may develop in several situations. The latest HP storage arrays may handle a large number of servers. (See the configuration rules in this Guide for detailed information.) A system manager may decide to distribute information over a wide set of storage systems, thus requiring each application to access multiple storage systems. This situation can arise when host-based mirroring is used. Another possibility is that it may be easier to manage the data if it is partitioned and stored on multiple storage systems. For example, Accounting Department data might be stored on one storage system, and Personnel Records data on another. A server requiring access to both data types generates I/O requests to both storage systems.

  Another important situation where data is distributed across a range of storage systems is when the HP VersaStor virtualization technology is used. VersaStor distributes data over all the available storage systems in a SAN.[2] In this case, I/O requests from a given application server are handled by one or more storage systems, in a pattern that is controlled by the virtualization management appliance. In this environment, many servers access many storage systems, which is a "many-to-many" pattern. Management traffic may occur between servers, storage systems, and management appliances.

  From the viewpoint of the SAN fabric, any port may send traffic to any other port, which is an "any-to-any" pattern. Furthermore, because the virtualization manager performs dynamic reallocation of storage system capacity, the traffic patterns vary continuously without manual intervention.

The optimum SAN configuration depends on the I/O traffic, whether it be a one-to-one, many-to-one, or any-to-any pattern.

2. Configuration details are controlled by management options.

Core and edge switch concept

The optimum fabric configuration uses a high-performance “core” surrounded by a number of edge switches. The core provides roughly equal connection performance between any pair of ports. The edge switches provide port aggregation to match the performance requirements of the servers and storage systems to the performance of the core.

Figure 81 shows a large configuration that uses core and edge switches.

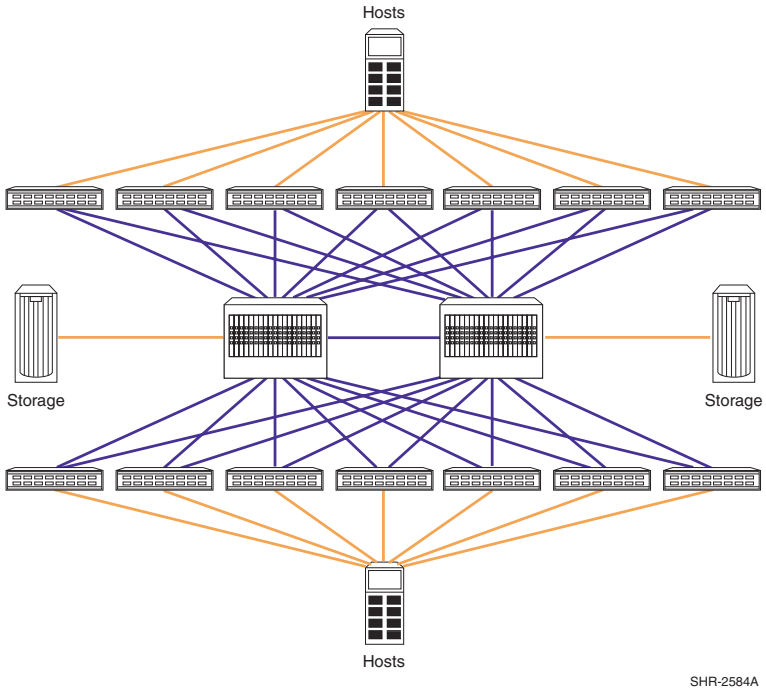

Figure 81: Core switch and edge switch configuration

Fabric core options

The simplest fabric core is a single switch. Fibre Channel switches support simultaneous full bandwidth connections between any combination of port pairs. A single switch fabric core guarantees support for any-to-any traffic.

Any combination of switches has less performance than a single switch, and the difference depends on the fabric topology. The best-performing topology is the “fat tree,” which has enough interswitch links (ISLs) to provide, on average, full bandwidth connections between any combination of port pairs. While it is possible to construct workloads that force traffic contention on the ISLs of a fat tree, which reduces the throughput, fat tree fabric core topologies provide full bandwidth any-to-any communication, on average, for random traffic patterns.

A related topology is the “skinny tree,” which has fewer ISLs and switches. This topology introduces an unavoidable performance limit to the fabric. In many cases, this limit exceeds what is required by the application servers. The process to upgrade a skinny tree topology to a fat tree topology is fairly straightforward, involving the addition of switches and ISLs to the existing tree

Edge switch options

The simplest edge switch is a single switch with one ISL connecting it to the fabric core. Each edge switch provides “User Ports” for connecting servers and storage systems.

The single ISL is a potential bottleneck. All the I/O traffic from the servers or storage systems connected to the edge switch must pass through just one ISL. More ISLs can be provided. Several combinations of ISL and user ports may be used. For example, with 16 port switches, the ISL to user port ratio could be 1:15, 2:14, 3:13, 4:12, etc. Each of these combinations represents a port aggregation ratio. The ratios are 1:15, 1:7, 3:13, 1:3, etc.

The workload of the servers and storage systems attached to an edge switch determines the required port aggregation ratio for the switch. For lightly loaded application servers, a 1:15 port aggregation ratio may be adequate. Heavily loaded servers may require a 1:7 or 1:3 ratio. High-performance servers, such as high-end HP Integrity or HP 9000 Series of servers, may be able to completely "fill up" a Fibre Channel connection. In this case, there is no advantage to using an edge switch, and the server should be connected directly to the fabric core. Storage systems may also be able to support a full bandwidth Fibre Channel connection.

To select the appropriate port aggregation ratio, see the I/O requirements of your applications and servers. This information is available for many situations by using the Active Answers application sizing tools. In other cases, measurements of an existing system may be required to determine the workload.

Designing a subsettable SAN

In many cases, the growth pattern for a storage installation is difficult or impossible to predict. Global economic growth, conditions in a given business market, the growth rate of your company, and internal reorganizations or reallocations of computing resources may all have a significant impact on the requirements that must be met by the SAN.

To accommodate this unpredictable variability, the SAN designer should plan for growth within a predefined design. The initial installation should be a subset of a larger pre-designed configuration.

The "core plus edge switch" approach supports this strategy for SAN design.

When the time comes to expand an existing installation, the system manager can make incremental changes to the configuration rather than a complete reconfiguration of the entire Fibre Channel fabric. Changes to the fabric core are isolated from the edge switches, which minimizes the impact of changes required to support core growth. Changes to a given server's connection to an edge switch are isolated from the core, which minimizes the impact of server-related changes. Furthermore, because two or more fabrics are in use, server I/O traffic may be temporarily forced to a single fabric while the other fabric is undergoing modification.

Start with a single switch core for a moderate-sized initial installation. When needed, the core can be expanded by replacing the switch with one that has more ports, or by reconfiguring the core to a skinny tree or fat tree topology. An existing fat tree core may be expanded by replacing it with a fat tree made up of switches with more ports, or by reconfiguring it to a wider fat tree configuration.

Use a generous estimate of the required I/O performance when selecting edge switches. A port aggregation ratio of 1:7 or 1:3 is adequate for most applications. Increasing bandwidth is a simple, localized modification, if it turns out that more is required.

The initial design should include spare ports on the core to support the future addition of edge switches. For example, consider a configuration that uses 16 port switches, a single switch core, and edge switches with a port aggregation ratio is 1:3. This design supports up to 4 edge switches and 48 user ports. This would be a suitable solution for a system where 36 ports are required now, requiring 3 edge switches. Future growth to 48 ports can be accommodated by adding another edge switch.

SAN design summary of recommendations

Enterprise-level SANs should include the following features:

- Multiple independent fabrics
- "Core plus edge switch" topology
- Appropriate port aggregation ratio, depending on application server requirements
- Appropriate core design, depending on number of ports required
- Subsettable design, with initial installation suitable for current needs

By following these guidelines for SAN planning, your design will support future storage technology and growth in your storage environment.

Configuring a SAN

When you complete the planning phase, you can begin to configure your SAN. As described in the planning phase, it is important that you document the configuration. During the configuration phase, you should record the details of the actual physical configuration.

- **Recording**—As you construct the SAN, record the cable connections and mark this information on the configuration layout diagram. Record the WWN of all nodes and devices and identify where they physically reside. It is recommended that you place a label on each Fibre Channel HBA, with the WWN clearly identified. HP storage systems are prelabeled with this information; however, you may wish to place an additional label on the front of the unit in plain view.
- **Cabling**—Define a system for cable labeling. Even a small SAN can include a high number of fiber optic interconnect cables. Label both ends of each cable with the same unique cable number or color-code scheme. This allows you to quickly identify each cable uniquely. Also consider placing a label at each end of the cables that identifies connection points at both ends, such as TO and FROM. Use label types that are easy to create and read, and ensure they are attached securely to the cable.
- **Switch Ports**—Protect unused or open switch ports with port plugs. Never leave ports exposed.
- **Cable Dressing**—Use care when routing fiber optic cables and ensure that you do not exceed the recommended minimum bend radius. For single-mode and multi-mode fiber cable the minimum bend radius is 25 mm. Where cables are bundled or hanging unsupported, use velcro tie wraps to group and support the cables. Never use plastic tie wraps as they can damage the internal fiber core if over-tightened.
- **Cable Symmetry**—When connecting cables, consider slot/port-numbering symmetry. Be consistent across similar servers with cabling in terms of HBA slot placement and cabling to switches. If configuring with two SAN fabrics and multi-pathing, connect HBA 1 to SAN fabric 1, HBA 2 to SAN fabric 2, etc. Cable symmetry is not a requirement but serves as an aid to troubleshooting.
- **Configure Fibre Channel Switches**—Although all HP Fibre Channel switches are preconfigured, verify that all Fibre Channel switches in the fabric have the same parameter settings and that each has a unique domain ID.

  Label switches using a relevant naming scheme particular to the topology. For example, if implementing a ring topology, label each switch in the ring as Ring1, Ring2. Although not a requirement in all configurations, HP highly recommends that all switches utilize the same switch firmware revision. Different switch code revisions running in the same fabric are supported during a rolling upgrade. This is considered a temporarily acceptable situation for the duration of the code update.
- **Configure Servers**—For each platform or operating system type, utilize the appropriate HP StorageWorks platform kit to ensure that the required server drivers and configuration settings are loaded. Ensure that servers are configured with the proper operating system versions and all required updates.

  Use a numbering scheme for naming multiple servers of the same type, such as NT01 and NT02 for Windows NT servers.
- **Configure Storage**—Use the storage map created in the planning phase to configure each of the storage systems. Verify server-to-storage connectivity and access one server at a time.

When initially defining storagesets, disable all access first, and then enable the desired individual access. For Enterprise/Modular RAID Array storage systems, define connection names to be consistent with zoning alias names. Be consistent with connection names relative to storage port and controller connection. Choose a scheme that is easily understood and quickly conveys the physical connectivity.

- Define Zones. Use the zoning map to configure zones. Consider starting with small zones that allow a smaller logical subset of a larger physical SAN to be tested initially.

  Always save old zoning configurations before and after making any zoning change. If possible, it is recommended that no zoning changes be made when an individual switch normally configured in the fabric is temporarily not available.

  You can zone by operating system or by storage system. Zoning by operating system is useful when the operating systems are accessing storagesets that are localized to specific RAID arrays. For example, NT1, NT2, and NT3 have access to storage on ARRAY1; VMS1, VMS2, and VMS3 have access to storage on ARRAY2.

ZONE NAME	**NT_ZONE**	**VMS_ZONE**
Members	NT1	VMS1
	NT2	VMS2
	NT3	VMS3
	ARRAY1	ARRAY2

  ARRAY1 will only have host connections for the NT1, NT2 and NT3 servers and ARRAY2 will only have host connections for the VMS1, VMS2 and VMS3 servers.

  Zoning by storage system will limit the connections to the G80 to those systems actually having storagesets on them. This is useful when the storagesets for a specific system are on multiple storage systems.

  The following example, shows the addition of three more NT servers and another storage system to the NT zone:

ZONE NAME	**NT_ZONE**	**VMS_ZONE**
Members	NT1	VMS1
	NT2	VMS2
	NT3	VMS3
	ARRAY1	ARRAY2
	NT4	
	NT5	
	NT6	
	ARRAY3	

  Array1 and Array2 will have host connections from all six NT systems. This may not be a problem in a small SAN, but as the SAN grows the connections will increase. Also, you do not know which of the NT servers are accessing storage on ARRAY1, and which ones are accessing storage on ARRAY2.

  The preferred method to zone is a combination of operating system and storage system zoning.

ZONE NAME	**NT_ARRAY1_ZONE**	**NT_ARRAY3_ZONE**	**VMS_ARRAY2_ZONE**
Members	NT1	NT4	VMS1
	NT2	NT5	VMS2
	NT3	NT6	VMS3
	ARRAY1	ARRAY3	ARRAY2

  Zoning this way makes it easier to troubleshoot, especially if servers access storage on multiple arrays.

Due to zoning restrictions, you may need more than one zone for a particular ARRAY. If AIX_ARRAY1 also has IBM AIX servers, you must zone that separately.

AIX_ARRAY1_ZONE1

ARRAY1

AIX_1

AIX_2

Zone and zone alias names

When setting up zoning, use meaningful names for zones and zone aliases and be consistent with the naming convention throughout the fabric.

Servers are identified by the WWN of the HBA. Name these by using the system name and the HBA number. For example, server NT1 with one Fibre Channel HBA would have an alias of NT1_HBA1. Server NT1 with a second HBA would have an alias of NT1_HBA2

RA8000 storage systems in a transparent failover configuration will have two WWNs on the fabric, one for port 1 and one for port 2. Give each RA8000 a unique number. RA8000 number 1 could have aliases of R1_P1 (port 1) and R1_P2 (port 2).

For a multiple-bus failover configuration, the RA8000 will present four WWNS to the fabric. If you have a multipathing NSPOF configuration, two WWNs will be in one fabric; the other two will be in the second fabric. Name the ports using an alias such as R2_A1 (Controller A Port 1), R2_A2 (Controller A Port 2), R2_B1 (Controller B Port1), and R2_B2 (Controller B Port 2).

Ports A1 and B2 will be cabled to the first fabric. Ports A2 and B1 will be cabled to the second fabric. The aliases in fabric 1 will be R1_A1 and R1_B2; the aliases in the second fabric will be R1_A2 and R1_B1. Keep the ports and HBAs the same throughout the setup. For example, always have HBA 1, R1_A1, and R1_B2 in fabric 1 and HBA 2, R1_A2, and R1_B1 in fabric 2.

Using this convention conveys the failover mode that the RA8000 is configured for. Any alias with a P1 or P2 is in transparent mode, any alias with A1, A2, B1, or B2 is in multiple-bus mode.

Define RA8000 host connection names for the adapter WWNs in the same manner as you defined the alias name in the fabric. For example, the fabric alias name for NT1, HBA1 will be NT1_HBA1. The host connections on the RA8000 controller should match this as closely as possible.

For example alias NT1_HBA1 in the fabric would have host connection names on the RA8000 of:

```
NT1-P1  WINNT   THIS   1      081200 OL this  30
HOST_ID=2000-0000-C922-8ADC     ADAPTER_ID=1000-0000-C922-8ADC

NT1-P2  WINNT   OTHER  2      081200 OL other 130
HOST_ID=2000-0000-C922-8ADC    ADAPTER_ID=1000-0000-C922-8ADC
```

Note: While storage system connection names are not case sensitive, switch alias names are. That means that the switch might have an alias name of TRU64_1 and another alias name of Tru64_1 that refer to two different things.

Upgrading a SAN

Upgrading a Fibre Channel switch

See the Installation and Hardware Guide for your switch.

Scaling a SAN

The information in this section applies to all SAN topologies, whether a custom design or HP design.

- Replace 8-port switches with 16-port switches.
- Add switches, up to the limits specified for a single fabric in "Fabric rules for B-Series switches" on page 88, "Fabric rules for C-Series switches" on page 102, and "Fabric rules for M-Series switches" on page 110.
- Add a second fabric as a high-availability no single point of failure (NSPOF) solution.
- Deploy multiple independent SANs.
- Migrate to a different topology.

Scaling specific SAN topologies

The information in this section is specific to the HP-defined topologies. See “General planning considerations” on page 351 for information about preventing fabric segmentation when adding new switches to an existing fabric.

Whenever you are expanding a topology, ensure that the new switch and device connectivity is consistent with the original SAN topology design requirements and goals. Avoid making changes to the topology that may serve to disrupt the original topology design goals. If you need to make topology changes based on a change in data access requirements, consider migrating to a different topology that is better suited to meet these needs. It is important in any expansion that the original data access needs be maintained.

If you have implemented a high availability fabric design (see "SAN fabric topologies" on page 41), it may be possible to expand your SAN in a nondisruptive manner. HP highly recommends, however, as a precaution, that all data be backed up and that I/O activity quiesced when adding new switches to the fabric.

Cascaded Fabric

Expand an existing cascaded fabric by connecting a new switch to an available port on an existing switch. If there are no available ports, remove a device or set of devices from an existing switch, connect the new switch to those ports, and connect the device or devices to the new switch.

Meshed Fabric

Expand an existing meshed fabric by connecting a new switch to available ports on an existing switch. If there are no available ports, remove a device or set of devices from an existing switch, connect the new switch to those ports, and connect the device or devices to the new switch. To maintain the meshed topology, you must ensure that there are multiple paths (ISLs) connecting the new switch to the existing meshed fabric.

Ring Fabric

Expand an existing ring fabric by breaking the ring and inserting another switch into the ring.

Add new switches cascaded off the ring, up to the maximum number of switches supported in a single fabric. When expanding outside the ring, ensure that no two devices that need to communicate are more than seven hops apart.

Tree Backbone Fabric

Add edge switches. Expand an existing Tree Backbone SAN fabric by adding edge switches. Connect these edge switches to available ports on one or two backbone switches.

Add a second backbone switch if your current design only contains one. Connect all of the edge switches to the new backbone switch.

Migrating SAN topologies

This section describes how to convert from one topology type to another if required. HP highly recommends that you thoroughly review your initial design to ensure that it meets your present and future requirements in order to avoid having to modify your initial topology design. There may be situations, based on changes in business requirements, that require you to consider converting to another topology type. For those circumstances, information is provided below that can help you gain an understanding of how the different topologies can be converted.

As described in the planning phase, it is important that the SAN fabric topology be well documented. If you are required to change from one topology type to another, use the existing topology diagrams to determine the most efficient manner in which to modify the topology. Create a new diagram that details the desired final connectivity scheme and use this as a map for the topology migration or conversion.

If you have implemented a high-availability fabric design, depending on the specific cabling changes required, it may be possible to migrate your SAN in a nondisruptive manner. It is highly recommended, however, as a precaution, that all data be backed up and that I/O activity be quiesed when migrating or reconfiguring any portions of the fabric.

If you have implemented a two-fabric, no single point of failure (NSPOF) SAN, you have the ability to failover over all operations to one fabric while you reconfigure the other fabric. This makes it possible to perform a totally nondisruptive topology migration.

- As a general rule, migrations that only require the addition or recabling of ISLs are less disruptive then migrations that require devices be moved from one switch to another. When planning a migration, try to avoid or minimize scenarios that require moving devices from one switch to another.
- **Cascaded to a Meshed Fabric**. Whether you have implemented a linear cascade or branched cascade of switches from one top switch, additional ISLs are required to connect all switches together as required in a mesh fabric design. Proper planning requires that you carefully calculate the number of additional ports that are needed for the additional ISLs. This may require that devices be moved from one switch to another.
- **Cascaded to Ring Fabric**. If you have implemented a linear cascade, connect the last switch in the cascade to the first switch to create a ring fabric. For a branched cascade, extensive ISL recabling may be required.
- **Cascade to Tree Backbone Fabric**. Whether you have implemented a linear cascade or branched cascade, determine which switch or switches will be utilized as the backbone (typically no devices) and which switches will be edge switches. Connect all edge switches to the backbone switch or switches; connect all devices to the edge switches. This conversion is less disruptive if you add new switches to the fabric for the backbone and use all of the existing switches as edge switches. In this case, you can simply connect one end of the existing ISLs to the new backbone switches.
- **Meshed to Ring Fabric**. A meshed fabric can be converted to a ring fabric by simply removing the cross-connected ISLs, leaving the outer connected ISLs connected as a ring. The available ports can be utilized as additional redundant ring ISLs or for additional devices.
- **Meshed to Tree Backbone Fabric**. Determine which switch or switches will be utilized as the backbone (typically no devices) and which switches will be edge switches. Connect all edge switches to the backbone switch or switches; connect all devices to the edge switches. This conversion is less disruptive if you add new switches to the fabric for the backbone and use all of the existing switches as edge switches. In this case, you can simply connect one end of the existing ISLs to the new backbone switches.

- **Ring to Meshed Fabric**. If you have implemented two ISLs between all switches in the ring, move one end from an ISL between any two switches to the appropriate switch based on the final mesh design. Repeat this for all of the second ISLs between any two switches. There may be an optimal place to "break" the ring relative to recabling. Evaluate different scenarios prior to performing the actual conversion.
- **Ring to Tree Backbone Fabric**. Determine which switch or switches will be utilized as the backbone (typically no devices) and which switches will be edge switches. Connect all edge switches to the backbone switch or switches; connect all devices to the edge switches. This conversion is less disruptive if you add new switches to the fabric for the backbone and use all of the existing switches as edge switches. In this case, you can simply connect one end of the existing ISLs to the new backbone switches. It is also less disruptive if you have implemented two ISLs between all switches in the ring in your original design.

Zoning rules and guidelines

Configure the zones based on the zoning map prepared during the SAN planning stage. There are several ways to configure zones.

You must understand the distinction between configuring the zones and zoning enforcement within the switches, and any correlation between the two. This may vary from product to product and hence the following paragraphs describe some general description about each of them, followed by specific details on all HP supported switches.

Note: When enabling a new configuration, it is strongly recommended that the fabric be quiesced. Zone membership should not be changed for devices that are actively performing I/O in a fabric. Once the new zoning is enabled, a state change notification is sent to all the nodes that have registered to receive the state change.

Zoning enforcement

Generally, there are three types of zoning enforcement/authorization techniques in use in Fibre Channel switches: access authentication, discovery authentication, and login authentication.

Access authorization

Access authorization provides frame-level access control in hardware. It verifies SID-DID combination of each frame and allows the frame to be delivered to the destination if that is a valid combination per the zone definition. This is more secure and generally referred to as "hard zoning," and requires hardware resources at the ASIC level to implement.

Discovery authentication

Zoning enforcement or protection of unauthorized access is provided during access to the Name Server directory where the switch or fabric presents only a partial list of devices from the NS directory corresponding to the partition in which the requesting device is part of. This type of zoning enforcement is generally referred to as "soft zoning." While this is secure enough in most of the cases, it is prone to security threats if malicious hosts attempt to access unauthorized devices violating FC-Protocols.

Login authentication

Some switches enforce authentication during FC-protocol login frame level such as PLOGI/ACC/ADISC/PDISC, in addition to providing discovery authentication. For example, if a host sends a PLOGI to a device that is not part of its zone, this frame gets dropped before reaching the destination. This type of enforcement has some additional protection compared to discovery authentication, but still not the same as access authorization at every frame level.

Zoning configuration

Domain ID and port numbers

Zones can be defined using the switch domain ID and port number to uniquely identify zone members. The advantage is ease of configuration and the zoning definition remains intact even if an HBA or target controller is replaced. The disadvantage is limited flexibility to move devices within the fabric, and as soon as any device is moved to a different port in the switch/fabric it will no longer be part of the zone.

WWN

Zones can be defined using the device WWN to uniquely identify zone members. The advantage is retaining the zoning definition even when the device is moved to a different port/switch in the fabric. The disadvantage is that the zoning definition has to be changed whenever an HBA is replaced with another, having a different WWN.

Combination of domain/port numbers and WWN

Zones can be defined using combination of switch domain ID /port number and WWNs to uniquely identify zone members. Advantages and disadvantages as described in the above two methods are applicable to individual zone elements based on their definition.

Note: Movement of devices within the fabric, as described above depending on zoning definition, is only applicable from a zoning perspective. However, there can be other restrictions that will not let movement of devices within the fabric, irrespective of zoning type in effect. For example, some OSs like HP-UX which create device filenames based on 24-bit fabric address will not allow moving the device to a different port because it will change the 24-bit port address and hence will be treated as a different device.

There should not be any dependency between the way zones are configured and the way zones are enforced. It should be possible to have any combination of zoning configuration/zoning enforcement from the above definitions. Due to implementation limits, certain switch products impose restrictions on the way zones are defined and enforced.

Fabric-based zoning

Zoning can be enabled in most storage units, in the switch fabric, and on the host. Fabric switches implement name server-based zoning, in which the zone members are identified by WWN or port location in the fabric.

Note: See the *HP StorageWorks Enterprise Backup Solution Design Guide,* http://h18004.www1.hp.com/products/storageworks/ebs/documentation.html for EBS zoning recommendations.

- **Storage-based zoning**—Storage units typically implement LUN-based zoning, commonly referred to as *LUN masking*. Its basic function is to limit access to the LUNs on the storage port to the specific WWN of the server HBA.

  This form of zoning is needed in most SANs. It functions during the probe portion of the SCSI initialization. The server probes the storage port for a list of available LUNs and their properties. The storage system compares the WWN of the requesting HBA to the defined zone list and returns the LUNs assigned to the WWN. Any other LUNs on that storage port are not made available to the server.
- **Fabric-based zoning**—Fabric switches implement fabric-based zoning, in which the zone members are identified by WWN or port location in the fabric.

  Fabric-based zoning is commonly referred to as *name-server-based* or *soft zoning*. With HP switches, there may also be additional hardware enforcement of the zone (see the switch model product documentation.)

  When a device makes a query to the fabric name server, the name server determines to which zones the device belongs. It then returns to the requesting device information on all members of the zones present in the fabric. Devices in the zone are identified by node WWN, port WWN, or domain/port of the switch to which the device is connected.
- **Host-based zoning**—Host-based zoning can implement WWN or LUN masking.

  There are several approaches for implementing host-based zoning; all will work, in most cases. However, there are pros and cons to each form. The primary forms are single HBA, grouping by application, grouping by operating system, port allocation, and no fabric zoning.

Zoning by single HBA

Zoning by single host bus adapter (HBA) most closely re-creates the original SCSI bus. Each zone created has only one HBA (initiator) in the zone; each of the target devices is added to the zone. Typically, a zone is created for the HBA and the disk storage ports added. If the HBA also accesses tape devices, a second zone is created, with the HBA and associated tape devices in it.

In the case of clustered systems, it might be appropriate to have an HBA from each of the cluster members included in the zone. This is equivalent to having a shared SCSI bus between the cluster members and presumes that the clustering software can manage access to the shared devices.

In a large fabric, zoning by single HBA requires the creation of possibly hundreds of zones; however, each zone contains only a few members. Zone changes affect the smallest possible number of devices, minimizing the impact of an incorrect zone change.

Zoning by application

Zoning by application typically requires zoning multiple, perhaps incompatible, operating systems into the same zones. This method of zoning creates the possibility that a minor server in the application suite could disrupt a major server, such as a Web server disrupting a data warehouse server.

Zoning by this method can also result in a zone with a large number of members, providing greater susceptibility to administrative errors, such as registered state change notifications (RSCNs) going out to a larger group than necessary.

Zoning by operating system

Zoning by operating system allows multiple HBAs of the same operating system type to be grouped with the storage ports being accessed. This helps keep the number of zones to a reasonable limited number in a given fabric. In some cases the zone sizes may become too large and can be possibly get broken into multiple zones within the operating system type depending on the application. Zoning by operating system type can limit any disruptions or fabric change notifications (like RSCNs) to a select few while at the same time retaining the advantage of grouping the HBAs of the same operating system type.

Also, zoning by operating system avoids HBAs from different operating systems from interacting with each other, potentially eliminating any incompatibilities between them.

This zoning philosophy is the preferred method, however, certain specific situations, such as when configuring for common server access to multiple storage types or zoning multiple paths in a single fabric, may require zoning by HBA. See “Data availability” on page 59, “Common server access” on page 135, “SAN fabric switch interoperability rules” on page 121, and “Fabric-based zoning” on page 365 for additional zoning requirements. See “Configuring a SAN” on page 356 for examples of this method of zoning.

Zoning by port allocation

Zoning by port allocation should be avoided unless the administration team has very rigidly enforced processes for port and device allocation in the fabric. It does, however, provide some positive features. One feature for instance, when a storage port, server HBA, or tape drive is replaced, the change of WWN for the new device is of no consequence. As long as the new device is connected to the original port, it continues to have the same access rights. The ports on the edge switches can be pre-associated to storage ports, and control of the fan-in ratio can be established. With this pre-assigning technique, the administrative team cannot overload any one storage port by associating too many servers with it.

No fabric zoning

Using no fabric zoning is the least desirable zoning option because it allows devices to have unrestricted access on the fabric. Additionally, any device attached to the fabric, intentionally or maliciously, likewise has unrestricted access to the fabric.

This form of zoning should only be utilized in a small and tightly controlled environment, such as when host-based zoning or LUN masking is deployed.

Maximum zone size

Generally, the supportable 'maximum number of zones' and 'maximum members in a zone' are very large and are usually constrained by memory usage. These numbers are far larger than the maximum devices that could be connected in fabric configurations supported and hence usually there are no limitations on zone sizes.

However, there is an exception in pure hardware enforced zoning environment on the 2Gb switch models where it's likely that we exceed some preset architectural limits, in which case those ports transition from hard zoning to soft zoning.

The current B-Series switches have a limitation of 64 unique SID entries per quad (predefined groups of four ports). Whenever this limit is exceeded, the affected port/ports will transition from hard to soft enforcement.

This transition is completely transparent to fabric operations, though switch administrator may see warning messages displayed in switch logs. However, data integrity is completely preserved during this transition and HP validated this in large SAN configurations.

The following CLI output indicates a port transitioning to soft zoning:

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

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

These two messages are related and indicate that the zoning configuration has outgrown internally preset architectural limits, thereby forcing the mentioned port be switched from hardware-enforced zoning to software-enforced zoning. It is important to note that only this specific port has turned "soft" and all other members that were zoned with the relevant port still remain hardware-enforced.These warning messages could be seen either statically at zoning configuration/setup time (in case if port-level zoning) or dynamically at run time (in case of WWN zoning).

The command `portzoneshow` shows the status of all ports as follows:

- Hard - Hardware enforcement
- Soft - Name Server plus ASIC-assisted authentication
- All - No zoning enforcement

Zoning guidelines (B-Series switches)

The following are suggested best practices only. However, other zoning configuration methods and zone types as appropriate for each switch are also supported.

- Define zones using WWNs always. All switch models support this type of definition, irrespective of zoning enforcement technique they use, whether it is hardware enforced or name server based or combination of both. Use port WWNs and not node WWNs.
- Exception: For all 1Gb fabric switches, define zones using domain/port numbers for selecting hardware enforced zoning.
- Define zones for all devices in a fabric whenever any zone is defined. In other words, do not define zoning partially for few devices in the fabric and leave others un-zoned.
- Overlapped zones can be defined and there is no upper limit on the number of zones and number of members in a zone.
- Configure zones based on operating environment, on a “per OS” basis,
  See the SAN/Platform zoning requirements for individual storage arrays for exact details, as defined in "SAN storage system rules" on page 171.
- Switch zoning provides security at the port level only and for maximum security in a SAN environment, it's required to use array based LUN security- Secure Manager for XP/VA arrays and SSP (Selective Storage Presentation) for HSG/HSV array controllers.
- To minimize/avoid “soft” port transitions in pure hardware enforced zoning environment (2Gb SAN fabric switches).

- Maintain locality as defined in your SAN design but avoid hosts/targets on the same quad. Quad is a group of predefined consecutive 4 ports (0-3,4-7,8-11,12-15 etc).
- Maintain a connectivity model that populates each quad with the members of the same zone or in other words avoid members of different zones on the same quad particularly when each of them are part of bigger zones. For example, if we have an UNIX zone and an WINDOWS zone, populate all UNIX zone members on one quad and WINDOWS members on a different quad.
- Minimize zone entries by including hosts and targets that practically need to talk to each other. For example, instead of combining all hosts of the same OS type into one zone, consider making smaller zones with only hosts and targets that need to talk to each other.
- Switch CLI command `portzoneshow` can be used to display and verify the individual status of each port whether it's “hard” or “soft” at any given time.

Primary Management Switch Recommendations (B-Series switches)

- Select one switch in the fabric as the “primary management switch”, and use it consistently for all management and control including zoning, Time Services, Fabric Manager, Web Tools, and general administrative access. Using one switch for access lessens the possibility of multiple administrators making changes to different switches in the fabric at the same time. If you have implemented a core-edge topology, it is recommended you use a core switch as the primary management switch. A core switch is is typically connected directly to all other switches in a core-edge fabric, providing optimal communication for administration.
- Configure the primary management switch in the fabric as the preferred principal switch by using the command `fabricprincipal`.
- The establishment of a principal switch in a fabric can vary based on upon the state of the fabric, switch WWN, and whether other switches/fabrics are merging into that fabric. A switch that is the principal switch in a fabric today may not be a principal switch after new switches are added to that fabric or if that fabric reconfigures. The implementation of the `fabricprincipal` command is based solely on mechanisms specified in the Fibre Channel standards. These mechanisms provide a preference for a switch requesting to be the principal switch in a fabric, but they do not provide an absolute guarantee that a switch requesting to be the principal switch will actually achieve this status.

Note: The `fabricprincipal` is only available in Fabric OS version 4.1.0 or later.

The primary management switch can also be used by Fabric Manager as the main access point for that fabric. Once selected, all the other portions of Fabric Manager will primarily access this switch for fabric information.

- Fabric Time Synchronization—The ability to synchronize time within a fabric is available as of Fabric OS versions 2.6.1, 3.1, and 4.1. All switches in a fabric where secure mode is not enabled synchronize their time with the principal switch in the fabric.The principal switch in the fabric can synchronize its clock with an NTP timeserver by identifying the timeserver to the principal switch with the tsclockserver command. In a fabric where secure mode is enabled, switches synchronize time with the Primary FCS, which may or may not be the principal switch

Maximum zone size

The supportable 'maximum number of zones' and 'maximum members in a zone' are very large and are usually constrained by memory usage. There are no limitations on zone size since the maximums are larger than the number of devices that could be connected in supported configurations.

Zoning guidelines (M-Series switches)

The following are suggested best practices only. However, other zoning configuration methods and zone types as appropriate for each switch are also supported.

- Define zones using WWNs always. All switch models support this type of definition. Use port WWNs and not node WWNs.
- Define zones for all devices in a fabric whenever any zone is defined. In other words, do not define zoning partially for few devices in the fabric and leave others unzoned.
- Overlapped zones can be defined and there is no upper limit on the number of zones and number of members in a zone.
- Configure zones based on operating environment, on a “per OS” basis.
  See the SAN/Platform zoning requirements for individual storage arrays for exact details as defined in "SAN storage system rules" on page 171.
- Switch zoning provides security at the port level only and for maximum security in a SAN environment, it's required to use array based LUN security- Secure Manager for XP/VA arrays and SSP (Selective Storage Presentation) for HSG/HSV array controllers.

Special considerations in zoning (for all switch models)

- In high-availability environments like HP-UX service guard, it is required to have homogeneous OS environments on a storage array port and this can be achieved by securing LUNs using array secure manager software and also by properly configuring zones.
- Software environments like OVSAM and Command View for XP/VA do not impose any restrictions on switch zoning. The same supportability exists in these environments as well.
- SANs with Data Protection (tape backup) may require separate rules. Contact your HP representative for more information.

Merging SAN fabrics

This section describes the process for merging two (or more) independent fabrics into a single, larger fabric. This is typically done when you:

- Have grown independent SAN islands to the point where more resources are needed
- Wish to share the resources in two or more fabrics
- Wish to make information in one SAN available to servers in another SAN

With support for longer distances you may also desire to connect geographically separated SAN islands together into a single SAN, spanning across very long distances.

Although StorageWorks SAN designs and components allow versatile configurations, HP highly recommends that you thoroughly review all SANs to ensure they will meet existing SAN rules after they are merged into a single fabric. The newly created fabric should not exceed any existing SAN rules.

Merging fabrics can be a complicated process, especially if the fabrics are large. The procedures in the document require a complete understanding of fabrics, zoning commands, and rules. They also require that the user understand how to use the telnet commands as well as the web-based GUI.

It is important to consider not only current SAN configurations but any future SAN needs that may be required. Most difficulties related to merging SANs are due to the fact that not enough planning was put into future SAN considerations at the time the initial SAN was designed and built. Another problem is that the SANs being merged may be implemented differently.

When fabrics discover each other they must go through basic login procedures, or sanity checks, to determine if they are compatible to work as one fabric. If the discovery process determines they are not compatible then the fabric will segment. This means that although they are physically connected, they will still run as separate fabrics.

When zoned fabrics merge they append their zone configuration database to include each fabric's zone configurations. If a non-zoned fabric merges with a zoned fabric, all zoning information is proliferated to the non-zoned fabric switches. If there was a zone configuration enabled at the time of the merge, then that zone configuration will be enabled on the non-zoned fabric switches as well. This means that any devices that were in the non-zoned fabric will be not accessible until they are added into the current enabled configuration.

Please review these causes of SAN segmenting prior to physically connecting multiple fabrics together:

- *The name of a zone object in one fabric should not be used for a different type of zone object in the other fabric (Zone type mismatch).* In other words, if you create a zone name on Fabric A, that same name should not be an alias or configuration name in Fabric B; otherwise the fabrics will not merge.
- *The definition of a zone object in one fabric is different from its definition in the other fabric (Zone content mismatch).* If an alias, zone or configuration name is the same on both Fabric A and B but the content or definition of that object is different between the fabrics the fabrics will not merge.
- *Zoning is enabled in both fabrics and the zone configurations that are enabled are different (Zone configuration mismatch).* Because of this mismatch the switches within each fabric are not going to assume one fabric has the correct zone configuration enabled. The fabrics will not merge until one of the merging fabrics has its zone configuration disabled.

- *Not only must each switch within a fabric have a unique domain ID but each switch within the multiple fabrics of the enterprise should have a unique id as well.* For example, If Fabric A has five switches with domain IDs 1 through 5 and Fabric B has five switches with the same domain IDs these two fabrics will not merge until all switches within both fabrics have a unique domain ID.

Note: If you use port-level zoning, changing the domain IDs may affect access to devices. Port level zones are based on the domain ID and the port number.

Note: When enabling a new configuration, HP strongly recommends that the fabric be quiesced. Zone membership should not be changed for devices that are actively performing I/O in a fabric. Once the new zoning is enabled, a state change notification is sent to all the nodes that have registered to receive the state change.

Merging fabric together can be accomplished by simply disabling the effective configuration on one fabric, then plugging both fabrics together. The problem with this method is that once you disable the effective configuration, you open up that fabric so all servers will see all storage. Also once you plug the fabrics together, devices from the second fabric will not be accessible until you add them into the effective configuration.

To merge these two fabrics without having to disable the effective configuration for the entire fabric, it is necessary to disable at least one switch in each fabric or have a spare switch available. This will be the switch used for merging the zones and creating the new configuration. Keep in mind that there can be multiple defined configurations, but only one can be the effective or enabled configuration.

Merging a SAN consisting of high-availability redundant fabrics

A redundant fabric provides flexibility to merge a fabric by bringing it offline while redirecting active I/Os to the other fabric. Current I/O operations are not impacted as a result of merge activity. Please note that during the merge the hosts are operating in degraded mode with no alternate data path protection. Any failure on an active path will result in failed I/Os. With proper planning downtime can be minimized. Once the fabric merge is complete and verified, it can be brought online by restoring the I/O paths. The merge process can be repeated for the second fabric after all I/O paths are successfully restored on the first fabric.

For example, your SAN consists of Fabric A and a redundant Fabric B for high availability, and you wish to expand this SAN by merging each of these fabrics with a SAN consisting of Fabrics C and D, also configured for high availability.

1. Review and correct any issues that may cause a fabric segmentation.
2. Verify that each fabric is configured to provide an alternate path to all fabric attached devices.
3. Verify both paths are open to each device that must remain online during the merge.
4. Select one of the two fabrics for merging. In this case we are merging Fabric A with Fabric C.
5. Close all active paths on the fabric selected for merging and prepare single attach devices for downtime. For example if you are running a multipathing software package SecurePath or equivalent, redirect I/Os by performing a failover to the alternate fabric path.
6. Verify that the fabric selected for merging is free from I/O activity.

7. Merge the selected fabric, (Fabric A and Fabric C)
8. Verify all switches are in the newly merged fabric and the zoning has merged correctly.
9. Restore I/O operations on the new fabric from the multipathing software console.
10. Verify both paths are open and restored for each device.
11. Repeat this procedure to merge Fabric B and Fabric D once you are sure all paths to the new fabric have been restored and I/O is running correctly.

Troubleshooting

This section describes troubleshooting steps for isolating problems related to storage access. When initially building a SAN, lack of access either to individual storagesets or entire storage systems is not uncommon. This can usually be traced to an incorrect device setting or an inadvertent cabling or configuration setup error in the initial hardware configuration. These steps will assist you in isolating access problems:

1. On the server:
   a. From the server, determine if lack of access is to all of the storage (the entire storage system) or only to a portion of the storage (specific storagesets). If there is no access to only a portion of the storage system, see step 3.
   b. If access is not available to the entire storage system, verify from the server that the correct driver versions are loaded and that all parameters for the driver are correct. For multi-path applications, verify that the multi-path software is set up correctly.
   c. Verify that all Fibre Channel cables are plugged in and that all green indicator LEDs are on.
   d. Examine the event or error logs on the system.
   e. verify the appropriate cable connection and that the port Link LED is on.
   f. Execute commands on the switch and verify that the server HBA is logged into the fabric correctly. Verify the correct port connection: F-Port, L-Port public, or L-Port private (see the specific HBA for more information on the correct login port types).
      — F-Port: Tru64 UNIX, HP OpenVMS, HP-UX Fabric, Linux, Microsoft Windows NT, Windows 2000, SGI IRIX, and Sun Solaris
      — L-Port, 1 public: Novell NetWare
      — L-Port, x private, x phantom: HP-UX FC-AL
   g. Verify all switch configuration and parameter settings.
   h. Verify that the switch is in the fabric and not segmented.
   i. Verify that all E-Ports are online.
2. On the Fibre Channel switch to which the storage is connected:
   a. Verify the appropriate cable connection and that the port green indicator LED is on.
   b. Execute commands on the switch and verify that the server HBA is logged into the fabric correctly. Verify the correct port connection: F-Port or L-Port private.
      — F-Port: MA6000, MA/RA8000, EMA/ESA12000, EMA16000 set to FABRIC Topology
      — L-Port, x private, x phantom: MA6000, MA/RA8000, EMA/ESA12000, EMA16000 set to LOOP_HARD Topology
   c. For MA6000, MA/RA8000, EMA/ESA12000, EMA16000:

      Verify the connections to the storage system. Execute a “show connections” command at the CLI and verify that the server connection is “online.” Verify the connections are named correctly.
3. On the storage system:
   a. Verify correct controller settings and configuration, “show this” and “show other.”
   b. Verify that the controller ports are online and configured for the correct topology setting.

c. Verify that the storagesets are online to the appropriate controller without errors.

d. Verify that the storagesets are correctly configured and enabled for access, "show unit dn."

e. Verify that unit offset parameters are correct. Also verify that the appropriate storage controller port is indicated in the connection name that will be accessed by the unit you have enabled.

f. Verify that the connection OS parameter type is set correctly for the operating system that is using the connection.

4. General Fibre Channel switch verification:

a. If zoning is in effect, verify that the effective zone matches the enabled zone.

b. Verify that all zone definitions are correct.

c. Verify that zoning alias/nick names are assigned to the correct WWNs.

d. Verify that the servers and storage being accessed are in the same zone. If zoning is in effect, the WWN must be in a zone that is in the enabled configuration or it will not have access to the fabric.

e. From the switch GUI, examine the name server table. Verify that the appropriate WWNs are listed and what zones they are in. Verify that the zones required are in the enabled configuration.

f. Fabric segmentation occurs when you connect together two switches or two fabrics and one of the following mismatch conditions exists between them:

— Zoning configuration mismatch

— Zoning type mismatch

— Zoning content mismatch

— Switch configuration parameter mismatches

Note: All switches in a fabric must have the same switch parameter settings with the exception of the following parameters:

switch name
IP address
domain ID

If you are experiencing fabric segmentation, carefully review and compare these settings in each of the two switches or fabrics.

5. QuickLoop verification:

Note: QuickLoop is only required for HP-UX private loop attachment.

   a. Verify that the QuickLoop license is installed.
   b. Verify that the switch ports are set to QuickLoop mode.
   c. If using QuickLoop with two Fibre Channel switches, verify that the switches are in a QuickLoop partnership.

glossary

This glossary defines acronyms and terms used in this guide and is not a comprehensive glossary of computer terms.

access authorization
A fabric security method that provides frame-level access control in hardware and verifies the SID-DID combination of each frame. Also known as *hard zoning*.

ACS
array controller software.
The firmware that runs HSG80-based storage systems.

arbitrated loop
See FC-AL.

array controller software
See ACS.

ATM
asynchronous transfer mode.
Communications networking technology for LANs and WANs that carries information in fixed-size cells of 53 bytes (5 protocol and 48 data)

B-Series
Fibre Channel switches manufactured for HP by Brocade Communications Systems, Inc.

C-Series
Fibre Channel switches manufactured for HP by Cisco Systems, Inc.

Continuous Access EVA
A storage-based HP StorageWorks product that consists of two or more EVA storage systems performing disk-to-disk replication, along with the Continuous Access management user interface that facilitates configuring, monitoring, and maintaining the replicating capabilities of the storage systems.

Continuous Access Storage Appliance
A storage appliance–based HP StorageWorks product that consists of two or more storage appliances with attached storage systems performing disk-to-disk replication, along with the management interface that facilitates configuring, monitoring, and maintaining the replicating capabilities of the storage appliances.

Continuous Access XP
A storage-based HP StorageWorks product consisting of two or more XP disk arrays performing disk-to-disk replication, along with the management user interface that facilitates configuring, monitoring, and maintaining the replicating capabilities of the storage systems.

controller pair
Two interconnected controller modules that together control a disk array.

corporate fabric
A SAN fabric that uses HP StorageWorks SAN Switch 8, 16, 8EL, 16EL, and Integrated 32/64 model switches.

director fabric
A SAN fabric that uses HP StorageWorks Director 64 or 2/64 model switches.

discovery authentication
A security method where the fabric presents only a partial list of authorized devices. Also known as *soft zoning*.

DRM
HP StorageWorks Data Replication Manager.
DRM is a storage-based HP StorageWorks product that consists of two or more storage systems performing disk-to-disk replication, along with the management user interface that facilitates configuring, monitoring, and maintaining the replicating capabilities of the storage systems.

enterprise
Any large organization where information technology is essential for continuing operations.

Enterprise/Modular RAID Array
Storage system based on an HSG60 or HSG80 controller. These systems include MA6000, MA8000, RA8000, EMA12000, EMA16000, and ESA12000 storage systems

entry-level fabric
A SAN fabric that uses HP Compaq C8 model switches.

EVA
HP StorageWorks Enterprise Virtual Array.
A high-performance, high-capacity, and high-availability storage solution for the high-end enterprise class marketplace. Each EVA storage system consists of a pair of HSV virtualizing storage controllers and the disk drives they manage.

fabric
A network of one or more Fibre Channel switches that transmit data between any two N_ports on any of the switches.

failover
An automatic method for transferring operations from a failed or down system to a secondary, identical system.

FC
Fibre Channel.
A comprehensive set of standards for concurrent communication among servers, storage systems, and peripheral devices.

FC-AL
Fibre Channel arbitrated loop.
A serial, full-duplex data transfer architecture for high-performance storage systems.

FCC
Fibre Channel Congestion Control.

FCIP
Fibre Channel Internet Protocol.

Fibre Channel
See FC.

GBE
Gigabit Ethernet.
Ethernet standards for transmitting data at 1 gigabit per second.

GBIC
gigabit interface converter.
A hardware module that connects fiber-optic cables to a device and converts electrical signals to optical signals.

GLM
gigabit link module.
A 1 Gbps fiber-optic transceiver.

Gbps
gigabits per second.

HBA
host bus adapter.
A hardware device that connects the host server to the fabric.

heterogeneous
A mixed environment that incorporates different operating systems, protocols, and architectures, and equipment from different vendors or product families.

hop
An interswitch link between a pair of Fibre Channel switches.

in-band communication
Communications that uses the same communications channel as the operational data.

ISL
interswitch link.
A connection from an E-port on one switch to an E-port on another switch.

iSCSI

Internet Small Computer System Interface.
A standard protocol that uses SCSI commands to transfer data over IP networks.

IT

information technology.

IVR
inter-VSAN routing.

LC
Lucent connector.

M-Series
Fibre Channel switches manufactured for HP by McDATA Corporation.

MTRJ
Connector.

NSPOF
no single point of failure.
A configuration where failure of a single component does not cause failure of the entire system.

out-of-band communication
Communication that uses a different communications channel than that used by operational data.

oversubscription
Transmissions from one or more devices exceed the capacity of an ISL.

platform
A supported combination of hardware components and operating system version on a server from a specific vendor.

SAN
storage area network.
An intelligent infrastructure that interconnects heterogeneous servers with shared, heterogeneous storage systems.

SC
Subscriber connector.

SFP
small form pluggable.

SSP
Selective Storage Presentation.
A feature that provides the ability to restrict access to a Fibre Channel LUN.

SFP
small form-factor pluggable GBIC.
A 2-Gbps GBIC.

SPOF
single point of failure.

SMB

small and medium business.
Any organization that uses on-site computer systems.

ST
straight tip connector.

topology
The physical structure of interconnected components that form a network.

VCS
Virtual Controller Software.
The firmware that runs the Enterprise Virtual Array storage systems.

VSAN
virtual storage area network.
A logical SAN partition that can be independently configured and managed.

WDM
wavelength division multiplexing.
The technique of placing multiple optical signals on a single optical cable simultaneously. Dense wavelength division multiplexing (DWDM) places many signals on a cable. Coarse wavelength division multiplexing (CWDM) places only a few signals on a cable.

zone
A collection of devices or user ports that are permitted to communicate with each other through a fabric. Any two devices or user ports that are not members of at least one common zone are not permitted to communicate through the fabric.

index

C

D

G

H

I

L

M

N

O

P

Q

R

S

T

U

V

W

X

Z