

SAN Extension

8

With the advent of extension technologies specifically developed for the transport of data it is now possible to consolidate, simplify, manage and integrate storage Fibre Channel SAN fabrics within the enterprise to further exploit its networking investments and lower the cost to manage global storage.

A SAN extension is considered an Inter Switch Link (ISL) connection between two Fibre Channel switches greater than 500 meters for 1 Gbps Fibre Channel switch pair or greater than 300 meters for a 2 Gbps Fibre Channel switch pair. Whether it's called SAN Extension, SAN Bridging or SAN Mirroring, HP seamlessly integrates these new technologies into the benefits of today's Fibre Channel SAN.

This chapter describes the current HP supported technologies and products available that provide SAN Extension in non-Continuous Access or non-DRM topologies. If disaster recovery protection extension is necessary, please read the section [SAN/Continuous Access EVA Integration](#) or [SAN/DRM Integration](#) in Chapter 4, and refer to the *HP StorageWorks Continuous Access And Data Replication Manager SAN Extensions* available at:

<http://h18000.www1.hp.com/products/storageworks/san/documentation.html>

This chapter covers the following major topics:

- [Why Extend the SAN?](#)
- [Supported SAN Extension Technologies](#)
- [Fibre Channel Long Distance Technologies](#)
- [TCP/IP Data Protocol Technologies](#)
- [IP Network Considerations](#)
- [HP StorageWorks SR2122-2 IP Storage Router](#)

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

Currently, HP supports the following technologies for Fibre Channel ISL SAN extension.

- Fibre Channel Long Distance Technologies
 - Long Wave Transceivers
 - Wavelength Division Multiplexing (WDM)
- TCP/IP Data Protocol Technologies
 - Fibre Channel over Internet Protocol (FCIP) using the SR2122-2 IP Storage Router and C-Series MDS IP Storage Services Module

Supported SAN Bridging Technology

- iSCSI to Fibre Channel Bridging using the SR2122-2 IP Storage Router

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 for short. The 2 Gbps transceivers use the “LC” style connectors that are known as Small Form Factor Pluggable transceivers, or SFP. 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 kilometers.

Currently HP supports the following long wave transceivers

- 10 kilometer GBIC
- 100 kilometer GBIC
- 10 kilometer SFP
- 35 kilometer SFP

Long wave transceivers are supported on HP B-Series, HP C-Series, and HP M-Series product lines. C-Series products currently only support 10 kilometer 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 becomes 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.

Refer to the individual switch product line WDM sections in this chapter for additional information about WDM support.

Note: SAN extension using WDM is supported on all WDM devices listed in the CA/DRM SAN Extension Reference Guide. Refer to ftp://ftp.compaq.com/pub/supportinformation/techpubs/user_reference_guides/aa-ru5cb-te.pdf.

Maintaining Performance beyond 5 or 10 kilometers

A primary consideration with extended fabrics is maintaining the performance of the Inter Switch 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 kilometers or 5 microseconds per kilometer. If you calculate the time it takes a frame to travel 100 kilometers and for the "RRDY" (frame acknowledgement) to travel back the same 100 kilometers 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 160 kilometers across the SAN between any two devices.

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 and Core switch 2/64 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 1Gbps only and 2Gbps capable switches at the "L0" portcflongdistance setting only.

"portcflongdistance" 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 "portcflongdistance" 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: 1 Gbps links up to 10 kilometers or 2 Gbps links up to 5 kilometers
No Extended Fabric license required
- LE: 2 Gbps links between 5 and 10 kilometers
No Extended Fabric license required
- L0.5: Extended links greater than 10 kilometer but not more than 25 kilometers.
Extended Fabric license required
- L1: Extended Links greater than 10 kilometer but not more than 50 kilometers
Extended Fabric license required
- L2: Extended Links greater than 50 kilometer but not more than 100 kilometers
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”.

Table 45: 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 and 4.1	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 45: Long Distance Port Matrix (Continued)

Fabric OS	Speed	Port A	Port B	Port C	Port D
HP StorageWorks FOS Version: 3.1 and 4.1	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 = Inter Switch 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 kilometers (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 then the Fabric Long Distance Bit must be off. In other words you cannot have an extended link of greater than 10 kilometers between a pair of StorageWorks 1 Gbps SAN series switches and an extended link greater than 5 kilometers 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 kilometers (portcfglongdistance = L0) between a pair of StorageWorks 1 Gbps SAN series switches and any length extended link between a pairs 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 kilometers between StorageWorks 1 Gbps SAN series switches as long as there are no Inter Switch Link (ISL) connections greater than 5 kilometers 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 3, with a maximum total distance of 160 kilometer across the SAN between any two devices.

For C-Series Fibre Channel switches, HP offers a Coarse Wave Division Multiplexing (CWDM) solution. CWDM is a technology involving similar concepts as Dense Wave Division Multiplexing (DWDM) but is less expensive, less expandable (8 channels max) and works over a shorter distance of 100Km. CWDM allows up to eight 1Gbps or 2Gbps 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. A typical CWDM SFP can reach up to ~100 km in a point-point topology or around 40 km in a ring topology.

Refer to Cisco 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 two C-Series directors, Two C-Series MDS 9200 Series fabric switches, or a combination of a 9500 Series director and a 9200 Series fabric switch. 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.

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 160 kilometers across the SAN between any two devices.

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 kilometers at 1 Gbps and up to 10 kilometers 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 kilometers 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.

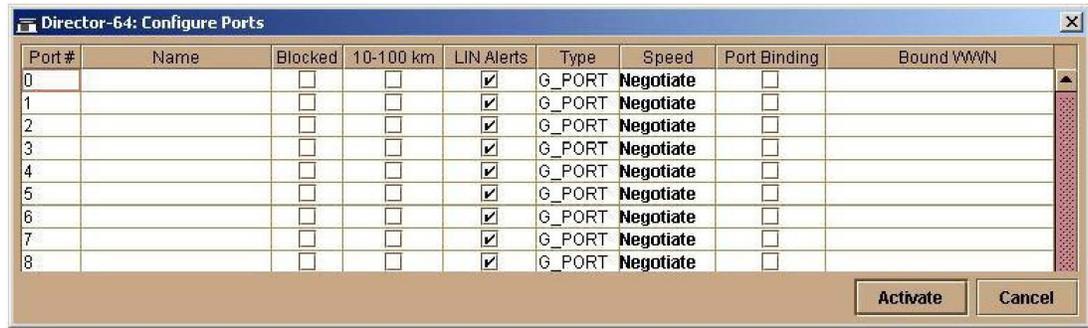


Figure 45: HAFM Configure Ports for 10-100 km setting

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.

FC Switches connect to the FCIP gateways through an E_Port for SAN fabric extension to remote locations.

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 FC SANs using an Internet Protocol (IP) intersite link.

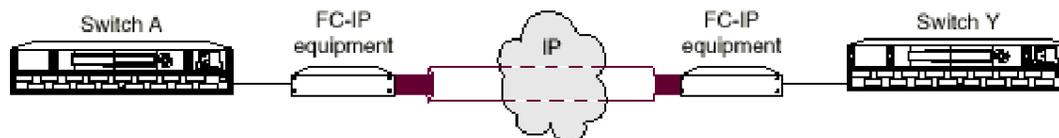


Figure 46: Connecting Fibre Channel SANs with an IP link

Using Internet Protocol over an IP-based network, FCIP can link sites over any extended distance. Local SANs are connected through an IP network to create an extended SAN. An FC-IP gateway is used at each end of the intersite link. Each FC-IP gateway box encapsulates received FC frames into IP packets for transmission over the IP network. Similarly, the FC-IP box extracts the original FC frame from received IP packets and retransmits them to the destination FC node. The FC-IP boxes also handle IP-level error recovery.

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 and M-Series Fibre Channel switches.

The Cisco MDS 9000 IP Storage Services Module is 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 manufacturers documentation for further configuration details.

IP Network Considerations

Considerations Relevant to Using the Existing IP Network

The ability to use your existing network with FC-IP 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 46: 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. Adhoc 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 FC-IP equipment supports Ethernet speeds 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 kilometers per second or 5 microseconds to travel one kilometer.

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 kilometer. Since 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 kilometer, 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. Since 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.

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:

$$\frac{1.0 \text{ MB/s}}{8 \text{ KB per I/O}} = \mathbf{125 \text{ I/Os per second}}$$

2. 50 Miles of Latency

Distance: 50 miles (80 kilometers)

Latency: 8 μ s/mile (5 μ s/kilometer)

Write size: 8 KB

Latency for 1 I/O per mile: 4 trips * 8 μ s/mile = 32 μ s per mileLatency for 1 I/O at 50 miles: 50 miles * 32 μ s/mile = 1.6 ms per I/O

Reciprocal of total latency indicates maximum I/Os:

$$\frac{1.0}{1.6 \text{ ms per I/O}} = \mathbf{625 \text{ I/Os per second}}$$

I/O's multiplied by size of I/O = bandwidth used:

$$625 \text{ I/O per second} * 8 \text{ KB per I/O} = \mathbf{5 \text{ MB/s}}$$

3. 150 Miles of Latency

Distance: 150 miles (241 kilometers)

Latency: 8 μ s/mile (5 μ s/kilometer)

Write size: 8 KB

Latency for 1 I/O per mile: 4 trips * 8 μ s/mile = 32 μ s per mileLatency for 1 I/O at 150 miles: 150 miles * 32 μ s/mile = 4.8 ms per I/O

Reciprocal of total latency indicates maximum I/Os:

$$\frac{1.0}{4.8 \text{ ms per I/O}} = \mathbf{208 \text{ I/Os per second}}$$

I/Os multiplied by size of I/O = bandwidth used:

$$208 \text{ I/O per second} * 8 \text{ KB per I/O} = \mathbf{1.6 \text{ MB/s}}$$

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

Currently 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 networks.
- Guarantee the bandwidth 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.

IP Storage Services Module

Cisco MDS 9000 Family IP storage (IPS) services modules extend the reach of Fibre Channel SANs by using open-standard, IP-based technology. The switch connects separated SAN islands using Fibre Channel over IP (FCIP). It integrates seamlessly into the Cisco MDS 9000 Family, and supports the full range of features available on other switching modules, including VSANs, security, and traffic management.

The IPS module can be used in any Cisco MDS 9000 Family switch and has eight Gigabit Ethernet ports.

- FCIP-FCIP transports Fibre Channel frames transparently over an IP network between two Cisco MDS 9000 Family switches or other FCIP standards-compliant devices. Figure 47 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 Inter Switch 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 Cisco MDS9000 Family enhanced capabilities including VSANs, advanced traffic management, and security across remote connections.

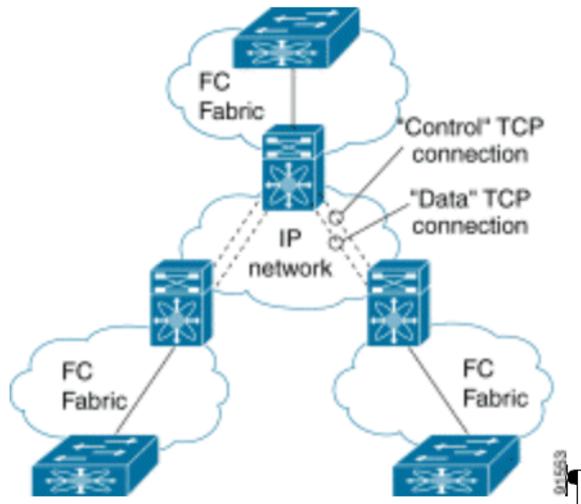


Figure 47: FCIP Scenarios

Table 47: 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

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 either 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

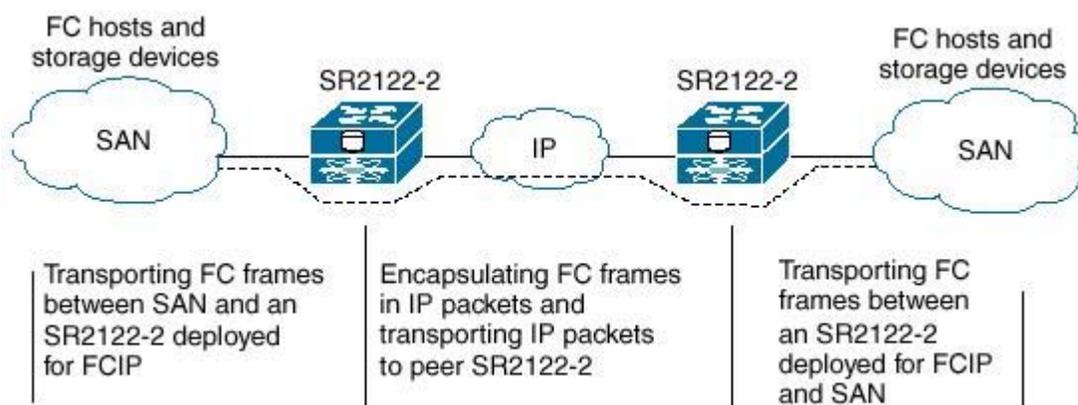
Fibre Channel over IP (FCIP) enables SR2122-2 Storage Routers to provide connectivity between FC hosts and FC storage devices over an IP network. To deploy FCIP, two SR2122-2 Storage Routers are required. Each system is configured for FCIP and connected to a SAN. An FC host or FC device needs no additional hardware or software to access storage devices via an SR2122-2 Storage Router deployed for FCIP.

Note: Refer to [Table 24](#) for a list of devices that are supported for FCIP heterogeneous SAN extension.

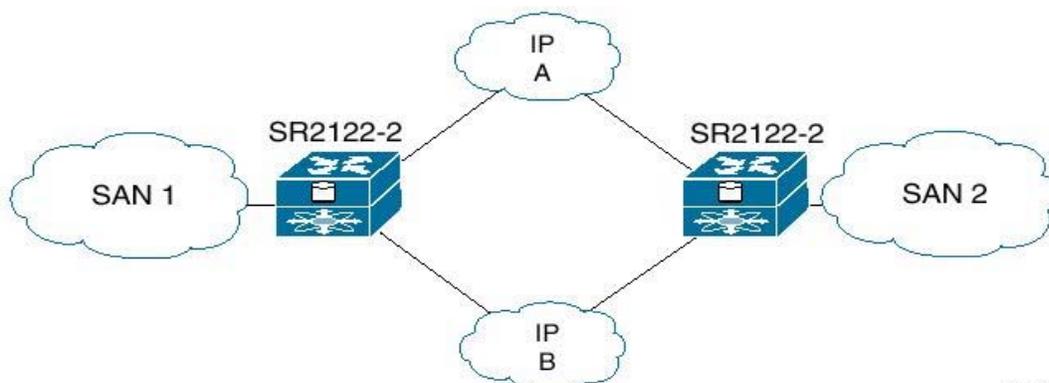
With FCIP, peer systems transport FC frames over an IP network. From the perspective of the SANs the storage devices accessed through the peer systems appear to be part of one unified SAN.

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 FC 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 FC frames over the FC interfaces to the storage devices.

The peer systems are connected to each other through an IP network.

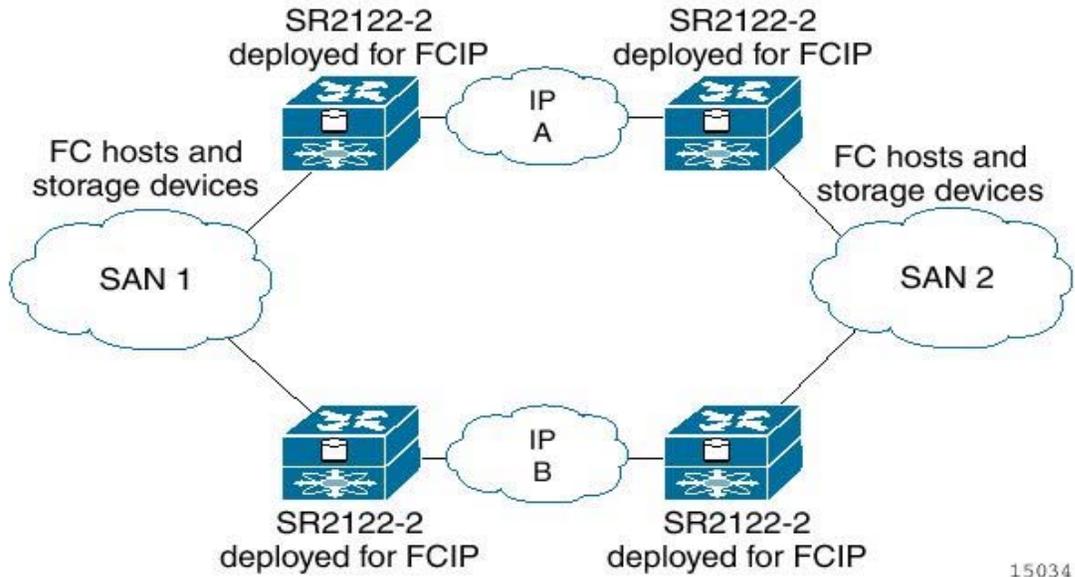


In this example a FC host or FC device connects to one or more Fibre Channel interfaces of each peer SR2122-2 Storage Router deployed for FCIP. Each SR2122-2 connects to the IP network through one of its Gigabit Ethernet interfaces. Through the IP network each FCIP instance accesses its peer, thereby connecting the SANs.



15033

In this example configuration, an FC host or FC 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 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 FC hosts or FC 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 FC host or storage devices).

HP StorageWorks SR2122-2 IP Storage Router - iSCSI Overview

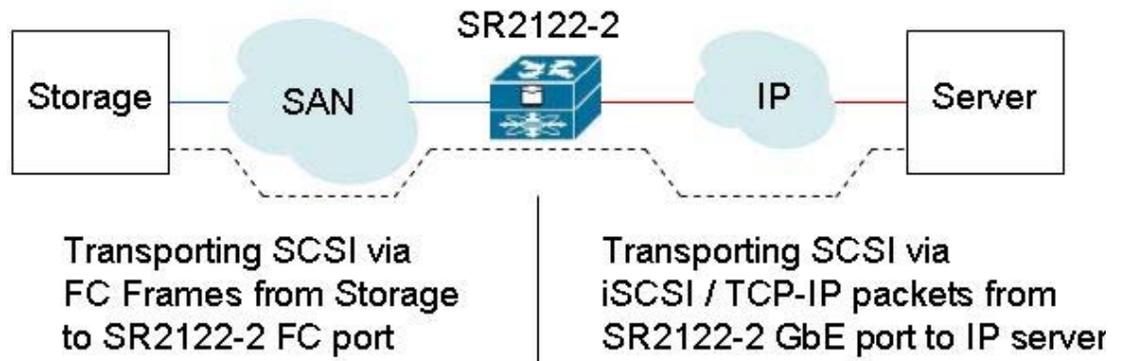
The SCSI transport protocol iSCSI maps block-oriented (CDB) storage data over TCP/IP networks. The iSCSI protocol enables universal access to storage devices and storage-area networks (SANs) over standard Ethernet-based TCP/IP networks.

These networks may be dedicated networks or may be shared with traditional Ethernet applications. IP LAN/WAN routers and switches can be used to extend the IP storage network to the wide area of applications such as synchronous and asynchronous remote disk copy or tape backup and restore.

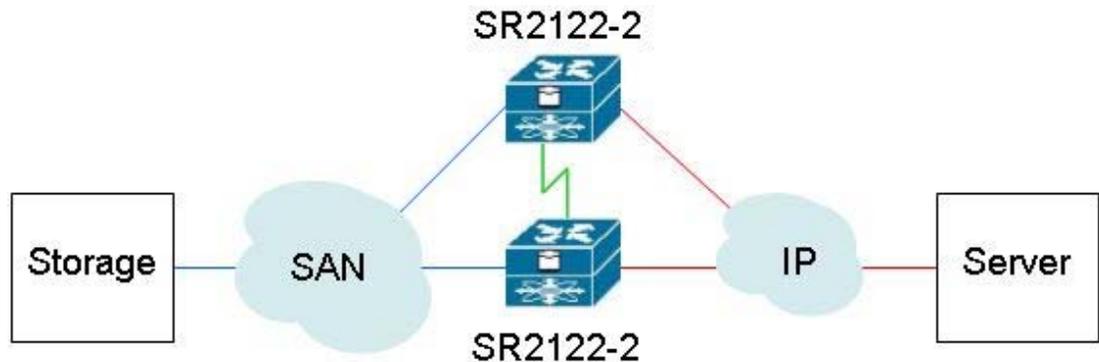
Note: Although the iSCSI protocol is written as a complete data transport from host to storage, this section will only discuss the current HP supported topology of iSCSI hosts to Fibre Channel storage using the HP SR2122-2 IP Storage Router.

SCSI routing provides IP hosts with access to FC storage devices as if the storage devices were directly attached to the hosts, with access to devices being managed primarily in the storage router. An iSCSI target (also called logical target) is an arbitrary name for a group of physical storage devices. The iSCSI targets are created and mapped to physical storage devices attached to the storage router. The SR2122-2 presents the iSCSI targets to IP hosts (iSCSI initiators) as if the physical storage devices were directly attached to the hosts.

With SCSI routing, storage devices are not aware of each IP host; the storage devices are aware of the storage router and respond to it as if it were one FC host.



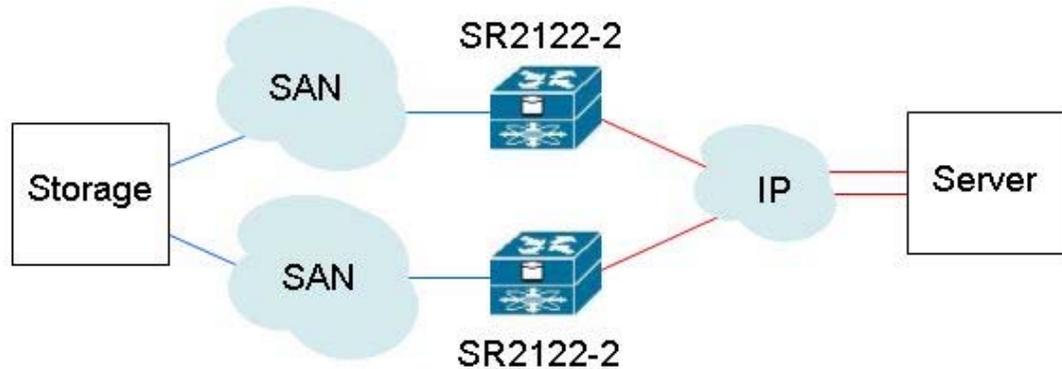
In this example FC storage connects to the Fibre Channel interface of the SR2122-2 Storage Router deployed for iSCSI. The SR2122-2 connects to the IP network through one of its Gigabit Ethernet interfaces. The server accesses the storage served from the router through the IP network.



In this example configuration, two storage routers are connected (clustered) together to allow the storage routers to back each other up in case of failure.

Note: A storage router can participate in a cluster only if it is deployed for SCSI routing.

In a cluster, storage routers continually exchange High Availability information to propagate configuration data to each other and to detect failures in the cluster. The storage routers exchange HA information through a separate network connected to the HA interface of each storage router.



This example shows an even more reliable iSCSI configuration in which pairs of SR2122-2 Storage Routers and SANs provide full redundancy. In this configuration, loss of a SR2122-2 or loss of connectivity through one of the SANs can be tolerated with no loss of connectivity between the storage and server.

Note: For multiple paths between SANs, it is assumed that the multipath management is being done by an entity outside the SR2122-2 (for example, by management applications on the FC host such as HP Secure Path).

SR2122-2 Hardware and Software Support

This section lists the hardware, devices, and operating systems that are compatible with this SR2122-2 IP Storage Router.

Storage Array Hardware Support

The following hp Storage Array products are supported:

- MSA1000
- RA/MA8000
- ESA/EMA12000
- EMA16000
- Enterprise Virtual Array
- VA7100
- VA7400/7410
- XP128/1024

Fibre Channel Switch Hardware Support

The SR2122-2 Storage Router is supported with the HP B-Series, C-Series and M-Series product line switches listed in Chapter 3.

Network Interface Controller (NIC) Hardware Support

The following Proliant Server Network Interface Controllers are supported:

- NC6136 Gigabit Server Adapter
- NC7131 Gigabit Server Adapter
- NC7770 PCI-X Gigabit Server Adapter

The following Blade Server Network Interface Controller is supported:

- NC7781 Gigabit Server Adapter

Operating System Software Support

- Microsoft Windows 2000 SP2 with either Microsoft hotfix Q302895 or Q248720 and Microsoft hotfix Q318271, SP3
- Microsoft Windows 2003
- MSCS support (HP iSCSI driver only)
- Red Hat Advanced Server 2.1
- Secure Path (Windows 2000, Windows 2003 only)

Compaq Network Teaming Software Support

- Compaq Network Teaming (Windows 2000, Windows 2003 only)

SR2122 Management Software Support

The following HP management software is supported:

- Compaq Insight Manager 7
- hp OpenView Storage Area Manager (SAM)

iSCSI Initiator Software Support

- HP
- Microsoft (MSCS, Secure Path not supported)

SR2122-2 iSCSI Configuration Rules

SR2122 Router Rules

Table 48: SR2122-2 Router Rules

Router Rule	Maximum
Maximum scsirouter instances per SR2122-2 Router (and per SR2122-2 Router Cluster)	12
Maximum iSCSI host connections per SR2122-2 SCSI Router instance	32
Maximum active logical units (LUNs) per SR2122-2 Router	30
Maximum active targets per SR2122-2 Router	30

- The SR2122-2's 2nd fibre Channel port (FC2) is not supported as a redundant iSCSI SAN port for FC1. FC2 can, however, be configured as a FCIP port
- Direct connect of the SR2122-2 FC ports to any HP storage array is not supported.
- The SR2122-2 Management port must be in a different subnet than the SCSI Router Instances.
- The SR2122-2 fibre channel ports appear as host bus adapters to the FC switches and to all storage arrays.

ISCSI Host Rules

Table 49: ISCSI Host Rules

iSCSI Host Rule	Operating System	Maximum
Maximum targets accessed per iSCSI host	Windows 2000, Windows 2003	30
	Red Hat Linux, HP-UX	4
Maximum active LUNs per target	Windows 2000, Windows 2003, Linux, HP-UX	30

Operating System Rules

- 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 both SCSI-3 Transparent Failover Mode and Multiple-Bus Failover Mode.
- Without hp Secure Path the RA8000/MA8000 and the Enterprise Virtual Array is supported with the SR2122 accessing only one controller port. This will disable controller failover protection.
- The MSA1000 is supported with the SR2122-2 accessing only one MSA controller port. This will disable controller failover protection.

Fibre Channel Switch/Fabric Rules

- The SR2122-2 is supported on the HP B-Series, C-Series and M-Series Switches.
- The SR2122-2 should only be zoned with the storage devices that it will access. Zoning the SR2122-2 with other servers is not supported.

Management Software Rules

- hp OpenView Storage Area Manager (SAM) support is limited to property support only. It will identify the device, and by clicking on it, one can launch the device embedded web server interface or telnet. A device specific plug-in for the SR2122 is available on the
- SAM Website.
- CIM 7 Supports the SR2122's SNMP management capabilities.
- Management of the storage arrays through the SR2122 is not supported. Please use the recommended application/element manager to configure the storage array.

SR2122-2 FCIP Configuration Rules

SR2122 Router Rules

- IP network speeds 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

is 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.

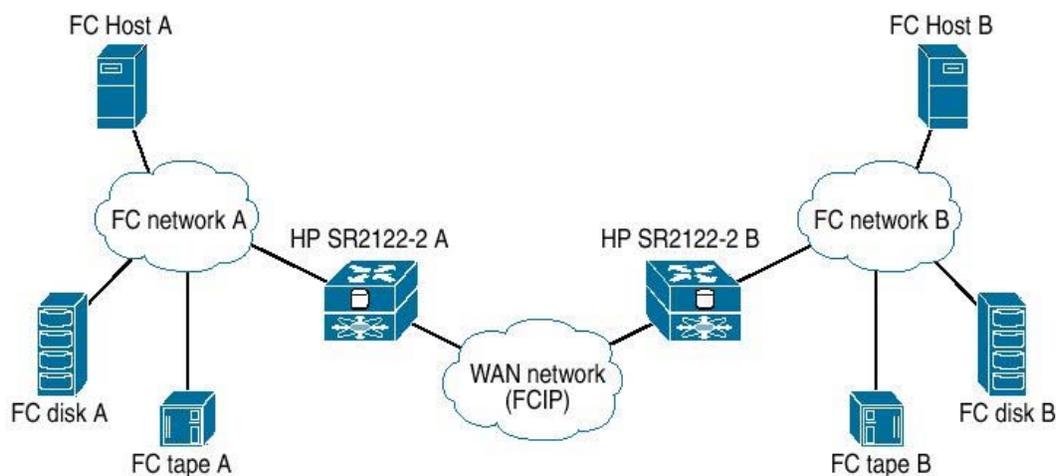


Figure 48: FCIP only

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. Since 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 SR2122-2s, to avoid a single point of failure, a total of four SR2122-2 units should be included in this configuration.

As shown in [Figure 48](#), 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.

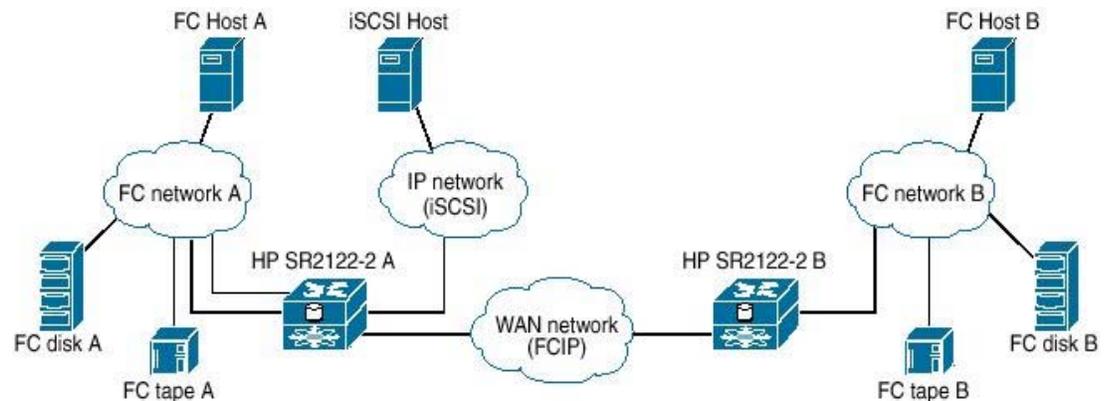


Figure 49: FCIP with Local iSCSI Hosts

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 connect to the large SAN through a WAN network as shown below.

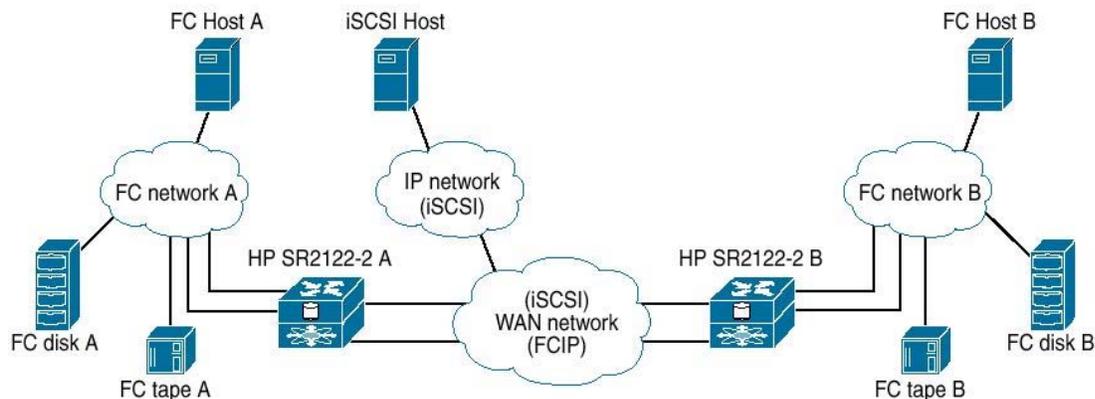


Figure 50: FCIP with Remote iSCSI Hosts

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 over-subscription of the WAN connection.

Sample Configurations

For maximum supported SAN and Storage configurations, see Chapter 3 and Chapter 4.

For maximum supported IP configurations please consult with your network administrator

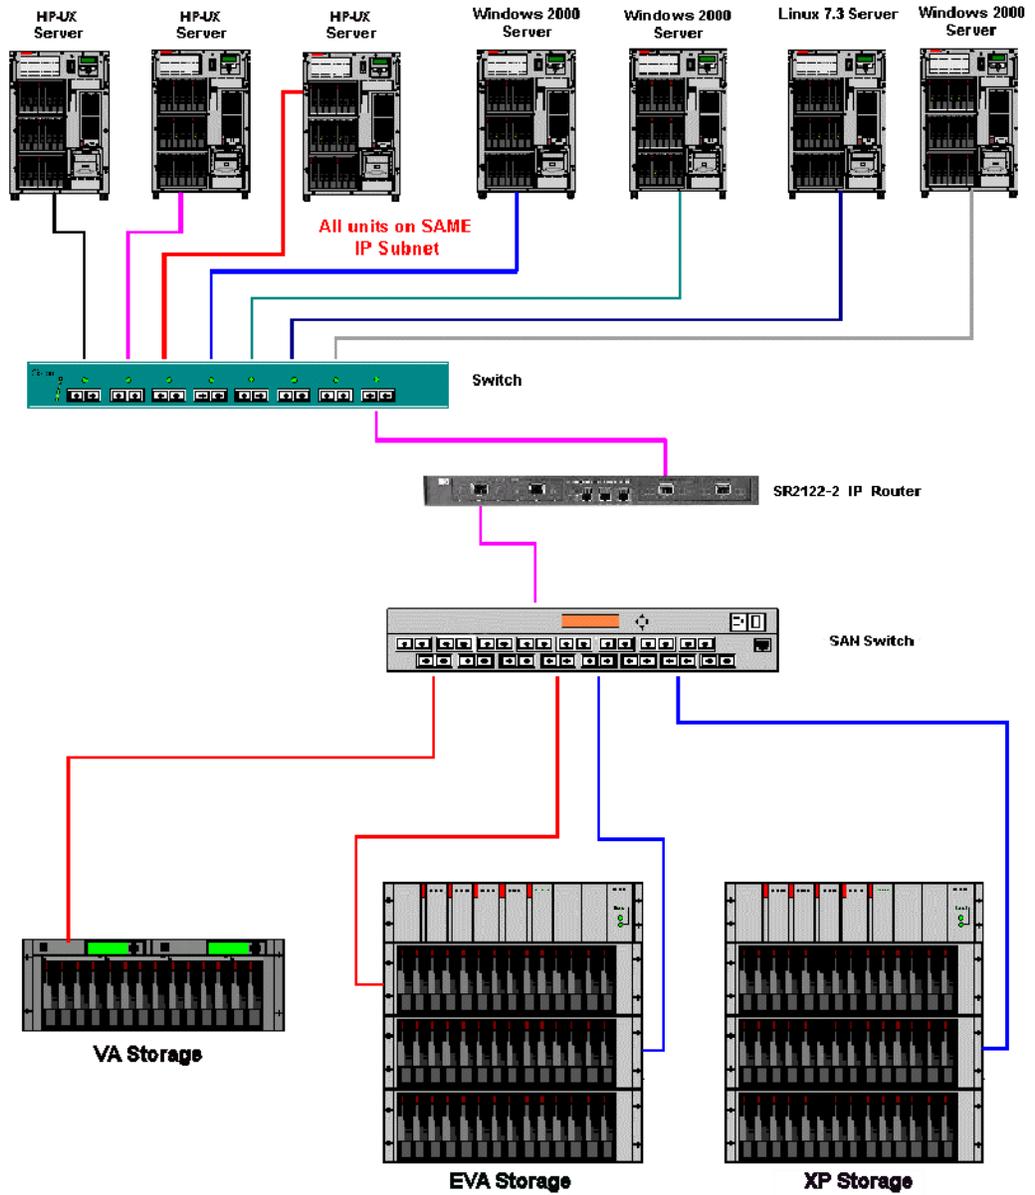


Figure 51: Example of Multiple OS Systems in a Non-Redundant Path Configuration

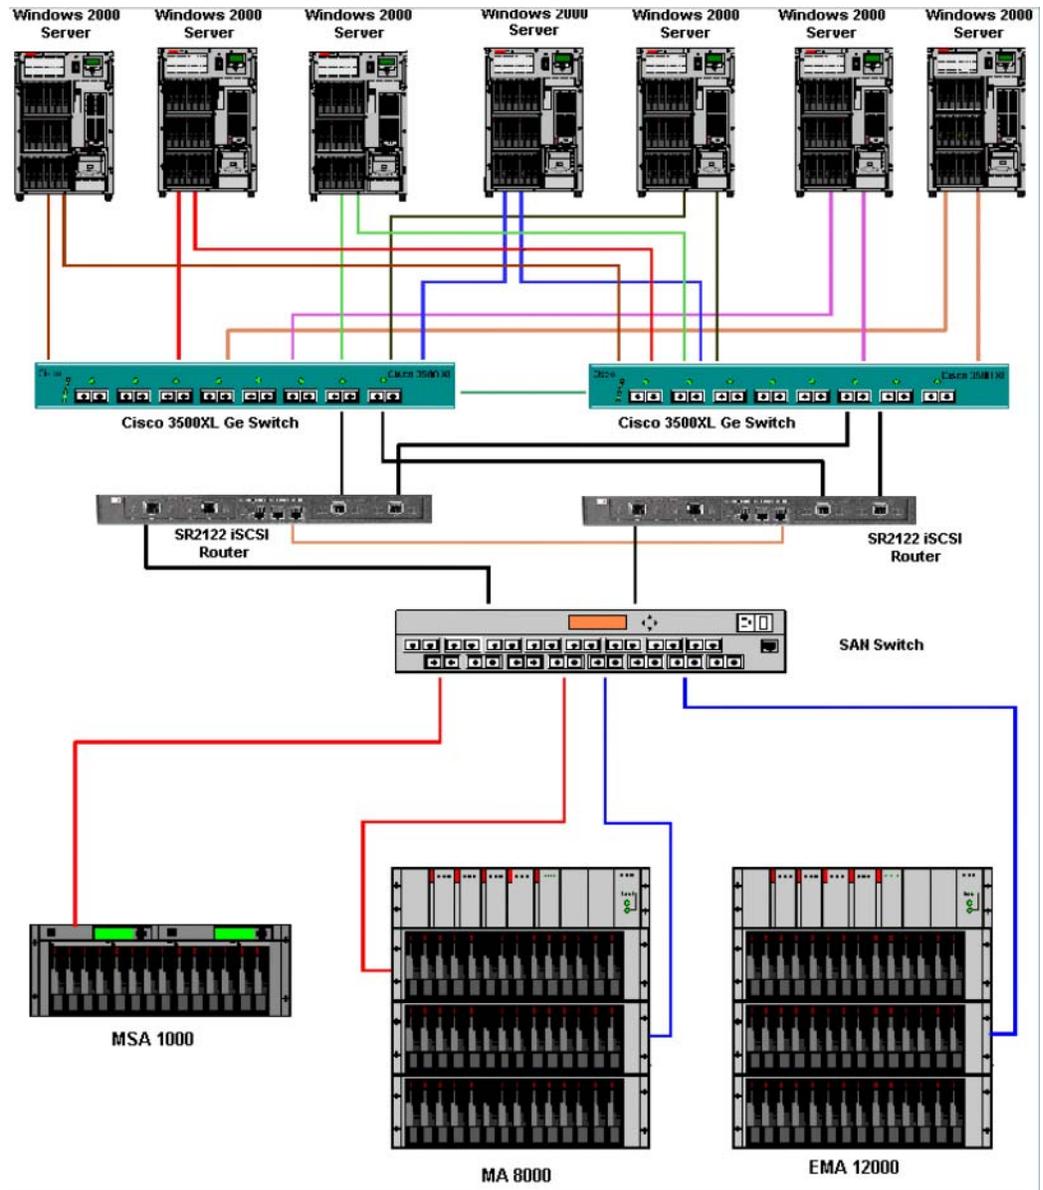


Figure 52: Windows 2000 Servers with NIC Teaming: 2 Node SR2122-2 Cluster

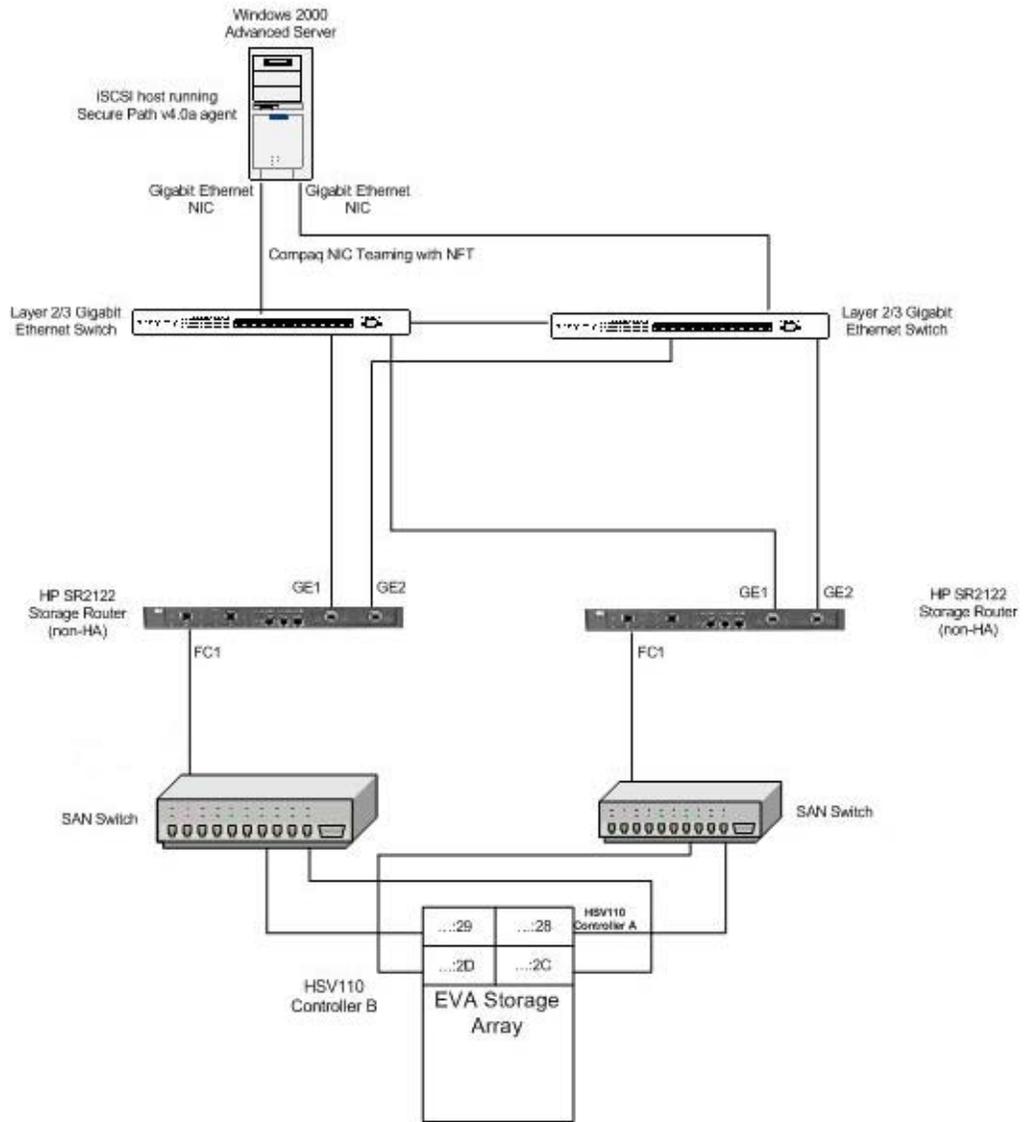


Figure 53: Secure Path Configuration

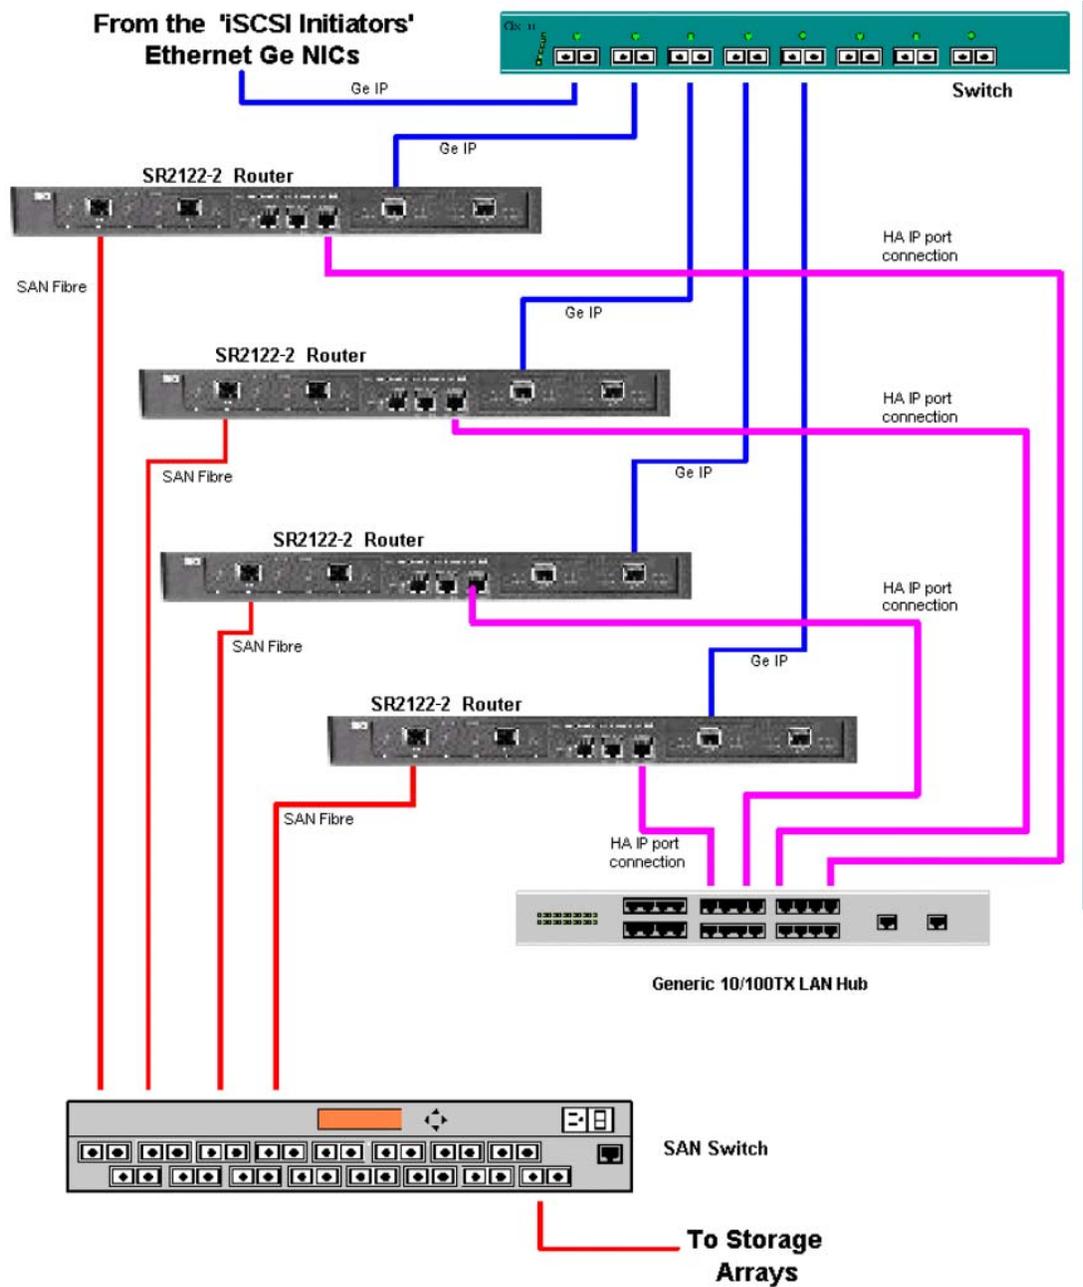


Figure 54: Maximum SR2122-2 Cluster Configuration Using HA Ports

