

# SAN Extension

## 8

With the advent of new 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, or see the *hp StorageWorks Continuous Access/DRM Extension 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?](#)
- [Supported SAN Extension Technologies](#)
- [Fibre Channel Long Distance Technologies](#)
- [TCP/IP Data Protocol Technologies](#)
- [IP Network Considerations](#)
- [SAN Valley SLI000 IP-SAN Gateway](#)
- [SR2122 iSCSI 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)

---

**Note:** Not all technologies are supported by all HP fibre channel switch families. 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 both HP M-Series and HP B-Series product lines.

### 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 increase the buffer-to-buffer credits in order to maintain line speed performance.

Wavelength Division Multiplexing is supported for both 1 Gbps and 2 Gbps.

---

**Note:** SAN extension using WDM is supported on all WDM devices listed in the DRM Over a WDM-enabled Inter-site Link document. Refer to <http://h18006.www1.hp.com/products/sanworks/drm/documentation.html>.

---

### Maintaining Performance beyond 5 or 10 kilometers

The primary consideration with extended fabrics is maintaining the performance of the Inter Switch Link or E-port 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.

The default buffer-to-buffer credit setting is capable of maintaining full line speed performance at 1 Gbps up to 10 kilometers and up to 5 kilometers at 2 Gbps per second.

---

**Note:** Performance may vary depending on the condition of the fiber optic cable and connections between switches. Losses due to connectors, cable splices, tight bends, cable cladding degradation and other conditions can affect the performance over the link and the maximum distance possible.

---

Once the distance between a pair of switches exceeds 10 kilometers for a 1-Gbps or 5 kilometers for a 2-Gbps ISL, the number of buffer-to-buffer credits should be increased. It is not required that you increase the buffer-to-buffer credits, but the penalty you pay may be a reduction in performance. At the speed of light it takes a full second for light to travel 200,000 kilometers or 5microsecond 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 60 buffer-to-buffer credits to keep the link running at full speed.

There are different limits on the number of extended links that are allowed between the two HP switch product lines as well as differences in the commands to configure the buffer-to-buffer credits for each switch product line. The following sections detail these limits and the procedures for configuring extended links for each of the two HP switch product lines.

## 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 300 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 49: HAFM Configure Ports for 10-100 km setting

## 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. Only one segment can be an Extended Link greater than 10 kilometers.

### 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 2.0.x installed,
- StorageWorks 2 Gbps SAN switch series with version 3.0.x installed, or
- hp StorageWorks core switch 2/64 (Core Switch 2/64) switches with version 4.0.x installed.

An extended fabric link can only exist between two switches of the same technology, meaning a StorageWorks 1 Gbps SAN series switch can only have an extended fabric link to another StorageWorks 1 Gbps SAN series switch.

The StorageWorks 2 Gbps SAN series switch can only have an extended fabric link to another StorageWorks 2 Gbps SAN series or an hp StorageWorks core switch 2/64 switch.

StorageWorks 2 Gbps SAN series or hp StorageWorks core switch 2/64 switches cannot have an extended fabric link to a StorageWorks 1 Gbps SAN series switch.

### “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**
- 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 38: Port Quad configuration limits**

Speed	Port a	Port b	Port c	Port d
1 Gbps	L2	L1/LO	Fx/LE	X <sup>1</sup>
1 Gbps	L2	Fx/LE	Fx/LE	Fx/LE
1 Gbps	L1/LE/LO/Fx	L1/LE/LO/Fx	L1/LE/LO/Fx	L1/LE/LO/Fx
2 Gbps	L2	X <sup>1</sup>	X <sup>1</sup>	X <sup>1</sup>
2 Gbps	L1	L1	X <sup>1</sup>	X <sup>1</sup>
2 Gbps	L1	LO/LE	LO/LE/Fx	X <sup>1</sup>
2 Gbps	L1	Fx	Fx	Fx
2 Gbps	LO/LE/Fx	LO/LE/Fx	LO/LE/Fx	LO/LE/Fx

1. In some configurations, not all ports in a quad are enabled. This factor must be considered when calculating the total number of switches required for the fabric.

Fx = Fabric port      L0, LE, 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 (portcflongdistance = 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, 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 (portcflongdistance = 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.

## 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 50: 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 Gateway Products supported for Heterogeneous SAN Extension

SAN Valley SL700/SL1000 IP-SAN Gateways are supported with B-Series and M-Series Product Line switch models for heterogeneous SAN extension.



## 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 39: 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 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  $\mu$ s/mile (5  $\mu$ s/kilometer)

Write size: 8 KB

Latency for 1 I/O per mile: 4 trips \* 8  $\mu$ s/mile = 32  $\mu$ s per mile

Latency for 1 I/O at 50 miles: 50 miles \* 32  $\mu$ 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  $\mu$ s/mile (5  $\mu$ s/kilometer)

Write size: 8 KB

Latency for 1 I/O per mile: 4 trips \* 8  $\mu$ s/mile = 32  $\mu$ s per mile

Latency for 1 I/O at 150 miles: 150 miles \* 32  $\mu$ 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.

## SAN Valley SL1000 IP-SAN Gateway

The SL1000 IP-SAN Gateway extends SANs over IP networks with wire speed throughput, enabling cost-effective deployment of remote-storage applications. The SL1000 connects Fibre Channel SAN switches and storage resources across networks with wire-speed connectivity over MANs and WANs.

The SL1000 connects storage islands over IP networks with over 8 Gbps of total chassis throughput, providing the fastest most scalable solution available today. With independent, dedicated, hardware-based Fibre Channel (FC)-to-Gigabit Ethernet (GbE) encapsulation engines, the SL1000 provides superior channel/performance, scalability, and throughput without microprocessor interrupts.

Support for standards-based QoS ensures that storage traffic is always prioritized for delivery across the IP network. The SL1000 also maps Fibre Channel buffer-to-buffer flow control to the Gigabit Ethernet channel, maximizing data throughput and efficiency.

### SL1000 Features

*IP-SAN connectivity.* Four Fibre Channel to Gigabit Ethernet channels in a 1-rack unit (1RU) chassis interconnect Fibre Channel SANs using Gigabit Ethernet within the data center or co-location facility.

*Fibre Channel tunneling.* Hardware-based encapsulation of Fibre Channel frames into Ethernet packets enables data transport across IP networks for distances up to 2000km at rates up to 1Gbps. Both Layer 2 and UDP Layer 4 encapsulation modes are supported.

*Management interface options.* The SL1000 can be managed via the CLI, a command line interface that resides on the SL1000 and provides two levels of user access: one Administrator and up to four non-Administrator users. The SL1000 can also be managed via SVS Management System, a Java-based SNMP v1 application that can be installed on a supported network workstation; the application uses SNMP Community Names to control access levels. Additional flexibility in management interfaces and security is available through an embedded SNMP agent that supports SNMP v1, v2, or v3 applications.

*System and network performance monitoring.* Continuous monitoring of redundant system power and cooling, network performance statistics, and channel integrity. If the device is managed from SVS Management System, system events can be recorded in a circular event log. If SVS Management System is integrated with a supported network management platform, traps from the SL1000 can be forwarded to the management application.

### SL1000 setup in an HP SAN


Because we are now taking advantage of two separate and distinct data transport protocols - Fibre Channel and UDP - to create one virtual Storage Network, care must be taken to follow the rules and guidelines for each.

### SL1000 Fibre Channel Setup

The SL1000 has been qualified in the supported maximum 28-switch topology on the HP B-Series product line only. Although switches may physically be in different locations throughout the Storage Network, these switches are in a single fabric with only one Principal switch.

Because of the FCIP transparency to the Fibre Channel SAN, the same infrastructure/storage rules apply as when creating any large Fibre Channel topology (i.e. maximum switches, maximum hops, zoning, etc.).

The SL1000 to SL1000 connection does not add to the hop count. Only the connect from the local switch to the remote switch is considered a hop.



**Figure 51: One hop total for SL1000 fabric connection**

Multiple FCIP Intersite Links are supported as long as the supported infrastructure/storage rules are followed.

The SL1000 Intersite Link may be between different HP B-Series FC switches as long as it is a supported ISL connection as described in this guide.

The HP M-Series switch port connected to the SL1000 uses the standard short wave GBIC or SFP transceiver and does not support long distance settings greater than the default of L0. This does not affect the overall IP distance.

## SL1000 IP Network Setup

It is highly recommended that the Storage Administrator work closely with the Network Administrator during the IP Network setup of the SL1000.

In summary, the steps necessary for initial SL1000 network setup are as follows:

- Read the current configuration of the channel.
- Name the channel.
- Configure the channel's encapsulation mode to UDP\_LAYER4.
- Configure the channel's Gigabit Ethernet port IP addresses: source IP address, source IP subnet mask, source gateway IP address (if used), and destination IP address.
- Configure the channel's IP TTL count based on the number of hops from the source to the destination.
- Configure the MTU size/upper bound for this channel.
- Configure the Maximum Data Transmit Throughput rate for this channel.
- If the GbE port will be connecting to a switch that does not support auto negotiation, disable channel auto negotiation.
- Enable the channel.

Details for each step can be found in SAN Valley's August 2002 edition of "*SL1000 Configuration and Operations Guide*", part number 72-02007-02 Rev. B.

## Controlling Data Flow between SL1000 IP-SAN Gateways

This section describes operator-configurable channel settings that cause the Gigabit Ethernet port to send packets formatted to utilize network options for controlling data flow.

These settings configure fields in the packet header so that traffic from that channel can be directed using optional data flow features supported on the network. These features implement traffic classification policies and quality of service (QoS) policies.

### Using the IP type of service (TOS) field

For a UDP Layer 4 channel connection, the 8-bit IP type of service field can be changed from the default value of 0; the field can be set to any value from 0 to 255. Layer 3 devices that manage traffic using a relative priority scheme can interpret the field. Such devices read the type of services (TOS) or DiffServ (DS) field in the IP header of each incoming packet and apply one of the following packet-prioritizing mechanisms:

**Type of service (TOS).** This standard is specified in RFC 791.

**Differentiated services (Diffserv).** This standard is proposed in RFC 2484.

In this proposal, the field is referred to as DS rather than TOS.

### Using the UDP port number fields

For a UDP Layer 4 channel connection, the 16-bit UDP source and destination port numbers can be changed from the default values of 1024; either port can be set to a value from 1024 to 65,535 to uniquely identify packets associated with that channel. There are two situations in which it is useful to change the UDP port numbers:

**Packet routing based on UDP port numbers.** If switches or routers are configured to filter and direct packets based on a particular set of UDP destination port numbers, two devices can communicate via these ports to use the packet routing services configured on the connection. In this situation, configurable UDP port numbers support a traffic classification policy.

**Avoiding conflict with another application or device.** Even when the IP-SAN gateway devices do not communicate via particular UDP ports (that is, when the devices do not utilize a packet routing service), there can be another use for changing the UDP port numbers. If another application or type of device uses 1024 as the UDP source or destination port, the IP-SAN gateway devices can be configured to use a different UDP port.

### Using the VLAN header fields (extension of the 802.3 header)

For a Layer 2 or UDP Layer 4 channel connection, packet headers can be formatted to include a 4-byte header tag that supports port-based VLANs. Two fields are operator configurable: the VLAN identifier (VID) can be set from 0 through 4095, and the user priority can be set from 0 through 7. These fields can be interpreted by VLAN-enabled switches and routers.

The result of using network options for controlling data flow is dependent upon the implementation of a packet-prioritizing mechanism on your network.

Details for controlling data flow can be found in SAN Valley's August 2002 edition of "*SL1000 Configuration and Operations Guide*", part number 72-02007-02 Rev. B.

## SR2122 iSCSI Storage Router

The HP StorageWorks SR2122 iSCSI Storage Router is an iSCSI to Fibre Channel Bridge that enables access to block storage on a Fibre Channel SAN across an Ethernet network. The SR2122 Storage Router is the first HP storage product delivered to this emerging market of customers looking to extend the benefits of their existing FC SAN's.

### iSCSI SR2122 Storage Router Documentation

The iSCSI SR2122 Storage Router kit also includes:

- hp StorageWorks iSCSI Storage Router 2122 Command Line Interface Reference Guide: Part Number 306001-002
- hp StorageWorks iSCSI Storage Router 2122 User Guide: Part Number 304835-002

Additional documentation, including white papers and best practices documents, are available via the HP website at: <http://www.hp.com>.

### iSCSI Protocol

The SCSI transport protocol iSCSI maps block-oriented (CDB) storage data over TCP/IP networks (through iSCSI PDUs). 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 chapter will only discuss the current HP supported topology of iSCSI hosts to Fibre Channel storage.

---

### iSCSI in a Heterogeneous SAN

The iSCSI protocol is written as a complete data transport from host to storage. This chapter will only discuss the current HP supported topology of iSCSI hosts to Fibre Channel storage using the HP SR2122 iSCSI Router.

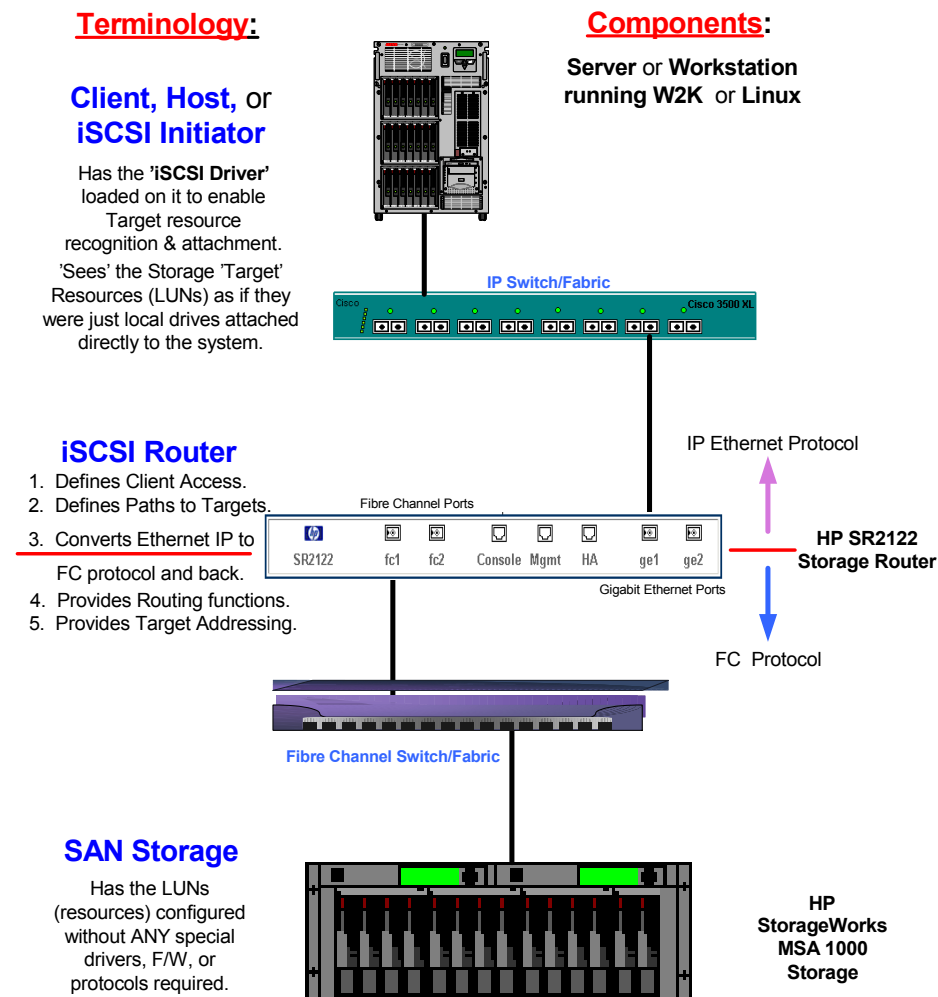
For further details please read “*hp StorageWorks iSCSI storage router 2122 User Guide* PN: 304835-002” Section 4 - SCSI Routing Overview.

### Typical Network Structure

The addition of iSCSI to the SAN means that two data protocols need to be configured and managed – TCP/IP and Fibre Channel. IP hosts with iSCSI drivers access the SR2122 through an IP network connected to the Gigabit Ethernet interface of each SR2122. The SR2122 accesses FC storage devices connected to the Fibre Channel interfaces of each SR2122.



A typical IP/FC network example:



**Figure 52: IP/FC Network Example**

### iSCSI Initiator

The IP host, or iSCSI Initiator uses the iSCSI Driver to enable target resource recognition & attachment to the SR2122 over IP. The iSCSI driver is configured with the Gigabit Ethernet IP address of each SCSI routing instance running on the SR2122 with which the host is to transport SCSI requests and responses. The iSCSI initiator sees the storage resources (LUNs) as if they were just local drives attached directly to the server.

On the SR 2122, IP host access is controlled through an access list and optional VLAN identifier (VID).

The iSCSI Initiator may use multiple NIC cards with teaming for failover. The term “team” refers to the concept of multiple network adapters working together as a single network adapter, commonly referred to as a Virtual Network Adapter.

### IP Network

The IP network is the infrastructure used to transfer data between the iSCSI initiator and the HP SR2122.

**SR2122 iSCSI router**

The HP SR2122 iSCSI Router is the bridge between the iSCSI Initiators and the SAN storage that converts IP to FC protocol and provides iSCSI Initiator access to FC storage.

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 Storage Router 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.

You can configure storage routers in a cluster to allow the storage routers to back each other up in case of failure.

A storage router cluster consists of two storage routers connected as follows:

- Connected to the same hosts
- Connected to the same storage systems
- Connected to each other through their management and high availability (HA) interfaces

In a cluster, storage routers continually exchange HA information to propagate configuration data to each other and to detect failures in the cluster. The storage routers exchange HA information through two separate networks: one connected to the management interface of each Storage Router and the other connected to the HA interface of each Storage Router. To make sure that HA information is exchanged reliably between storage routers, the storage routers balance the transmission of HA information between the management and the HA interfaces.

SR2122 access is controlled through an access list and a VLAN identifier (VID) number of the hosts. An access list enables access to storage devices attached to the Storage Router with any combination of host IP address(es), CHAP user name(s), or iSCSI name(s). An access list contains these combinations. Host VID enables access to storage devices according to the VID of each host. You can use a combination of access lists and VIDs to configure access in the Storage Router; that is, you can specify that certain hosts according to IP address in a VLAN can access storage devices attached to the Storage Router.

**Fibre Channel SAN**

A Storage Area Network is a dedicated, centrally managed, secure information infrastructure, which enables any-to-any interconnection of servers and storage. SANs are built to incorporate the best of both storage and networking interfaces: fast and efficient communications optimized for movement of large amounts of data, but with access to a wide range of other servers and storage devices on the network.

The SR2122 iSCSI Router is supported on any currently supported hp B-Series or M-Series switch fabric.

**SAN Storage**

The SR2122 supports the RA/MA8000, ESA/EMA12000, EVA, MSA, VA, and XP storage arrays.

## SR2122 Hardware and Software Support

This section lists the hardware, devices, and operating systems that are compatible with this SR2122 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 Storage Router is supported with the HP B-Series Product Line switches listed in [Table 9](#) or with the HP M-Series Product Line switches listed in [Table 10](#).

### 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
- MSCS (Windows 2000 only)
- Red Hat Linux v7.3, 8.0 Red Hat Advanced Server 2.1
- Secure Path (Windows 2000 only)

### Compaq Network Teaming Software Support

- Compaq Network Teaming (Windows 2000 only)

### SR2122 Management Software Support

The following HP management software is supported:

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

## SR2122 Configuration Rules

### SR2122 Router Rules

**Table 40: SR2122 Router Rules**

Maximum scsirouter instances per SR2122 Router (and per SR2122 Router Cluster)	12
Maximum iSCSI host connections per SR2122 SCSI Router instance	32
Maximum active targets per SR2122 Router	100

- The SR2122's 2<sup>nd</sup> fibre Channel port (FC2) is not supported as a redundant SAN port for FC1.
- Direct connect of the SR2122 FC ports to any HP storage array is not supported.
- The MSR2122 Management port must be in a different subnet than the SCSI Router Instances.

### iSCSI Host Rules

**Table 41: iSCSI Host Rules**

Maximum targets accessed per iSCSI host	Windows 2000	8
	Red Hat Linux	16
Maximum LUNs per target	Windows 2000	255
	Red Hat Linux	256

### 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.
- Multi Linux iSCSI initiators connected to a single SR2122 is not supported.

### Storage Array Rules

- The HSG80 is supported in both SCSI-3 Transparent Failover Mode and Multiple-Bus Failover Mode.
- Without hp Secure Path the Enterprise Virtual Array is supported with the SR2122 accessing only one EVA controller port. This will disable controller failover protection.
- Without hp Secure Path the MSA1000 is supported with the SR2122 accessing only one MSA controller port. This will disable controller failover protection.
- The SR2122 fibre channel ports appear as host bus adapters to the FC switches and to all storage arrays.

### Fibre Channel Switch/Fabric Rules

- The SR2122 is supported on the HP B-Series SAN Switches and HP M-Series Switches.

- The SR2122 should only be zoned with the storage devices that it will access. Zoning the SR2122 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.

### Configuring the SR2122 iSCSI Storage Router in a SAN

Configuring the SR2122 in a SAN takes on five distinct steps or sections:

- Setup/configuration of the SR2122 itself
- Adding the SR2122 to a SAN – standalone or clustered
- Serving storage to the SR2122 from the storage arrays over Fibre Channel
- Serving storage to the iSCSI Initiators from the SR2122 over IP.
- Setup/configuration of the iSCSI Initiators

For the details of setup, see the addendum [“using the iSCSI SR2122 storage router bridging from TCP/IP hosts to Fibre Channel storage.”](#)

### 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

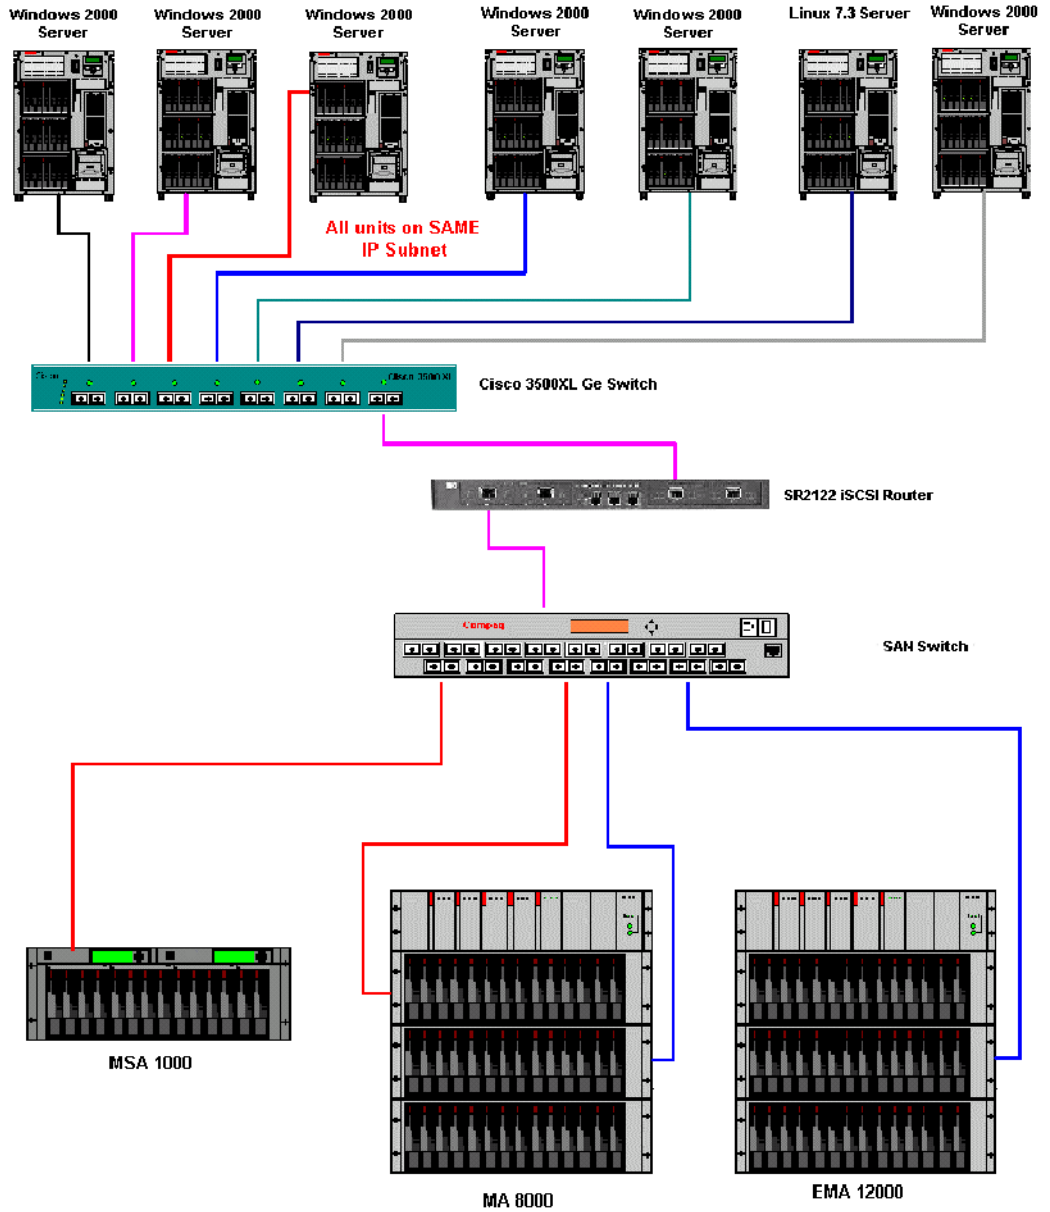


Figure 53: Example of Multiple OS Systems in a Non-Redundant Path Configuration

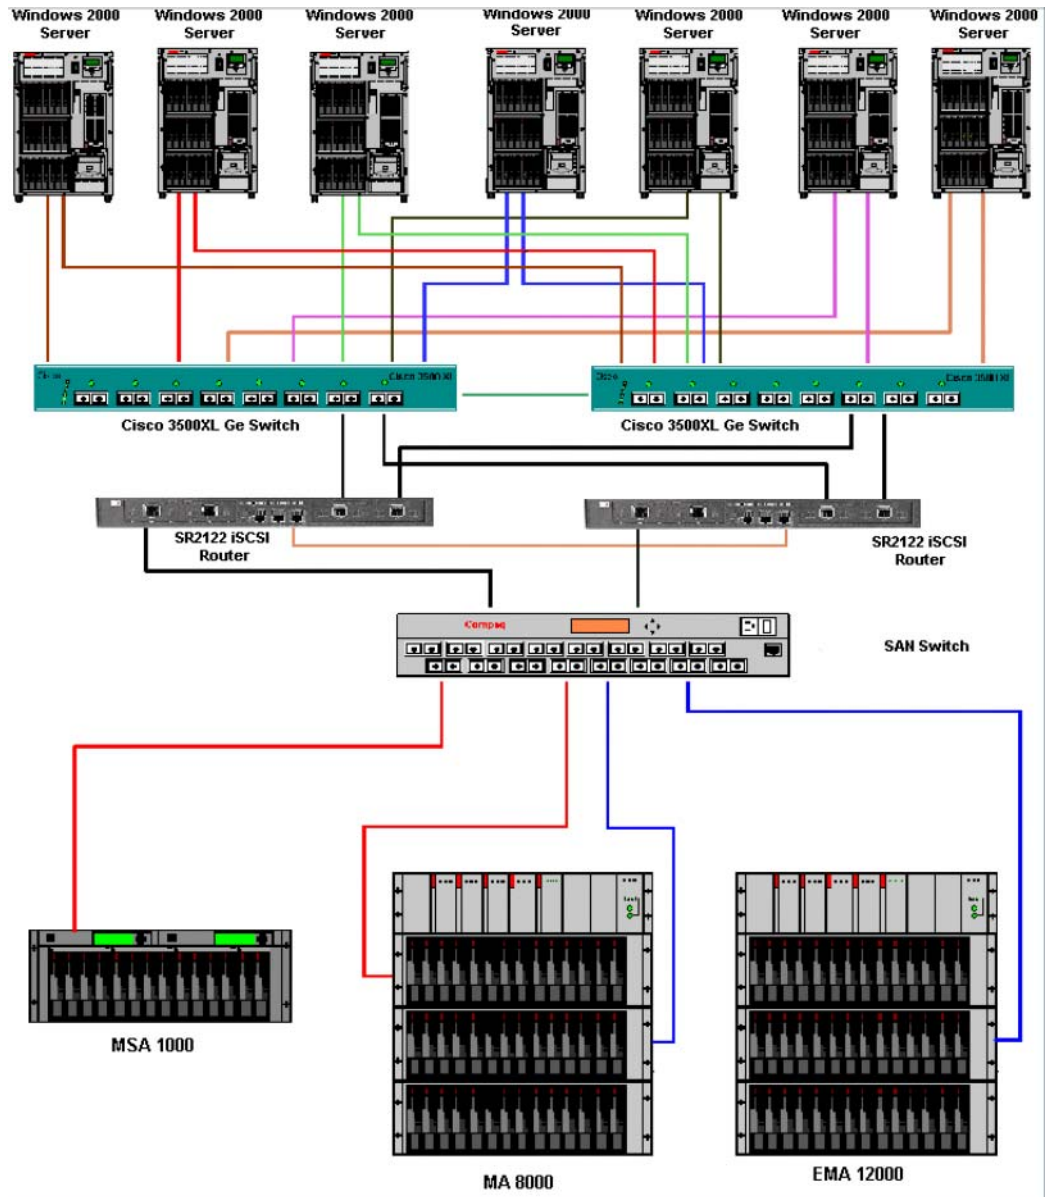
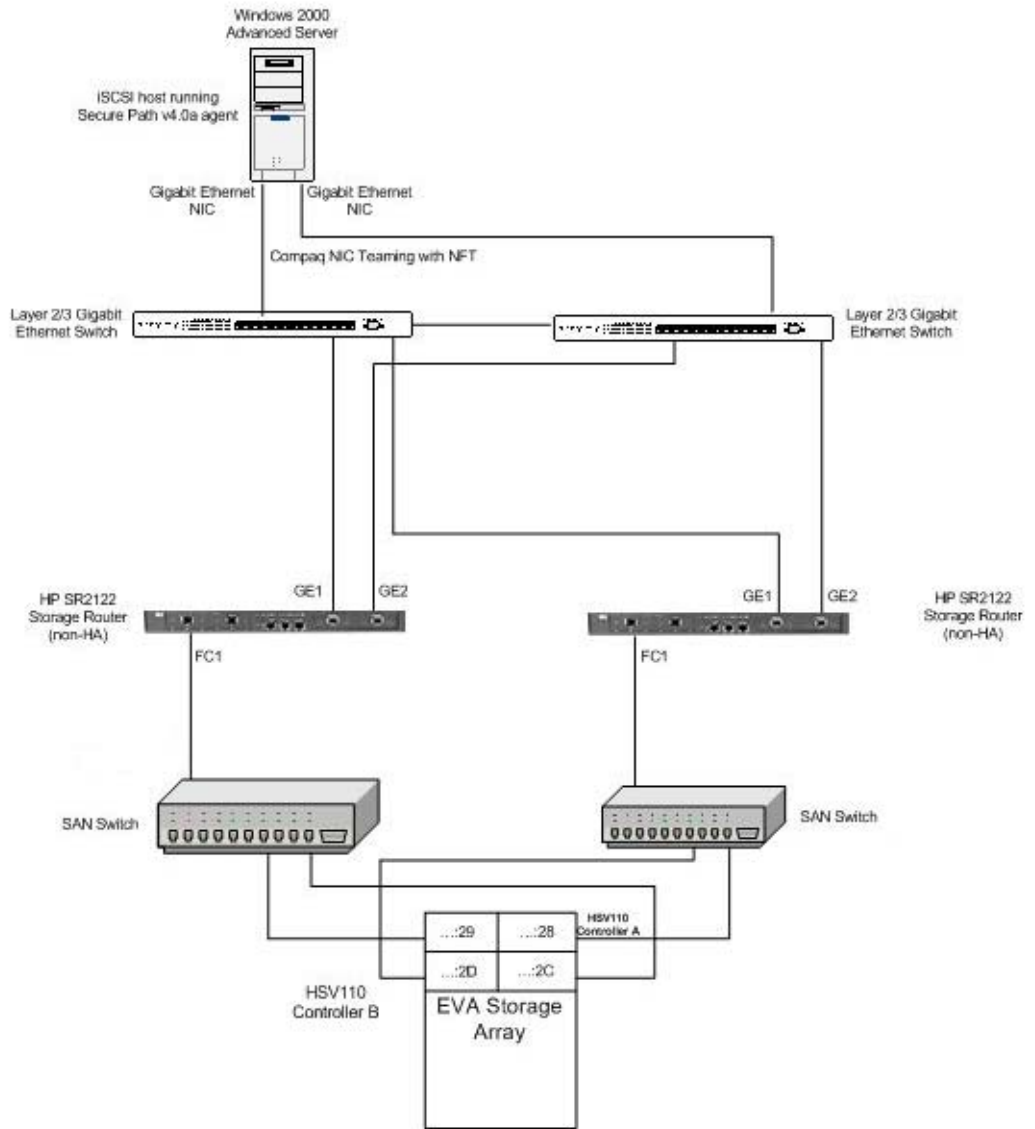


Figure 54: Windows 2000 Servers with NIC Teaming: 2 Node SR2122 Cluster



**Figure 55: Secure Path Configuration**



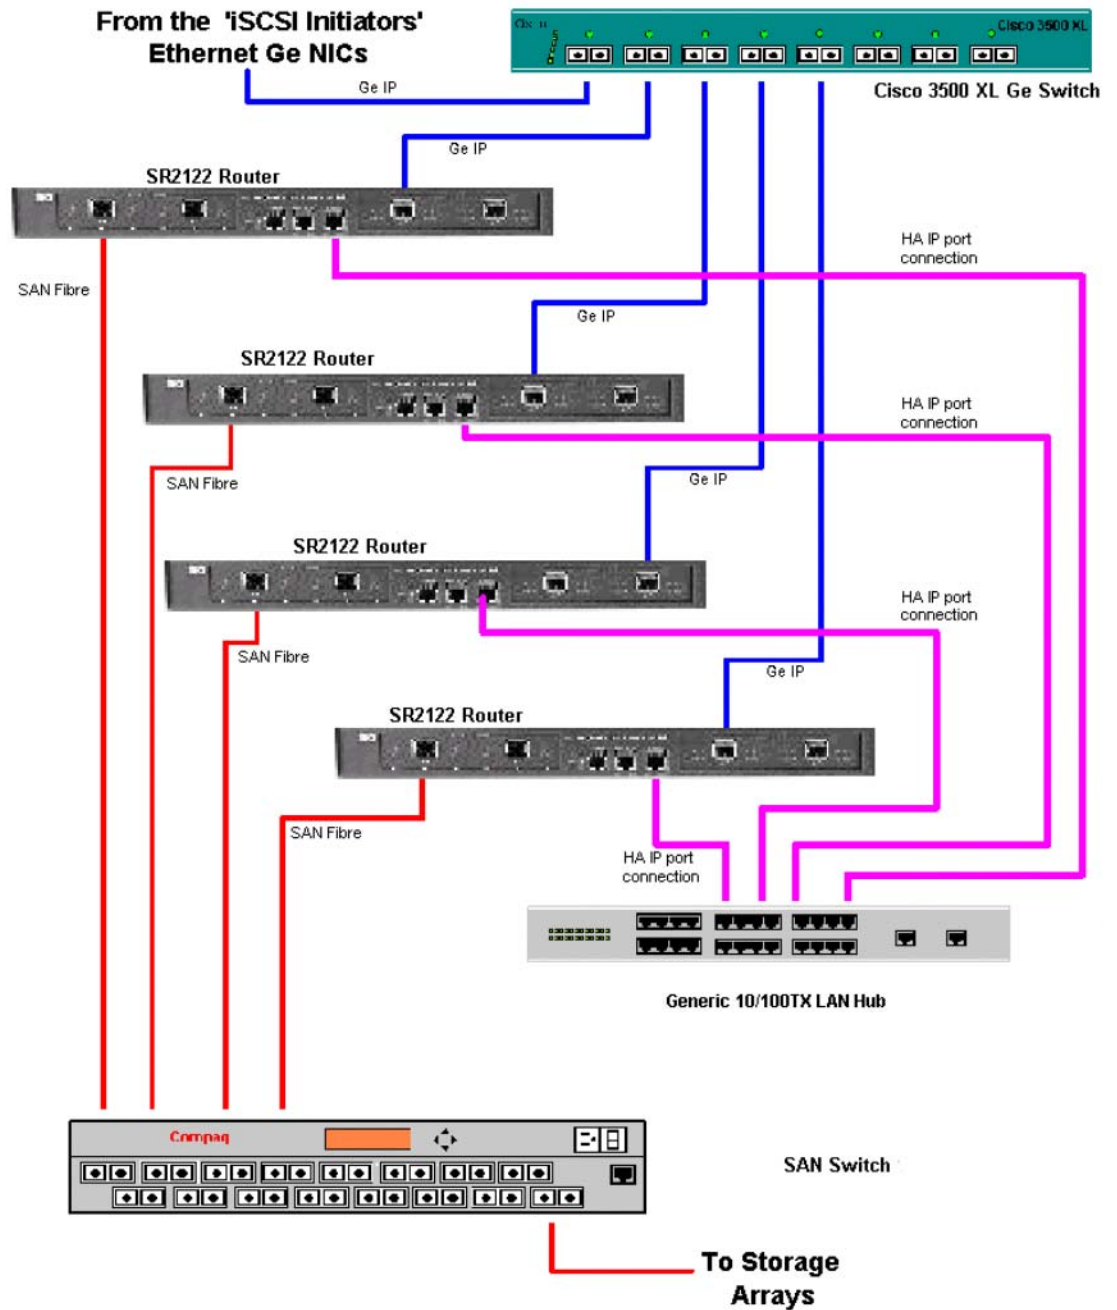


Figure 56: Maximum SR2122 Cluster Configuration Using HA Ports

