

Storage Area Network Configurations for RA8000/ESA12000 on Sun Solaris

Application Note
EK-SMA37-AN. A01

Introduction

This application note serves as a guide for Storage Area Network (SAN) configurations for Compaq StorageWorks RAID Array 8000 (RA8000 FC) and Enterprise Storage Array 12000 (ESA12000 FC) storage system(s) in conjunction with the Sun family of servers running the Solaris operating system.

Clearly, there are many levels of complexity in dealing with servers in light of performance, usage and load as well as the nature and configuration of fibre channel fabric and switches, and equally important, the actual configuration of the storage sets that are used by the servers. Each installation or site may have needs that are beyond the scope of this application note, and to that end, we recommend contacting your local Compaq Services representative for on-site consultation.

Visit Our Web Site for the Latest Information

Compaq is continually making additions to its storage solutions product line. Please check the Compaq Web site for more information on the complete line of fibre channel storage products, product certification, technical information, updates, and documentation. This information can be accessed at:

www.compaq.com/products/storageworks

NOTE

This application note describes RA8000/ESA12000 FC-Switched Fabric SAN configurations for Solaris. For Fibre Channel Arbitrated Loop configurations, refer to Application Note "RA8000/ESA12000 FC-AL Configurations for Solaris", part number EK-FCALS-AA

Enterprise Network Storage Architecture (ENSA)

Compaq StorageWorks RAID Array 8000 (RA8000 FC) and Enterprise Storage Array 12000 (ESA12000 FC) products are the latest in fully integrated fibre channel RAID storage solutions.

These products provide another key step in delivering on Compaq's ENSA vision. ENSA addresses the issues that our customers expect to face now and in the future. Both products address today's issues including: economical capacity growth, data protection, high availability, increased distance, multi-vendor platform support, and investment protection as the base for building the ENSA vision.

With the addition of fibre channel switched fabric support in the RA8000 and ESA12000 FC storage products, Compaq has taken the next evolutionary step in delivery of the ENSA vision. Our

customers can now realize the full benefits of a Storage Area Network (SAN) providing exceptional connectivity while increasing the effective bandwidth of the network. Fibre Channel Fabric provides the robust backbone needed to address dispersed server operations with shared storage access across the enterprise.

Product Overview

The RA8000 FC and ESA12000 FC storage systems are based on a common architecture and offer customers centralized management, high availability, exceptional performance and scalability in open systems environments. These products provide a Fibre Channel storage solution with industry leading Storage Area Network (SAN) support utilizing both fibre channel Switch (FC-SW) and fibre channel Arbitrated Loop (FC-AL) topology and technology.

The RA8000 FC is a mid-range storage system available in a pedestal enclosure for departmental use in the office. It can also be installed in server racks for data centers. An easy-to-deploy, flexible solution for open systems, the RA8000 currently supports up to 2.6 TB of storage with two expansion cabinets using a total of 72 disks, each with a capacity of 36GB.

The ESA12000 FC, designed for the data center, is available in three easy-to-order building blocks; each optimized for specific, high-capacity customer requirements. The ESA12000 FC offers unprecedented levels of storage, scaling in performance and capacity as user requirements increase.

These storage systems incorporate the latest in RAID technology providing RAID levels 0, 1, 0+1, adaptive 3/5, and non-RAID disks (JBOD). Compaq's implementation of RAID capability assures that data availability will be maintained despite hardware failure. Features such as read-ahead cache and mirrored write-back cache improve performance while preserving data and high availability support.

Platform kits are components of software and documentation created for each specific operating system that can be used with the Compaq StorageWorks RAID arrays. Platform kits exist for the SUN (Solaris) operating system, Hewlett-Packard (HP-UX) operating system, and other supported operating systems.

The RA8000/ESA12000 FC platform kits contain the necessary operating system specific software and documentation needed to install, configure, and monitor your storage system. All platform kits include the StorageWorks Command Console (SWCC) for storage management operations including configuring, status inquiry and notification facilities. Additionally, the SWCC provides a graphical interface, simplifying the most complex of storage management operations.

Product Details

The RA8000/ESA12000 storage systems utilize the Compaq HSG80 RAID controller running Array Controller Software (ACS) V8.5F. The ACS software is designed to support multiple platforms providing features including: dual controller operation, two controller failover modes – Transparent and Multiple-Bus, mirrored write back cache, read ahead cache, RAID storage implementation, disk mirroring, and disk partitioning capabilities. In addition, ACS manages server interconnect and protocol services to provide data for event notification and status such as displayed by SWCC.

The HSG80 controller has two FC host ports providing up to a total of nearly 200 Mbytes per second of available bandwidth. Disk drives are connected to the controller through 6 UltraSCSI channels providing up to 40 Mbytes per second per channel of available bandwidth. Servers can use multiple host bus adapters (HBAs) to multiple RA8000 FC systems for scalable storage capacity.

The RA8000 FC is supplied in two basic building blocks, a pedestal style cabinet and what has been termed a "rackable" model. The pedestal is a self-contained desk height cabinet; it is available in Digital Classic Top Gun Blue and in Compaq Opal. The rackable unit is shipped ready to be mounted in a 19" Compaq rack or a 19" cabinet/rack using an included universal RETMA/metric mounting kit.

RA8000 FC controllers come with 64 MB of cache. Cache options allow for 128 MB, 256 MB and 512MB of cache in each controller. The ESA12000 FC controllers come with 256 MB of cache. Cache options allow for 512MB of cache in each controller. Cache features include read-ahead and mirrored write-back cache for use in optimizing performance where redundant capabilities are required for high availability system configurations. In the rare event that a storage controller fails the mirrored data is immediately available to the functioning controller for continuing operation.

The RA8000 FC pedestals and rackables come in both single controller and in dual controller models to allow the customer to select the level of high availability required for their application. A single controller model can be upgraded to a dual controller model when warranted by the application.

High availability is supported by the use of redundancy and hot swappable components. The RA8000 FC disk enclosures include redundant cooling, and have five (5) power supplies providing N+1 redundancy. Power can be upgraded to full redundancy with the addition of three (3) more power supplies and a second power distribution unit for AC input redundancy. StorageWorks components are "hot swappable" including storage controllers, when operating as a redundant pair.

You can configure and monitor the RA8000 and ESA12000 FC using the StorageWorks Command Console (SWCC). SWCC is included in the HSG80 Software Solutions Platform kit. Using this client/server tool, you can configure RAID sets locally or over a TCP/IP network. The Command Console Client provides an easy to use graphical user interface. It can be used to monitor your storage system and notify you of events by updating a graphical display, sending electronic mail, or alerting via a pager. The Command Console Client includes applications (or Storage Windows) that communicate with the Command Console Agent software running on the server platform.

The Agent can also be used as a standalone application without the Client. In this mode, the Agent monitors the status of the RAID systems and provides local and remote notification in the event of a failure.

The Command Console client is available on Windows 95/98 and on Intel/Alpha platforms running Windows/NT. A Storage Window for the HSG80 controller supports the Array Controller Software (ACS) V8.5F and offers integration with Compaq's Insight Manager (CIM). Insight Manager can receive SNMP traps from the Command Console Agent.

Veritas FirstWatch

Veritas FirstWatch is high availability (HA) software supported on the Solaris operating system. FirstWatch enables two Sun servers to share an RA8000 or ESA12000 FC storage system through a fibre channel bus. If a server failure should occur, ownership of shared storage is migrated to the functioning server. HA applications are started on the functioning server and access to data is restored automatically with no downtime.

Assumptions

This application note assumes familiarity with configuring RAID Storage Arrays, specifically the use of the Command Line Interface (CLI) as well as experience and use of the Sun Solaris operating system. Additionally, an understanding of Fibre Channel terminology and concepts is assumed.

Document References
Table 1 of this document lists the various documents related to the HSG80 Controller and the configuration of the RA8000/ESA12000 RAID Arrays. Refer to the documents contained therein for more detailed help on a specific topic.

SAN Concepts

A Storage Area Network interconnects multiple shared storage environments and servers into a storage network. All storage can be managed from a single location or from multiple locations. All of the consolidated storage becomes available to any host server, regardless of physical location.

SAN Management

SAN management tools and product features provided with the RA8000/ESA12000 storage systems include the following:

- SAN Storage and Switch Management

SWCC – Utilized to manage and monitor storage systems, storagesets, and SAN fibre channel switches including configuration, status inquiry and notification facilities.

- SAN Access Management

Switch Fabric Zoning – The FC switch Zoning feature provides a means to control SAN access at the node port level. Zoning can be used to separate one physical fabric into many virtual fabrics consisting of selected server and storage ports. This capability allows you to set up barriers between different operating environments, to deploy logical Fabric subsets by creating defined server and/or storage groups, to create temporary server and storage zones. Zones may be configured dynamically and the number of zones and zone membership are effectively unlimited. Nodes can be in multiple zones to allow for overlapping depending on the desired access control. Although not required for the Solaris environment, use of Zoning is supported in both homogeneous Solaris SAN configurations and heterogeneous mixed platform SAN configurations.

Selective Storage Presentation (SSP) – The RAID Array SSP feature provides a means to control access at the storageset level. SSP is an exclusive RA8000/ESA12000 storage system feature that allows multiple servers or hosts to share access to the same RAID array safely, with each server's storagesets (LUNs) presented exclusively to the server or servers that are allowed access to that storage.

Additionally, SSP allows the setting of host modes for each FC HBA connected to the array and LUN offsets. The host mode is specially tailored to the storage communication techniques of the operating system. The LUN offset feature allows higher numbered LUNs in a RAID Array to be presented in a range required by specific operating systems. The SSP feature also provides a means to track the numerous FC HBAs within servers attached to a SAN, by identifying each FC HBA using the World Wide Identification (WWID).

Additional information about these products and features is available in the documentation listed in Table 1.

Configuration References

Table 1 lists the documents that provide further information on the configuration of RA8000 and ESA12000 Fibre Channel Storage Systems.

Table 1 Configuration Reference Material

Topic	Document Title	Order Number
<i>Heterogeneous SANs</i>	<i>Heterogeneous Storage Area Networks</i>	<i>EK-SMA30-AN</i>
<i>RA8000/ESA12000 Standard FC-AL Configurations Application Note</i>	<i>RA8000/ESA12000 FC-AL Configurations for Solaris</i>	<i>EK-FCALS-AA</i>
<i>Release Notes</i>	<i>Compaq StorageWorks Release Notes, RA8000/ESA12000 HSG80 Solution Software V8.5 for Sun Solaris</i>	<i>AA-RFBPE-TE</i>
<i>Solution Kit Overview</i>	<i>Compaq StorageWorks Solution Software Kit Overview</i>	<i>EK-SOLSR-AA</i>
<i>Sun Solaris RA8000/ESA12000 Installation</i>	<i>RA8000/ESA12000 Fibre Channel Solution Software V8.5F for Solaris – Installation Reference Guide</i>	<i>AA-RFBQE-TE</i>
<i>HSG80 Array Controller CLI Reference Guide</i>	<i>HSG80 Array Controller ACS Version 8.5 CLI Reference Guide</i>	<i>EK-HSG85-RG</i>
<i>HSG80 Array Controller Configuration</i>	<i>HSG80 Array Controller ACS Version 8.5 Configuration Guide</i>	<i>EK-HSG85-CG</i>
<i>HSG80 Array Controller Maintenance</i>	<i>HSG80 Array Controller ACS Version 8.5 Maintenance and Service Guide</i>	<i>EK-HSG84-SV</i>
<i>StorageWorks Command Console</i>	<i>Command Console V2.2 (HSG80) for RAID Array 8000/ESA12000, User's Guide</i>	<i>AA-RFA2D-TE</i>
<i>Warranty Information</i>	<i>DIGITAL StorageWorks Warranty Terms and Conditions</i>	<i>EK-HSXSU-WC</i>

General Configuration Guidelines

Fibre Channel storage technology offers more configuration flexibility than parallel SCSI storage. This application note describes typical RA8000/ESA12000 FC Sun Solaris configurations based on the configuration maximums listed below. For configurations not shown here, the configuration maximums should be used to determine the viability of a proposed configuration. Exceeding the maximum will create unsupported configurations. Additional information specific to a particular configuration is provided by section.

Configuration Maximums

The HSG80 Controller

1. Maximum 128 visible LUNS (64 per port)
2. Maximum 200 assignable unit numbers
3. Maximum 72 physical devices
4. Maximum 20 RAID-3/5 Storageesets
5. Maximum 30 (RAID-3/5 and RAID-1) Storageesets
6. Maximum 45 (RAID-3/5 and RAID-1 and RAID-0) Storageesets
7. Maximum 8 partitions of a Storageeset or individual disk
8. Maximum 6 members of a Mirrorset
9. Maximum 14 members per RAID-3/5 Storageeset
10. Maximum 24 members per Stripeset
11. Maximum 48 physical devices per Striped-Mirrorset
12. Maximum 32 Connections per Port
13. Maximum 9 Characters per Connection Name

The Solaris Operating System - Switch Support

1. Maximum 512 GB file system per LUN (Solaris Versions: 2.6, 2.7 (with appropriate patches loaded))
2. Maximum 64 LUNs per FC Target
3. Patches Refer to *RA8000/ESA12000 HSG80 Solution Software V8.5F for Sun Solaris Reference Guide (AA-RFBQE-TE)* for required patch levels.

The Fibre Channel Storage Area Networks (SAN)

1. Minimum 2 meters per optical cable segment
2. Maximum 500 meters per optical cable segment, 50 micron, multi-mode using short-wave length laser GBICS; 1.5 km total distance nominal, 2.0 km worst case (single fault re-route) between server and storage.

200 meters per optical cable segment, 62.5 micron, multi-mode using short-wave length laser GBICS; 600 meters total distance nominal, 800 meters worst case (single fault re-route) between server and storage.

10 km per Inter-Switch Link (ISL) cable segment, 9 micron, single-mode, using long-wave length laser GBICS; 11 km total distance nominal, 21 km worst case (single fault re-route) between server and storage. 10 km links are only supported in switch to switch configurations.
3. Maximum Number of ISLs between switches depend on the number of available ports, required redundancy protection and data loads from servers to storage systems. A reasonable minimum is two links for redundancy.
4. Maximum 1 Unique Domain Number for each switch in a single fabric.
5. Maximum 16 fibre channel Host Based Adapters (HBA) per server.
6. Maximum 4 fibre channel switches in a single fabric, i.e., SAN with all switches interconnected.
7. Maximum A hop is defined as one (1) or more connections between two FC switches. One (1) switch hop between servers and storage, maximum of two (2) switch hops (worst case with a single fault. A fabric can be re-routed to a 2-hop path on a single link failure. Two (2) switches cascaded is one hop. Server to FC switch segments and FC switch to storage segments are not counted as hops.
8. Maximum 8 Solaris servers (assumes one FC HBA per server) per RA8000 or ESA12000 dual controller storage system, with a maximum of 4 servers configured on each port pair.
9. Maximum 2.6 TB of storage, assuming 72 disks, 36GB each in a cabinet with two (2) expansions.

General Configuration Guidelines

1. **This application note describes configurations with dual (redundant) controllers per storage system.** All single controller storage configurations can be upgraded to dual controller configurations.
2. Controller pairs can be configured for either Dual Redundant Failover Mode *or* Multiple-bus Failover Mode. When controllers are in Dual Redundant Failover Mode, this is referred to as 'Transparent Failover Mode'.

Note: This application note describes Transparent Failover Mode configurations only.

3. All configuration diagrams show a frontal view of the storage system.
 - Controller Host Port 1 is the left port; Controller Host Port 2 is the right port.
 - Controller A is the top controller; Controller B is the bottom controller.
4. Each storage system can be configured using the Storage Works Command Console (SWCC) or the Command Line Interface (CLI) configured through either the HSG80 controller maintenance serial port or a TCP/IP interface.
5. All configurations **require** Array Controller Software (ACS) V8.5F for the FC-Switch (fabric) topology. This is fabric support at the controller level.

All configurations **require** RA8000/ESA12000 FC Solution Software V8.5 for the Sun Solaris operating system.

6. All configurations utilize GBICS of both short and long-wave length lasers, depending on the constraints noted above for distances traversed. Inter-Switch Links (ISLs) are supported with long-wave GBICs and 9 micron single mode fibre channel optical cables.
7. All configurations **require** the Connection Name *Operating System* parameter set to 'SUN'. (CLI: "SET *connection-name* OPERATING_SYSTEM = SUN")
8. By default the HSG80 Command Console LUN (CCL) is enabled. This may be disabled. (CLI: "SET *controller* NOCOMMAND_CONSOLE_LUN")
9. Controller Units: The HSG80 controllers support a range of 200 unit numbers. In Transparent Failover Mode, the units D0-D99 are assigned to Port 1 and units D100-D199 are assigned to Port 2.
10. Each active controller host port presents one SCSI Target ID with up to 64 LUNs to each server (FC HBA) on the same switch. By default, the Solaris RAID Manager Platform kit platform kit defines 16 LUNS per target. The number of targets, LUNS, etc., can be modified using the Solaris Configuration Script option 20, add/change adapters. This script can be found in:

/opt/steam/bin/config.sh (if /opt was chosen as the default base directory area)
or
your_base_directory/steam/bin/config.sh (if /opt was not chosen)
11. When configuring controllers for access from Solaris servers, the controller SCSI mode should be set to SCSI-2.

SAN Configuration Guidelines

The following sections provide information about configuring servers and storage systems in a SAN.

Thinking About Connections.

For our purposes, "Connections" are defined as the endpoints of the physical path between a server and a storage system within a SAN. Obviously, a FC-Switch may serve to connect, direct or protect a path. Yet, our focus is to treat the connection as the end points of the path.

There are two points of view when considering fibre channel SAN connections. Managing connections from these two points of view is the key to the definition, support and maintenance of the servers and systems in a SAN.

The First View

The first view of connections is taken from the server side. The server itself can be configured for specific targets and logical units in a SAN. The targets and logical units may be configured as desired within the maximum values for the operating system and restrictions of the fibre channel specifications. This means that while there may be 64 targets in a fabric, the configuration file may not contain all 64 targets. In short, a specific server may not look for or try to access a specific unit.

The Second View

The second view of connections is from the controller side. In a fabric, the controller knows about each host-adapter combination (HAC) that it can communicate with. The ACS V8.5F software provides for up to 64 server connection name entries, 32 per port.

Each storage system in the SAN will automatically add connection name entries upon initial power up, assuming that the servers are configured and ready to communicate with the storage system. In Transparent Failover Mode, the controller adds one (1) connection per server on each active port. Thus, if there are four servers (4 HACs) visible to the controller in the fabric, the controller will create four connections. If there is one active port, the controller will establish 4 connections. If there are two active ports, the controller will create 8 connections.

As a result, in certain SAN configurations it is possible to exceed the number of available connection name entries prior to the addition of entries for the specific servers that you want access to.

By combining the server point of view and the controller point of view, we have the flexibility of managing the available connections. That is, if we define all the targets in the server's configuration file and map these targets to all available connections we will have a complete use of the available connections. On the other hand, if we wish to restrict access from the server point of view, we may edit the server configuration so that there are fewer targets than connections available, leaving unused connections. Those connections are still available but for reasons of load balancing or access control, we choose not to use them for this server.

When implementing a configuration, care must be taken to specify the connections that are needed per server and per system.

Managing Connections - Server Strategies

1. Prior to installation of a FC HBA, record the World Wide Node Name (WWNN) address that is affixed to the HBA/adapter. This document will explicitly name each WWNN as either Adapter WWNN or Controller WWNN to avoid confusion.

If the HBA is already installed, or if the WWNN is not affixed to the adapter, the HBA's WWNN will be displayed on the system console when the Fibre Channel driver attaches to the adapter. The WWNN is a hexadecimal value of the form, nnnn-nnnn-nnnn-nnnn.

This first step is necessary for managing the connections at the storage system controller as these adapter WWNNs are used for identifying the specific connections.

2. The server needs to know the system(s) that it will connect to. The identification of the system is facilitated with the World Wide Node Name (WWNN) of each controller in the SAN.

At each storage system, make note of the controller WWNN from the label that is affixed to each controller.

Since each controller has two ports, each port is designated with a World Wide Port Name (WWPN). It is these WWPNs that are used by the server for communicating with the specific storage system(s). These WWPNs are used by the Solaris implementation to create persistent bindings from server to storage system.

The controller's World Wide Port Names are derived from the World Wide Node Name as follows: If the WWNN is A234-B234-C234-D1200

then the WWPNs are A234-B234-C234-D1201 for Port 1
A234-B234-C234-D1202 for Port 2

where the last digit for the WWPN is the Port identification, either '1' or '2'. Record all the necessary WWPNs associated with the controllers. These will be used for defining the Targets.

3. Visit each server in the SAN and perform the following steps:
 - Assuming that the HBAs and the Solaris Platform Kit have been installed, invoke the configuration script, `# /opt/steam/bin/config.sh`
 - When invoked, this script will allow the configuration of each HBA to the specific storage system by specifying the WWPNs. This is implemented in Option 20, Add/Change Adapter. Briefly, the specific HBA is chosen; the mode for the driver is set to fabric; the Target Ids are specified; the number of LUNs specified; finally, the WWPNs are entered for each Target.
 - The last step is performed for each HBA that is to be configured on each server. Once these values have been supplied to the script, the system files are updated and instances are created in the `/kernel/drv/sd.conf` file. An inspection of this file will show that the correct number of LUNs have been created for the specified targets (which include the WWPNs provided.)
 - When all HBAs have been configured, the new mappings and instances must be acquired by the system. The system will acquire the new mappings with a reboot of the server using the reconfigure switch, `-r`. If the storage system has been prepared with at least one unit that is defined, that device will be visible under the `# format` utility as a storage device.

Note that this process of identification of the mode, targets, LUNs and WWIDs is used for additional systems or servers as they are added to the SAN fabric.

Managing Connections - Controller Strategies

1. The HSG80 Controllers operate in two fibre channel modes, loop and fabric. These configurations assume fabric mode. We set the controller mode as follows:

If the controller (CLI: SHOW_THIS) reports that the controller is not in FABRIC MODE then do the following:

CLI: SET THIS_CONTROLLER PORT_1_TOPOLOGY = OFF_LINE

CLI: SET THIS_CONTROLLER PORT_2_TOPOLOGY = OFF_LINE

CLI: SET THIS_CONTROLLER PORT_1_TOPOLOGY = FABRIC

CLI: SET THIS_CONTROLLER PORT_2_TOPOLOGY = FABRIC

Perform the same commands for OTHER_CONTROLLER

and then

CLI: RESTART OTHER_CONTROLLER

CLI: RESTART THIS_CONTROLLER

2. When a controller acquires the server-HBA connection in the fabric, it uses the adapter's WWIN as part of the identification for the connection names. The complete list of connections can be displayed by using the CL command: "SHOW CONNECTIONS". The specific connection can then be identified by the adapter WWID that was recorded in the Server Strategies above.

Connections that are not required may be deleted.

NOTE
The CLI provides the ability to delete a connection. This deletion is good ONLY until the controller reinitializes. When a controller reboots, it rebuilds its connection table for all devices that it can see.

3. Rename the connection names. By default, new connection names are automatically added to the controller connection name table as "!NEWCONnn" (where nn is a number from 1 to 32). The controller adds a connection when a path to an adapter from each active controller host port is detected. Once the connection is identified using the FC HBA WWID, rename the connection names to a name meaningful to the specific configuration.

(CLI: "RENAME *old-name new-name*")

Connection names may contain nine (9) characters. In some cases, server name and Port value may suffice (e.g., HEARTP1, DIAMONDP2, SPADEP1, CLUBP2, etc.).

At this point, we have the connections we need/want and have identified them for ease of management and documentation.

4. Selective Storage Presentation (SSP) may be used for configurations with more than one server accessing the same storage using the same HSG80 controller port.

Perhaps the configuration has two servers accessing storage units in the same storage system from the same port. Without selective storage presentation or access protection, there would be a conflict with access to the specific units. From each server's point of view, the first unit would become the first target/lun path for I/O. To prevent this, we use offsets to declare access to

specific units by connection(name).

Offsets are the designation of the mapping of the LUN 0 from a specific server to a unit on the storage system. Our problem started with two servers accessing the same units from the same port. Thus, from server A, its LUN 0 maps to the lowest unit, for example D3. From server B its LUN 3 maps to D3. To solve this conflict problem as we set the base (LUN 0) to be all units above D40 for server A, by doing the following:

CLI: “SET *connection-name* UNIT_OFFSET = *n*”

From our example,

SET SERVERAP1 UNIT_OFFSET = 40 forces server A to map its LUN 0 to D40 and above.

The impact of this action allows server B to access the lower unit values and server A, the higher unit values on the same port.

We can define a second level of access control from server to storage system by enabling an ACCESS_PATH. An ACCESS_PATH maps a specific server to a specific unit, disallowing all other server connections.

The steps are:

- Disable/Clear all current access_paths
(CLI: “SET *unit-number* DISABLE_ACCESS_PATH = ALL”)
- Set a specific server access (through the connection name)
(CLI: “SET *unit-number* ENABLE_ACCESS_PATH = *connection-name*”)

The impact of this action allows only server, *connection_name* to access this unit. No other servers have access, even if enabled at the server.

Configuration Descriptions

Table 2 lists four (4) SAN configuration examples for RA8000 and ESA12000 storage systems in a FC-Fabric topology when using Solaris platforms. The configurations are based on the maximums listed in the General Configuration Guidelines section. They show the limits with regard to:

- Number of servers and storage systems per FC switch configuration
- Resilient Fabric – Alternate Fabric paths
- Controller in Transparent Failover Mode

The configurations are listed in order based on the number of switches in the SAN. This follows a logical progression of an initial configuration of a small SAN containing a single FC switch, the addition of a second switch for high availability or increased connectivity, and the addition of more switches, as capacity needs increase.

All configurations allow the flexibility to trade-off server and storage system count based on FC switch port availability. All configurations support a mix of Sun Servers. This mix includes the m, d and u architectures as well as the Solaris OS versions of 2.6 and 2.7. All configurations are supported with 16-port and 8-port FC switches.

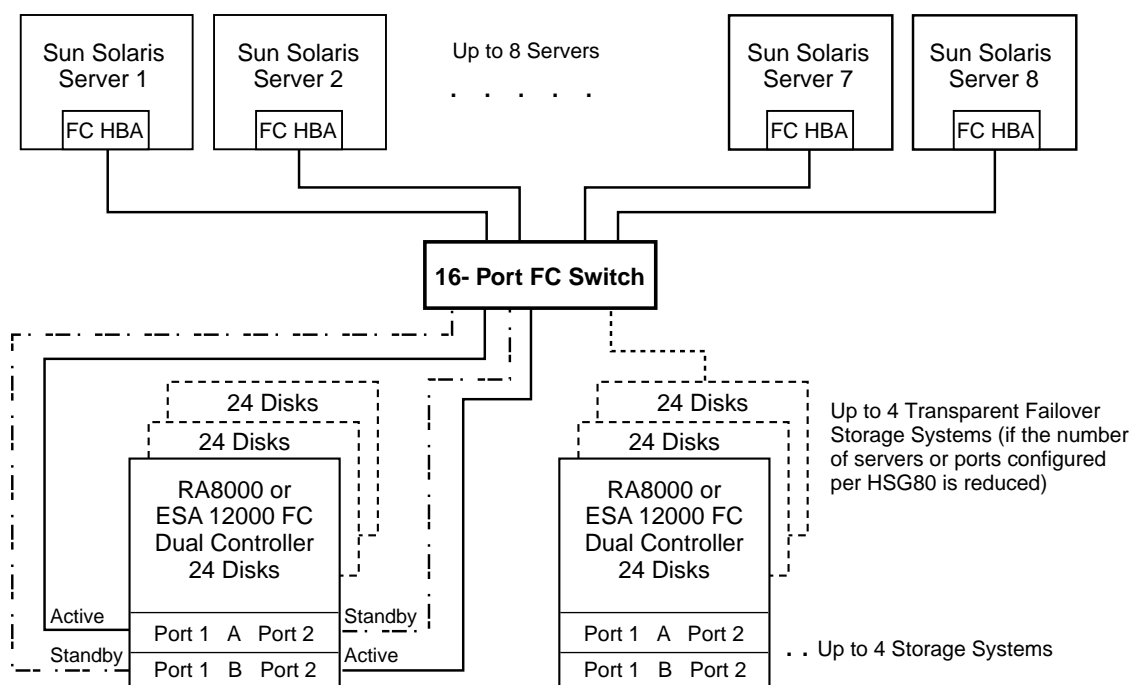
Table 2 RA8000 / ESA12000 FC-Switch SAN Configurations Solaris - Index

SAN Configuration #	Description	Pages
1	One FC Switch - Up to 8 Servers or Up to 4 Storage Systems, Controller Transparent Failover	15 - 16
2	Two FC Switches, Cascaded - Up to 16 Servers or Up to 7 Storage Systems, Controller Transparent Failover	17 – 18
3	Three FC Switches, Cascaded Meshed Fabric - Up to 24 Servers or Up to 10 Storage Systems, Controller Transparent Failover	19 – 20
4	Four FC Switches, Cascaded Meshed Fabric - Up to 24 Servers or Up to 13 Storage Systems, Controller Transparent Failover	21 – 22

SAN Configuration 1 - One FC Switch - Up to 8 Servers or Up to 4 Storage Systems, Controller Transparent Failover

SAN Configuration 1 (Figure 1) shows the maximum server and storage connectivity available through a single FC switch. Up to 8 separate servers can be connected through a 16-port switch providing access from all servers to a single storage system. With an 8-port FC switch up to 4 servers can be connected to one (1) storage system. With either switch the number of storage systems can be increased if the number of servers is reduced to provide switch ports for storage. The HSG80 array controllers within the RA8000 or ESA12000 storage enclosure are configured in Transparent failover mode providing full controller redundancy in the unlikely event of either controller failing.

Figure 1 One Switch SAN Configuration Example



SHR-1592

SAN Configuration 1 Server/Storage Connectivity Rules

- Any mix of Solaris servers using these maximums (assumes 1 FC HBA per server and an even storage port count)

16-port FC Switch:

- Up to 8 servers and 1 storage system (4 storage ports)
[16 – 4 storage ports = 12 server ports] or
- Up to 4 storage systems (14 storage ports) and 2 servers
[16 – 2 server ports = 14 storage ports]
- Maximum of 4 servers configured on one pair of controller ports

8-port FC Switch:

- Up to 4 servers and 1 storage system (4 storage ports)
[8 – 4 storage ports = 4 server ports] or
- Up to 2 storage systems (6 storage ports) and 2 servers
[8 – 2 server ports = 6 storage ports]

SAN Configuration 1 Fabric Rules

- All HSG80 array controllers configured in Transparent Failover Mode
- Total switches = 1, 16-port or 8-port

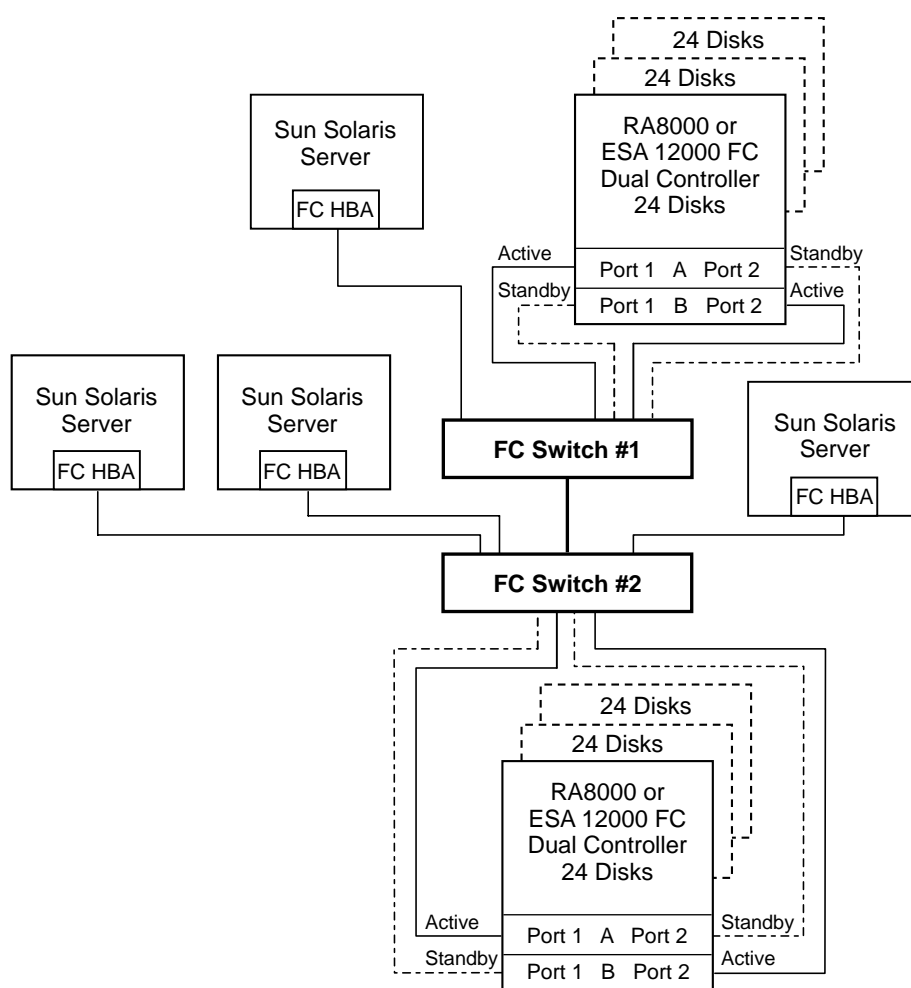
- Maximum switch hops (nominal) = 0
- Maximum switch hops (worst case) = 0
- Maximum distances:

Up to 500 meters per FC cable segment, 1 km total server to storage using 50 micron multi-mode fibre optic cable

SAN Configuration 2 – Two FC Switches, Cascaded - Up to 16 Servers or Up to 7 Storage Systems, Controller Transparent Failover

SAN Configuration 2 (Figure 2) shows the maximum server and storage connectivity available through two (2) cascaded FC switches. Up to 16 separate servers can be connected through a 16-port switch providing access from all servers to two (2) storage systems. With an 8-port FC switch up to 8 servers can be connected to one (1) storage system. With either switch the number of storage systems can be increased if the number of servers is reduced to provide switch ports for storage. A minimum of one (1) Inter-Switch Link (ISL) is required between the 2 switches. More ISLs may be required based on the specific server and storage configuration counts as well as the applications utilized. The HSG80 array controllers within the RA8000 or ESA12000 storage enclosure are configured in Transparent Failover Mode, providing full controller redundancy in the unlikely event of either controller failing.

Figure 2 Two Switches, Cascaded SAN Configuration Example



SHR-1593

SAN Configuration 2 Server/Storage Connectivity Rules

- Any mix of Solaris servers and storage systems using these maximums (assumes 1 FC HBA per server and an even storage port count)

16-port FC Switches:

Up to 16 servers and 2 storage systems (8 storage ports)

[32 – 2 ISLs – 8 storage ports = 22 server ports] or
Up to 7 storage systems (28 storage ports) and 2 servers
[32 – 2 ISLs – 2 server ports = 28 storage ports]
Maximum of 8 servers configured on 1 storage system (one controller pair)
Maximum of 4 servers configured on one pair of controller ports

8-port FC Switches:

Up to 8 Servers and 1 storage system (4 storage ports)
[16 – 2 ISLs – 4 storage ports = 10 server ports] or
Up to 3 storage systems (12 storage ports) and 2 servers
[16 – 2 ISLs – 2 server ports = 12 storage ports]
Maximum of 8 servers configured on one storage system (one controller pair)
Maximum of 4 servers configured on one pair of controller ports

SAN Configuration 2 Fabric Rules

- All HSG80 array controllers configured in Transparent Failover Mode
- Total switches = 2, 16-port or 8-port
- Maximum switch hops (nominal) = 1
- Maximum switch hops (worst case) = 1
- Maximum number of ISLs between the two switches is dependent on the number of servers/storage systems
- Maximum distances:

Up to 500 meters per FC cable segment, 1.5 km total server to storage when using 50 micron multi-mode fibre optic cable

Up to 10 km per inter-switch link, 11 km total server to storage when using 9 micron single mode fibre optic cable and long wave GBICs

SAN Configuration 2 Notes

- Each FC Switch must have a unique domain number (Domain_ID)
- Use these general performance rules to determine the optimum number of ISLs required between both switches.

For the highest available performance, whenever possible, devices that exchange the highest amount of data should be connected to the same FC switch, for example, servers and the storage assigned to them should be configured on the same FC switch, otherwise:

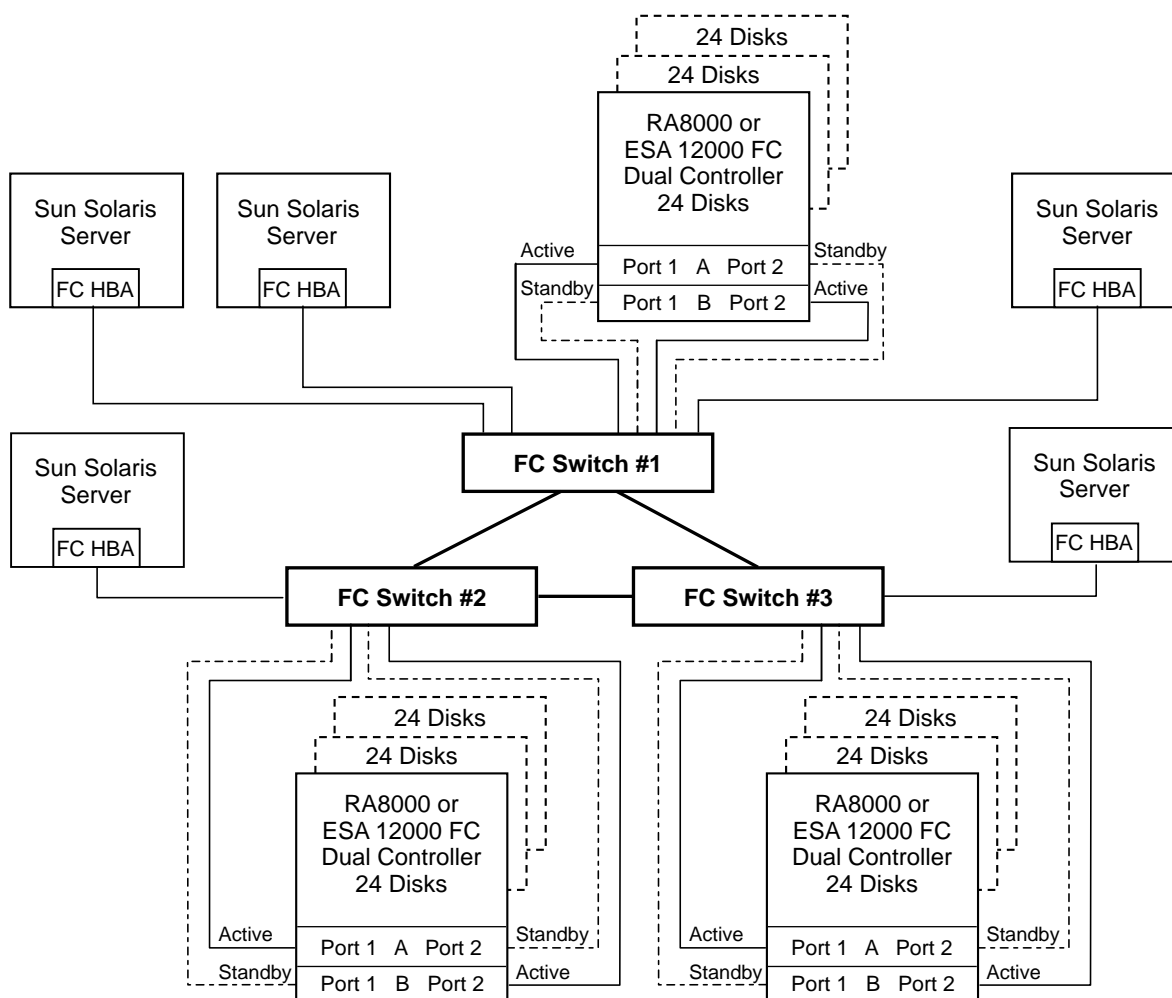
For high bandwidth applications – One ISL between switches for every 2 storage controller ports on one switch being accessed by a server on the other switch

For high throughput applications – One ISL between switches for every 6 storage controller ports on one switch being accessed by a server on the other switch

SAN Configuration 3 – Three FC Switches, Cascaded Meshed Fabric, - Up to 24 Servers or Up to 10 Storage Systems, Controller Transparent Failover

SAN Configuration 3 (Figure 3) shows the maximum server and storage connectivity available through three (3) cascaded FC switches in a meshed fabric configuration. Up to 24 servers can be connected through three (3) 16-port FC switches to three (3) storage systems. With 8-port FC switches up to 8 servers can be connected to one (1) storage system. This SAN configuration provides the fault resilient features of a meshed fabric, which automatically reconfigures a route between server and storage, should a component in the fabric path fail.

Figure 3 Three Switches, Cascaded Meshed Fabric SAN Configuration Example



SHR-1594

SAN Configuration 3 Server/Storage Connectivity Rules

- Any mix of Solaris servers and storage systems using these maximums (assumes 1 FC HBA per server).

16-port FC Switches:

- Up to 24 servers and 3 storage systems (12 storage ports)
[48 – 6 ISLs – 12 storage ports = 30 server ports] or
- Up to 8 storage systems (40 storage ports) and 2 servers
[48 – 6 ISLs – 2 server ports = 40 storage ports]

- Maximum of 8 servers configured on one controller pair
- Maximum of 4 servers configured on one pair of controller ports

8-port FC Switches:

- Up to 8 servers and 1 storage system (4 storage ports)
[24 – 6 ISLs – 4 storage ports = 14 server ports] or
- Up to 4 storage systems (16 storage ports) and 2 servers
[24 – 6 ISLs – 2 server ports = 16 storage ports]
- Maximum of 4 servers configured on one pair of controller ports

SAN Configuration 3 Fabric Rules

- All HSG80 array controllers configured in Transparent Failover Mode
- Total switches = 3, 16-port or 8-port
- Maximum switch hops (nominal) = 1
- Maximum switch hops (worst case) = 2
- Maximum number of ISLs between any two switches is dependent on the number of servers/storage systems
- Maximum distances:

Up to 500 meters per FC cable segment, 1.5 km total server to storage (2.0 km worst case) when using 50 micron multi-mode fibre optic cable

Up to 10 km per inter-switch link, 11 km total server to storage (21 km worst case) when using 9 micron single mode fibre optic cable and long wave GBICs

SAN Configuration 3 Notes

- Each FC Switch must have a unique domain number (Domain_ID)
- Use these general performance rules to determine the optimum number of ISLs required between both switches.

For the highest available performance, whenever possible, devices that exchange the highest amount of data should be connected to the same FC switch, for example, servers and the storage assigned to them should be configured on the same FC switch, otherwise:

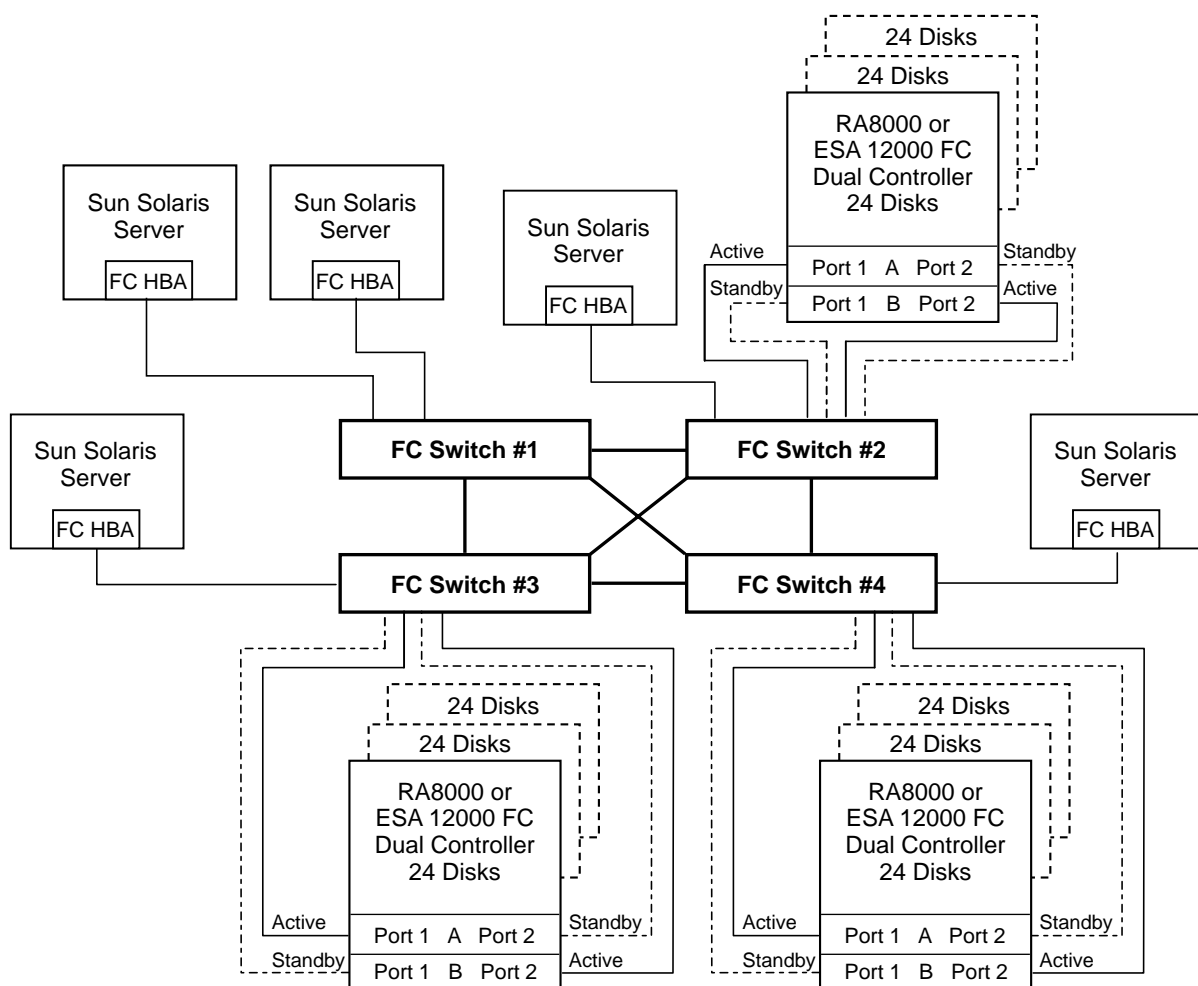
For high bandwidth applications – One ISL between switches for every 2 storage controller ports on one switch being accessed by a server on the other switch

For high throughput applications – One ISL between switches for every 6 storage controller ports on one switch being accessed by a server on the other switch

SAN Configuration 4 – Four FC Switches, Cascaded Meshed Fabric, - Up to 24 Servers or Up to 13 Storage Systems, Controller Transparent Failover

SAN Configuration 4 (Figure 4) shows the maximum server and storage connectivity available through four (4) cascaded FC switches in a meshed fabric configuration. Up to 24 servers can be connected through four (4) 16-port FC switches to three (3) storage systems. With 8-port FC switches up to 8 servers can be connected to one (1) storage system. This SAN configuration provides the fault resilient features of a meshed fabric, which automatically reconfigures a route to a good path, should a component in the fabric path fail.

Figure 4 Four Switches, Cascaded Meshed Fabric SAN Configuration Example



SHR-1595

SAN Configuration 4 Server/Storage Connectivity Rules

- Any mix of Solaris servers and storage systems using these maximums (assumes 1 FC HBA per server).

16-port FC Switches:

- Up to 24 servers and 3 storage systems (12 storage ports)
[64 – 12 ISLs – 12 storage ports = 40 server ports] or
- Up to 13 storage systems (50 storage ports) and 2 servers
[64 – 12 ISLs – 2 server ports = 50 server ports]

Maximum of 8 servers configured on one controller pair

Maximum of 4 servers configured on one pair of controller ports

8-port FC Switches:

Up to 8 servers and 1 storage system (4 storage ports)

[32 – 12 ISLs – 4 storage ports = 16 server ports] or

Up to 5 storage systems (18 storage ports) and 2 servers

[32 – 12 ISLs – 2 server ports = 18 server ports]

Maximum of 4 servers configured on one pair of controller ports

SAN Configuration 4 Fabric Rules

- All HSG80 array controllers configured in Transparent Failover Mode
- Total switches = 4, 16-port or 8-port
- Maximum switch hops (nominal) = 1
- Maximum switch hops (worst case) = 2
- Maximum number of ISLs between any two switches is dependent on the number of servers/storage systems
- Maximum distances:

Up to 500 meters per FC cable segment, 1.5 km total server to storage (2.0 km worst case) when using 50 micron multi-mode fibre optic cable

Up to 10 km per inter-switch link, 11 km total server to storage (21 km worst case) when using 9 micron single mode fibre optic cable and long wave GBICs

SAN Configuration 4 Notes

- Each FC Switch must have a unique domain number (Domain_ID)
- Use these general performance rules to determine the optimum number of ISLs required between both switches.

For the highest available performance, whenever possible, devices that exchange the highest amount of data should be connected to the same FC switch, for example, servers and the storage assigned to them should be configured on the same FC switch, otherwise:

For high bandwidth applications – One ISL between switches for every 2 storage controller ports on one switch being accessed by a server on the other switch

For high throughput applications – One ISL between switches for every 6 storage controller ports on one switch being accessed by a server on the other switch

Parts List

- Servers with Solaris Version 2.6, 2.7

Compaq Part #	Description
380554-001	RA8000/ESA12000 FC Solution Software V8.5 for Sun Solaris (QB-6BUAA-SA)
128697-B21	ACS V8.5 Controller Software
380575-001 123504-001 123505-001	SBUS FC HBA (32bit) PCI FC HBA (32bit) SBUS FC HBA (32/64 bit)
380560-B21 (Blue) 380560-B22 (Opal) 380670-B21 380580-001 (Blue) 380580-002 (Opal) 380590-B21 (Blue) 380590-B22 (Opal) 380600-001 (Blue) 380600-002 (Opal) 380610-B21 (Blue) 380610-B22 (Opal) 380620-001 (Blue) 380620-002 (Opal) 380630-B21 (Blue) 380630-B22 (Opal)	RA8000 Pedestal w/dual HSG80 RA8000 Pedestal w/dual HSG80 RA8000 Rackable w/dual HSG80 ESA12000 w/dual HSG80 24 Slot 60HZ ESA12000 w/dual HSG80 24 Slot 60HZ ESA12000 w/dual HSG80 24 Slot 50HZ ESA12000 w/dual HSG80 24 Slot 50HZ ESA12000 w/dual HSG80 48 Slot 60HZ ESA12000 w/dual HSG80 48 Slot 60HZ ESA12000 w/dual HSG80 48 Slot 50HZ ESA12000 w/dual HSG80 48 Slot 50HZ ESA12000 w/2 pairs/dual HSG80 48 Slot 60HZ ESA12000 w/2 pairs/dual HSG80 48 Slot 60HZ ESA12000 w/2 pairs/dual HSG80 48 Slot 50HZ ESA12000 w/2 pairs/dual HSG80 48 Slot 50HZ
380570-B21 (Blue) 380570-B22 (Opal) 380568-B21 380640-001 (Blue) 380640-002 (Opal) 380650-B21 (Blue) 380650-B22 (Opal)	Pedestal Expansion 24 slots Pedestal Expansion 24 slots Rackable Expansion 24 slots ESA12000 Expansion 48 Slot 60HZ ESA12000 Expansion 48 Slot 60HZ ESA12000 Expansion 48 Slot 50HZ ESA12000 Expansion 48 Slot 50HZ
380578-B21 (Blue) 380578-B22 (Opal) 380591-B21 (Blue) 380591-B22 (Opal) 158223-B21 158225-B21 158222-B21 158224-B21	FC Switch 16 port (no GBICs) FC Switch 16 port (no GBICs) FC Switch 8 port (no GBICs) FC Switch 8 port (no GBICs) FC Switch 16 port - fabric FC Switch 16 port - quickloop FC Switch 8 port - fabric FC Switch 8 port - quickloop
380561-B21 127508-B21	FC Optical GBIC (Short wave) FC Optical GBIC (Long wave)
234457-B21 234457-B22 234457-B23 234457-B24 234457-B25	FC 2 Meter Optical Cable FC 5 Meter Optical Cable FC 15 Meter Optical Cable FC 30 Meter Optical Cable FC 50 Meter Optical Cable
380691-B21 380595-B21 380694-B21 380588-B21 380589-B21 147599-001 159138-001	4GB UW 7200 RPM Disk 9GB UW 7200 RPM Disk 18GB UW 7200 RPM Disk 9GB UW 10000 RPM Disk 18GB UW 10000 RPM Disk 36GB UW 7200 RPM Disk 36GB UW 10000 RPM Disk

