Storage Area Network Configurations for RA8000/ESA12000 on HP-UX

COMPAQ STORAGEWORKS

Application Note
EK-SMA32-AN. A01

Visit Our Web Site for the Latest Information

At Compaq we are continually making additions to our storage solutions product line. Please check our web site for more information on the complete line of Fibre Channel storage products, product certification, technical information, updates to this document, and other available documentation. This information can be accessed through our web page at:

www.compaq.com/storageworks

Introduction

This application note describes Storage Area Network (SAN) configurations for RA8000/ESA12000 storage systems on HP-UX platforms. For mixed platform heterogeneous SAN configurations refer to the "Heterogeneous Storage Area Networks" application note, EK-SMA30-AN, and the applicable platform specific application notes listed in Table 4.

Enterprise Network Storage Architecture (ENSA)

The Compaq Enterprise Network Storage Architecture is key to supporting Compaq's NonStop ™ eBusiness strategy; through ENSA, Compaq leverages industry standards to allow deployment of storage where applications need it. ENSA uses the Compaq StorageWorks product family to deliver the storage solutions that address non-stop computing requirements like availability, reliability, performance, scalability, manageability, and data backup/restore capabilities. ENSA addresses the storage issues that our customers expect to face now and in the future. Compaq SANs address today's issues including:

- Data protection
- High availability
- Increased distance
- High connectivity
- High bandwidth
- Multi-vendor platform support
- Economical capacity growth
- Scalability
- Investment protection

With the addition of multi-switch fibre channel Fabric support in the RA8000/ESA12000 FC storage products and the integration of Enterprise Backup Solutions into the SAN, Compaq has taken the next evolutionary step in delivery of the ENSA vision. Our customers can now realize the full benefits of a scalable Storage Area Network providing the robust backbone needed to address dispersed server operation with shared storage access and backup across the enterprise.

Product Overview

The RA8000 and ESA12000 FC storage systems are based on a common architecture and offer customers centralized management, high availability, and exceptional performance and scalability in open systems environments. These products provide a Fibre Channel storage solution with industry leading SAN support utilizing both FC-Fabric (switch) and FC-AL (arbitrated loop) 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 small data centers. An easy-to-deploy, flexible solution for open systems, the RA8000 supports up to 2.6 TB of storage.

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. 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 or minimize the affect on performance while preserving data availability and supporting high availability.

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. SWCC provides a user oriented graphical user interface (GUI) simplifying the most complex of storage management operations.
Controllers in Multiple-Bus failover mode monitor each other and automatically failover storage units from the failed member of a controller pair. PV Links detects the failure of I/O operations to complete on a failed path and automatically re-routes all traffic to the surviving path.

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 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 setup barriers between different operating environments, to deploy logical Fabric subsets by creating defined server and/or storage groups, or to create temporary server and storage zones for tape backup. 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. Use of Zoning is supported in both homogeneous HP-UX 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 host bus adapters (HBAs) to share access to the same RAID array safely, with each server's or HBA's storagesets logical units (LUNs) presented exclusively to those that are allowed access. 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 by WWID.

Additional information about these products and features is available in the documentation listed in Table 4.

SAN Performance Considerations

The performance of an application on a system that utilizes RA8000 or ESA12000 FC storage can be limited by many different components of the system and the configuration of the SAN. Some of the possible component limiting factors include the host CPU(s), memory size, host bus adapters (HBA), RAID controllers, or the specific configuration of disks used behind the controllers. At the SAN level, performance can be limited by the number and arrangement of FC switches and inter-switch links in the Fabric, and the way servers and storage systems are connected to the Fabric. The limiting factor can move to any of these areas depending on the workload. Identifying the limits will assist in determining the best configuration for a given application.

Table 1 lists the upper limit performance specifications for the components of the RA8000 and ESA12000 FC storage systems based on testing using standard storage performance tests and methodologies. These numbers should be used to compare component level performance as a means to determine the best configuration from a performance perspective. User application tests may not necessarily reach these levels of performance as applications may perform additional levels of processing for each I/O. The controller specifications listed show both cache (no disk access) and media (with disk access) limitations.

The limits are based on I/O performance (I/Os per second) - typical of small transfer applications such as databases and mail, and bandwidth performance (MBs per second) - typical of large transfer applications such as video and graphics.

Table 1 Performance Limits of FC-SAN Components

Transfer Size	IO/sec (small transfer sizes, random access)		MB/sec (large transfer sizes, sequential access)	
Operation Type	**Read**	**Write**	**Read**	**Write**
HP FC HBA – A3740A	**21,837**	**18,715**	**97**	**78**
FC Switch (see note) **FC Switch ISL (see note)**	**400,000 25,000**	**400,000 25,000**	**1600 100**	**1600 100**
Single HSG80 Controller (1 Active Port)	**12,000 Cache 4,100 Media**	**9,000 Cache 4,000 Media**	**77 Cache 54 Media**	**50 Cache 44 Media**
Single HSG80 Controller (2 Active Ports)	**12,000 Cache 4,100 Media**	**9,000 Cache 4,000 Media**	**98 Cache 54 Media**	**100 Cache 47 Media**
Dual HSG80 Controllers (2 Active/2 Standby Ports)	**24,000 Cache 8,200 Media**	**18,000 Cache 8,000 Media**	**154 Cache 102 Media**	**101 Cache 88 Media**
Dual HSG80 Controllers (4 Active Ports)	**24,000 Cache 8,200 Media**	**18,000 Cache 8,000 Media**	**195 Cache 102 Media**	**178 Cache 88 Media**

NOTE

Fibre channel switch and ISL performance limits are theoretical. Actual measured performance is lower due to limits unrelated to the switch.

SAN Configuration Guidelines

The following sections provide information about configuring SANs including the base Fabric design selection, fibre channel Fabric rules, platform and disk storage connectivity rules, The configuration information is presented in the form of rules to provide for flexibility and customization depending on the specific customer need. In addition to the listed rules, examples of recommended SAN configurations based on the number and arrangement of fibre channel switches in a Fabric/Quickloop are shown. The example "SAN Configurations" illustrate proper application of the listed rules with regard to fibre channel switch arrangement, switch interconnection in a Fabric, and platform/storage connections. The exact SAN Fabric/QuickLoop configuration used for your specific application should begin with one of the base SAN configurations shown. In choosing a Fabric configuration you should consider capacity, connectivity, availability, distance, backup, and performance needs as well as future growth requirements.

It is not required that you configure server and storage in the SAN exactly as depicted in the illustrations. However, it is required that the switch interconnection rules and platform connectivity rules be strictly adhered to.

This application note describes homogeneous HP-UX SAN configurations. It is expected that customers may desire to implement a heterogeneous SAN consisting of many different operating system platforms. Whether implementing a homogeneous or heterogeneous SAN, it is necessary to adhere to the platform specific rules and maximums for each given platform within the SAN.

SAN Design Selection Process

It is suggested that the following general steps be followed when initially designing a SAN.

1. Considering capacity, connectivity, availability, distance, performance, and backup requirements, select the SAN configuration that best fits your needs. Refer to the configuration illustrations, notes, Fabric rules, and platform server and storage connectivity rules. Using the maximum server and maximum storage counts shown as the upper bound limits, determine the specific number of servers or storage required number for your particular installation.
   - If you want a lower server count than the maximum listed: You may increase the storage count, but only to the upper limit indicated for maximum storage, and provided you do not exceed the platform limits listed in Table 2 for HBAs/Server and Controller Ports/HBA.
   - If you want a lower storage count than the maximum listed: You may increase the server count, but only to the upper limit indicated for maximum servers, and provided you do not exceed the platform limits listed in Table 2 for Servers/Storage System.
   - If capacity or connection needs are greater than provided in a 4 switch Fabric, consider implementing multiple 4-switch Fabrics (future releases will support inter-connection of multiple Fabrics into a single Fabric to provide for growth and scaling of the SAN).
2. If you are configuring a homogeneous HP-UX SAN, refer to the HP-UX Server/Storage configuration rules in this document to determine how best to configure servers and storage. Line 1 in Table 2 lists the configuration maximums for HP-UX platforms in a homogeneous SAN configuration.
3. If you are configuring a heterogeneous SAN, refer to the appropriate platform application note(s) and the heterogeneous SAN application note to ensure platform interaction rules are understood.

   Line 2 in Table 2 lists the configuration maximums for HP-UX platforms in a heterogeneous SAN configuration.

4. Select the desired method(s) of SAN management and access control based on your specific needs. Use SWCC for storageset management, and controller based SSP or switch based Zoning (or both) for disk storage access control. SAN Configuration Rules and Maximums

The following sections list the SAN design rules as they apply to Fabric configurations, HP-UX platform specific attachment of servers and storage, and high availability storage.

Fabric Rules

- Up to 4 fibre channel switches total in a single Fabric (Two separate Quick Loops)- a SAN with all switches interconnected. All Compaq FC 16-port and 8-port switch models are supported inter-mixed. Note: A Quick Loop consists of up to two switches within a single fabric.

- Within a single Fabric, maximum of 1 switch hop between servers and storage, maximum of 2 switch hops worst case with a single fault - a Fabric can re-route to a 2-hop path on a single link failure. A hop is defined as 1 or more connections between two FC switches, for example; 2 switches cascaded = 1 hop. Server to FC switch segments and Storage to FC switch segments are not counted as hops
- Within a single Fabric where switches are interconnected, each FC switch must have a unique domain number (Domain_ID)
- Up to 15 inter-switch links (ISLs) between any 2 switches
- Minimum cable segment length is 2 meters
- Up to 200 meters per cable segment using short-wavelength laser GBICs and 62.5 micron multi-mode fiber optic cable. With multiple cable segments 600 meters total distance nominal, 800 meters worst case with a single fault re-route, between server and storage

- Up to 500 meters per cable segment using short-wavelength laser GBICs and 50u multi-mode fibre optic cable. With multiple cable segments 1.5 km total distance nominal, 2.0 km worst case with a single fault re-route, between server and storage
- Up to 10 km per ISL cable segment using long-wavelength laser GBICs and 9u single-mode fibre optic cable. With multiple cable segments 11 kilometers total distance nominal, 21 km worst case with a single fault re-route, between server and storage. 10 km links are only supported in switch to switch connections

Figure 1 One Server to Multiple 4-Switch Fabrics/QuickLoops

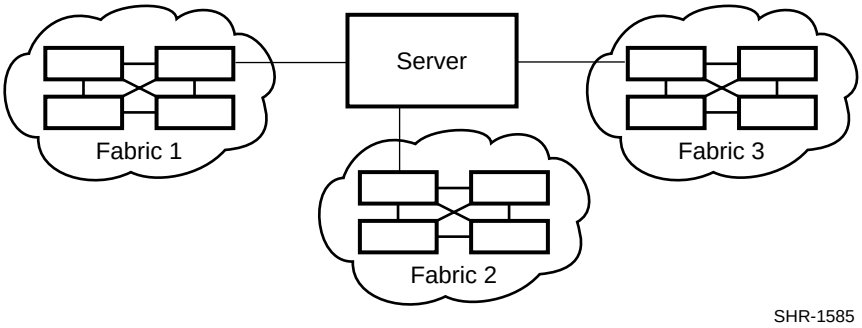

HP-UX Server/Storage Rules

Table 2 describes the HP-UX platform maximums when in a homogenous HPUX SAN and when using HPUX platforms in a heterogeneous mixed platform SAN.

Table 2 Guidelines for HP-UX Platform Supported Maximums

Maximum Supported Number of:	Host Bus Adapters per Server (1)	Active Controller Host Ports (HBA Targets) per HBA(2)	LUNs per HBA Target(3)	HBAs per Server per Switch Zone (4)	Active Hosts per Storage System (5)
HP-UX	**4**	**4**	**8**	**4**	**16**
Heterogeneous SAN					**8**

Table Notes:

1. The actual maximum number of host bus adapters per server is dependent on the specific server model.
2. The recommended maximum number of active HSG80 controller host ports a single HBA can simultaneously access, the actual number of Targets available is 16.
3. The maximum number of LUNs per SCSI target. The maximum number of LUNs per target a host bus adapter can address in a SAN may be less.
4. The maximum number of HBAs per server that can be connected to the same switch zone.
5. This column gives the maximum number of active hosts for one storage system in a Homogeneous SAN using the specific operating system. This assumes 1 FC HBA per host for non-HA Transparent Failover configurations and 2 FC HBAs per host for HA Multiple-bus Failover configurations. The Heterogeneous Access entry is the maximum total number of active host connections when multiple operating systems are accessing the same storage system.

- All configurations are supported under HP-UX 10.20 and 11.0
- For HP-UX, all server connections into the SAN use a FC-AL based HBA.
- The maximum number of FC HBAs per server is 4, or total allowed in the server if less than 4.

Use Selective Storage Presentation (see below) to assign specific LUNs exclusively to each HBA

- The maximum number of active controller ports per FC HBA is 4
- Any combination of HP-UX Servers - standalone or clustered, and RA8000/ESA12000 storage systems is allowed per SAN configuration, provided these rules are followed:
    1. A maximum of 16 HP-UX servers (assumes 1 FC HBA per server for non-HA and 2 FC HBAs per server for HA) per RA8000 or ESA12000 dual controller storage system, with a maximum of 8 servers configured on each active port.
    2. All specific SAN configuration rule limits are followed. Refer to the individual SAN configurations
- All configurations require the *Connection Name Operating System* parameter set to "HP"
- When configuring HSG80 controllers for access from HP-UX servers only, the controller SCSI mode should be set to "SCSI-2", **SCSI-3 mode is not supported on HP-UX platforms**.
- Each active controller host port presents one SCSI Target ID with up to eight LUNs to each server (FC HBA) that has a connected path to the port.
- The maximum number of LUNs per HSG80 controller pair is 128
- Dual redundant controllers can be configured for Transparent Failover Mode or Multiple-Bus Failover Mode. Multiple-bus Failover mode requiresVolume Groups to be created using the MKVGLINKS utility. For details see Appendix B.
- Each storage system can be configured using SWCC or the Command Line Interface (CLI) through either the HSG80 controller maintenance serial port or a TCP/IP interface
- All configurations utilize ACS V8.5F (Array Controller Software) configured for LOOP topology.
- ACS V8.5 provides for up to 64 connection name entries. For dual controllers in Transparent Failover Mode, up to 2 connection entries are created for each server HBA, thus a total of 32 server HBAs are visible to a storage system at any given time. In Multiple-Bus Failover Mode, up to 4 connection entries are created for each server HBA, so a total of 16 server HBAs are visible. Note: The maximum number of server HBAs supported on a storage system is 16. Each storage system in the SAN will automatically add connection name entries upon initial power up. As a result, in certain SAN configurations it is possible to exceed the number of available connection name entries in a storage system prior to the addition of entries for the specific server HBAs that you want access to. This could result in the inability of a storage system to properly connect to a desired server HBA in the SAN. To avoid this, you can configure specific server HBAs and specific storage systems into FC Switch Zones of up to 32 Servers when in Transparent Failover and up to 16 Servers when in Multiple-bus Failover.
- Rename connection names. By default new connection names are automatically added to the controller connection name table by the controllers when they detect a path to an adapter from each active controller host port. The default connection name assigned by the controllers will be "!NEWCONnn", where nn is a number from 1 to 64. After the controllers detect all paths, rename each connection name to be more meaningful to the configuration i.e., SERVER1P1, SERVER1P2, etc. (connection names can be a maximum of 9 characters).
- Selective Storage Presentation - for configurations with more than 1 FC HBA in a server or more than one server accessing through the same HSG80 controller host port,

    Use the Connection Name *Unit Offset* value to set the range of unit numbers to be visible from each server or HBA (using the CLI, "SET *connection-name* UNIT_OFFSET = *n*")

    Use the Unit *Connection Name* parameter to allow exclusive access to units from each server or HBA (using the CLI, "SET unit-number DISABLE_ACCESS_PATH = ALL ENABLE_ACCESS_PATH = *connection-name*")

NOTE
Shared access between different servers to the same storage unit (LUN) requires specific application software (i.e., MC/ServiceGuard) to ensure proper data preservation.

- LUNs (logical units) can consist of single disks (JBOD), a storageset of multiple disks or a partition (up to 8), configured for a specific RAID level of 0, 1, 0 + 1, or 3/5
- The maximum storage capacity stated (2.6 TB) assumes 36GB disks, increases in disk drive capacity will increase overall storage capacity for the same enclosure footprint

MC/Service Guard

A HP-UX MSCS cluster enables two servers to share an RA8000 or ESA12000 FC storage system through a Fibre Channel Fabric Attached Arbitrated Loop (Quick Loop). If a server failure should occur, the storage I/O is re-routed through to the next functioning server

In addition to cluster failover, the RAID array controllers can be configured in transparent failover mode to protect against the unlikely event of a single controller failure. This feature eliminates controllers as single points of failure and is invoked without the need for server intervention.

The highest availability can be achieved by configuring the RA8000 or ESA12000 in Multibus_Failover mode, and configuring four hardware paths between each of the cluster servers and the FC storage system. The HP-UX Logical Volume Manager (LVM) PV Links utility provides the failover mechanism. This configuration provides at least two independent physical paths between the servers and the FC storage system.

Multibus/PV Links General Configuration Notes

- An HP-UX server connects to the switch as an Arbitrated Loop (FC-AL) Private Loop device. All storage systems it accesses must also be Arbitrated Loop devices. The server and storage systems must be in the QuickLoop (QuickLoop enabled) on the switches. For more information on QuickLoop, please see Appendix A.
- Dual redundant controllers can be configured for Failover Mode (referred to as Transparent Failover Mode) or Multibus_Failover Mode.
- All configuration diagrams show a front view of the storage system, Controller A is the top controller; Controller B is the bottom controller. Controller Host Port 1 is the left port; Host Port 2 is the right port
- Each storage system can be configured using SWCC through a TCP/IP interface on the host
- All configurations utilize ACS V8.5 (Array Controller Software) configured for FC-AL topology
- All configurations shown utilize short wavelength laser/50 micron multi-mode fibre channel optical cables in all cable segments

All configurations require that the *Arbitrated Loop Physical Address* (ALPA), for each loop device, be unique (refer to the *RA8000/ESA12000 FC Solution Software V8.5 for HP-UX - Installation Reference Guide*, Appendix D, for a list of valid addresses), and is set to *LOOP_HARD*.

- To support Multibus_Failover, PV Links **must** be configured for full redundancy in the manner described in the *Configuring PV Links using the mkvglinks Utility* paragraph contained in Appendix B.
- When using Multibus_Failover, hsg_vgmon should be run. During installation , an entry is made in the /etc/inittab to start hsg_vgmon when the system is booted. ‘mkvglinks’ will ensure ‘hsg_vgmon’ is running each time you configure a volume group for PV Links.

Configuration Descriptions

Table 3 lists six SAN Fabric/QuickLoop configuration examples for RA8000 and ESA12000 storage systems when using HP-UX platforms. The configurations are based on the rules listed in the configuration guideline section. They show the limits with regard to:

- Total number and arrangement of fibre channel switches in a single Fabric, or multiple Fabrics for high availability
- Maximum number of servers and storage systems per Fabric configuration
- Recommended Fabric/QuickLoop switch configurations optimized for general purpose, maximum capacity, maximum connectivity, highest availability, backup, or highest performance

The configurations are listed in order based on the number of switches in the Fabric. This follows a logical progression where you may desire to initially configure a small SAN containing a single FC switch, then add a second switch for high availability or increased connectivity, and additional switches as capacity and connectivity needs increase.

The maximum server counts and maximum storage counts listed for each configuration show the upper bound limits - either maximum servers or maximum storage, for that switch arrangement. Depending on your connectivity or storage capacity needs you may choose to implement your SAN maximized for one or the other, or a lesser number of either. All configurations allow the flexibility to trade-off server and storage system count based on total FC switch port availability. However, you cannot exceed the upper bound limits listed or the platform limits in Table 2 for each configuration. All configurations support a mix of both clustered and non-clustered servers. All configurations are supported with all Compaq supplied 16-port and 8-port FC switches.

Table 3 RA8000 / ESA12000 SAN Configurations for HP-UX - Index

SAN Configuration #	Description	Pages
1	**One FC Switch, Controller Transparent Failover Mode, QuickLoop**	**10 - 11**
2	**Two FC Switches, High Availability, PV Links, Controller Multiple-Bus Failover Mode, Quick Loop**	**12 - 13**
3	**Two FC Switches, Cascaded, Controller Transparent Failover Mode, Quick Loop**	**14 - 16**
4	**Four FC Switches, Cascaded High Availability Fault Tolerant Fabric, PV Links, Controller Multiple-Bus Failover Mode, Quick Loop**	**17 - 19**

Additional Information/References

A table of all available product documentation is included at the end of this application note for reference. Refer to the documents listed for more details on the product installation and configuration.

SAN Configuration 1 - One FC Switch, Quick Loop, Controller Transparent Failover Mode

SAN Configuration 1 (Figure 2) is a general-purpose configuration that provides connectivity for multiple servers, disk storage, and tape storage all on a single SAN. Up to 12 separate non-cooperating 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 1 storage system. With either switch type the number of storage systems can be increased if the number of servers is reduced to provide additional 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 2 One Switch Quick Loop Configuration Example

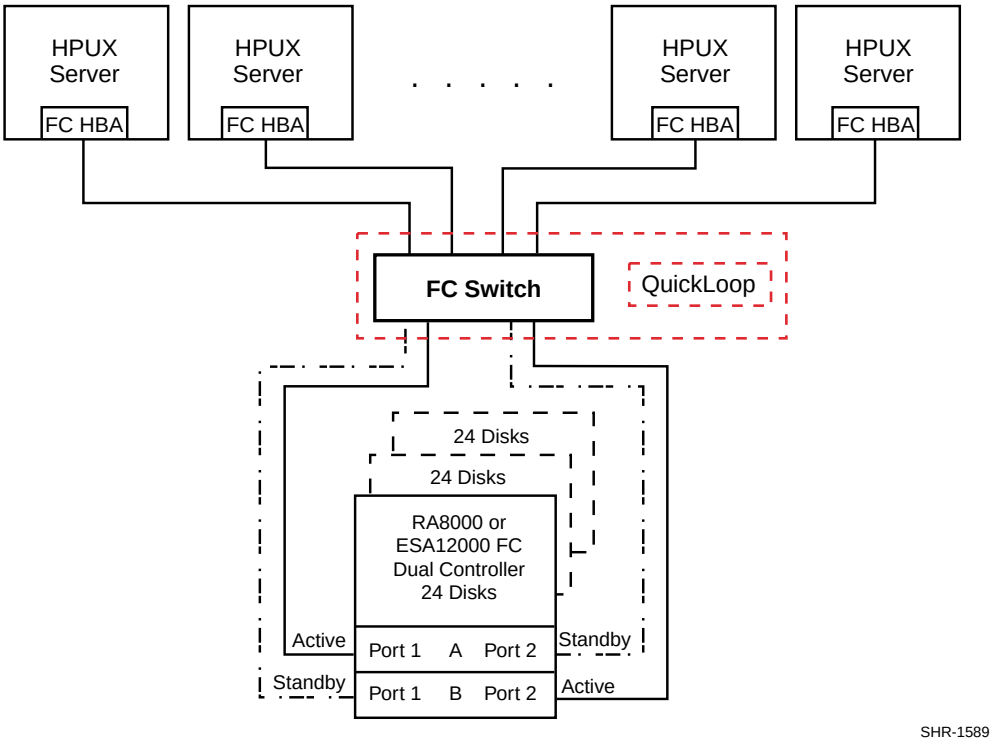

SAN Configuration 1 Server/Storage Connectivity Rules

- An HP-UX server connects to the switch as an Arbitrated Loop (FC-AL) Private Loop device. All storage systems it accesses must also be Arbitrated Loop devices. The server and storage systems must be in the QuickLoop on the switches. **For more information on QuickLoop, please see Appendix A.**
- Any mix of HP-UX Servers - standalone or clustered, and storage systems using these maximums. Assumes 1 FC HBA per server, an even storage port count.
- 16-port FC Switch, total number of switch ports available for Servers and Storage is 16
  - Maximum Server connectivity,
    - Up to 12 Servers and 1 Storage System (4 storage ports/2 Active, 2 Standby)
      [16 – 4 storage ports = 12 server ports]
      Maximum of 8 servers configured on any one active controller host port
  - Or Maximum Storage capacity,
    - Up to 4 Storage Systems (14 storage ports/7 Active, 7 Standby) and 2 Servers
      [16 – 2 server ports = 14 storage ports]
- 8-port FC Switch, total number of switch ports available for Servers and Storage is 8

Maximum Server connectivity,

Up to 4 Servers and 1 Storage System (4 storage ports/2 Active, 2 Standby)
[8 – 4 storage ports = 4 server ports]

Or Maximum Storage capacity,

Up to 2 Storage Systems (6 storage ports/3 Active, 3 Standby) and 2 Servers
[8 – 2 server ports = 6 storage ports]

- All HSG80 array controllers configured in Transparent Failover Mode

SAN Configuration 1 Fabric Rules

- 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. Refer to general Fabric rules for other cable types

SAN Configuration 1 Notes

- For servers accessing controller host port 1 use LUN number and offset values in the range of 0 – 99, for controller host port 2 use number and offset values in the range 100 -199
- Performance considerations, assuming relatively equal server load:

If using a single storage system, for balanced performance across all servers configure half of the servers on each of the controller port pairs. If using multiple storage systems configure an equal number of servers on each storage system.

The recommended maximum number of controller host ports configured per FC HBA is 2 for high bandwidth applications and 4 for high throughput applications.

SAN Configuration 2 – Two FC Switches, High Availability, PV Links, Controller Multiple-Bus Failover Mode

SAN configuration 2 (Figure 3) is a high availability storage configuration that uses two separate Fabrics and Quick Loops to provide two data paths between servers and disk storage configured in Multiple-Bus failover mode. Up to 15 separate (non-cooperating) servers or up to 7 pairs of clustered (cooperating) servers can be connected through two 16-port FC switches with 1 storage system. With 8-port FC switches up to 6 servers can be connected to 1 storage system. With either switch type the number of storage systems can be increased if the number of servers is reduced to provide switch ports for storage. This SAN configuration utilizes the high availability features of controller Multiple-Bus failover and LVM PV Links. With 2 FC HBAs in each server, and 2 FC switches, a second separate path to the storage is provided to each server. The 2 switches form isolated Quickloops providing the highest level of storage path fault tolerance should a failure occur with any path component - FC HBA, FC Switch, Fabric path interconnect, or storage controller. Non-HA configured servers and storage configured in Transparent Failover Mode are supported in this configuration, however access may be limited to one or the other of the two Quick Loops unless 2 FC HBAs are utilized, each connected to a different Quick Loop.

Figure 3 Two Switch HA Quick Loop Configuration Example

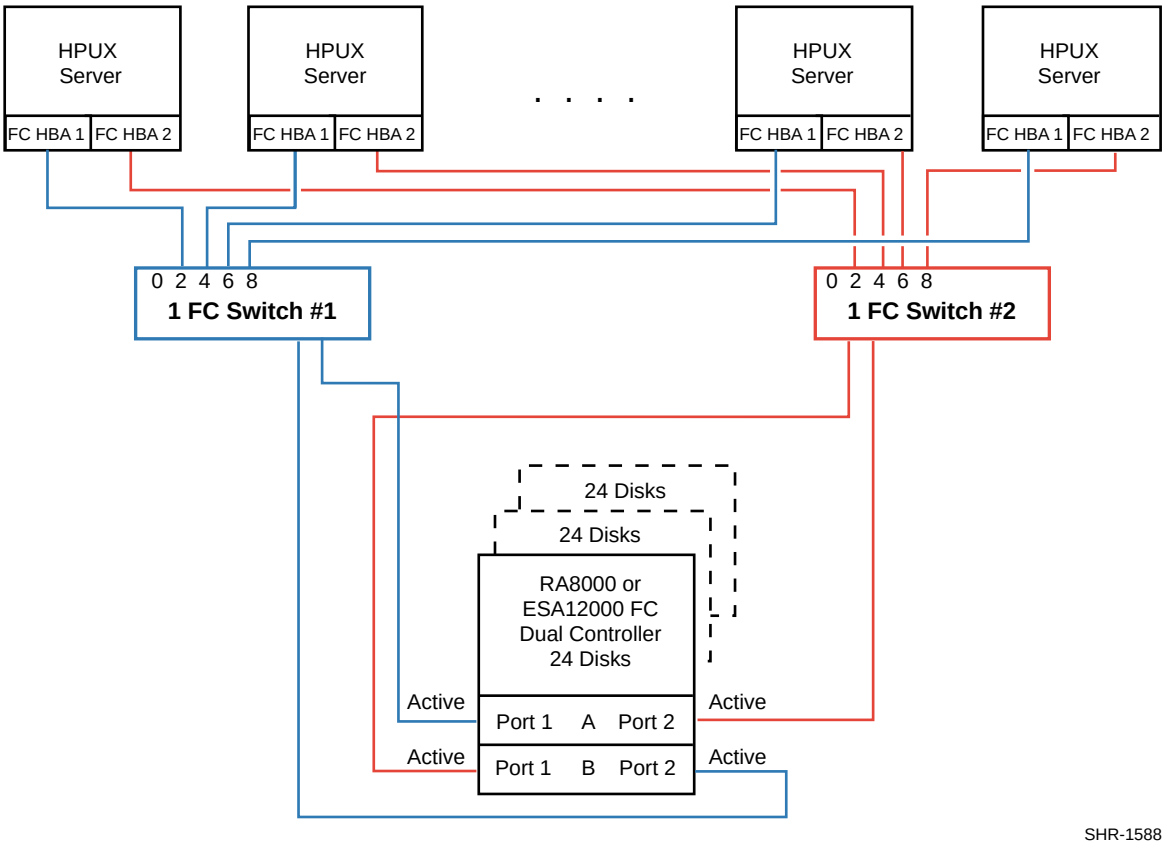

SAN Configuration 2 Server/Storage Connectivity Rules

- An HP-UX server connects to the switch as an Arbitrated Loop (FC-AL) Private Loop device. All storage systems it accesses must also be Arbitrated Loop devices. The server and storage systems must be in the QuickLoop (QuickLoop enabled) on the switches. **For more information on QuickLoop, please see Appendix A.**
- Any mix of HP-UX - standalone or clustered, configured for HA or non-HA, and any mix of Multiple-bus Failover and Transparent Failover storage systems using these maximums. Assumes all servers configured for HA, 2 FC switches and 2 FC HBA's per server, an even storage port

count.

- 16-port FC Switches, total number of switch ports available for Servers and Storage is 32
  Maximum Server connectivity,
  Up to 14 Servers and 1 Storage System (4 Active storage ports)
  [(32 – 4 storage ports)/2 = 14 servers]
  Maximum of 8 FC HBA's configured on any one active controller host port
  Or Maximum Storage capacity,
  Up to 6 Storage Systems (24 Active storage ports) and 4 HA Servers
  [32 – 8 server ports = 24 storage ports]
  Maximum of 4 active controller host ports per FC HBA
- 8-port FC Switches, total number of switch ports available for Servers and Storage is 16
  Maximum Server connectivity,
  Up to 6 Servers and 1 Storage System (4 Active storage ports)
  [(16 – 4 storage ports)/2 = 6 servers]
  Or Maximum Storage capacity,
  Up to 3 Storage Systems (12 Active storage ports) and 2 HA Servers
  [16 – 4 server ports = 12 storage ports]
- For dual path high availability, HSG80 controller pairs are configured in Multiple-Bus Failover Mode. For non-HA single path servers, controllers are configured in Transparent Failover Mode

SAN Configuration 2 Fabric Rules

- Total Switches = 2, 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. Refer to general Fabric rules for other cable types

SAN Configuration 2 Notes

- For servers accessing controller host port 1 use LUN number and offset values in the range of 0 – 99, for controller host port 2 use number and offset values in the range 100 -199
- Performance considerations, assuming relatively equal server load:

  If using a single storage system, for balanced performance across all servers configure half of the servers on each of the controller port pairs. If using multiple storage systems configure an equal number of servers on each storage system.

  The recommended maximum number of controller host ports configured per FC HBA is 2 for high bandwidth applications and 4 for high throughput applications.

SAN Configuration 3 – Two FC Switches, Cascaded, Controller Transparent Failover Mode, QuickLoop

SAN configuration 3 (Figure 4) is a general purpose single Fabric cascaded switch configuration that provides increased server, disk, and tape storage connectivity capabilities. Up to 22 separate (non-cooperating) servers or up to 11 pairs of clustered (cooperating) servers can be connected through a 16-port switch providing access from all servers to 2 storage systems. With an 8-port FC switch up to 10 servers can be connected to 1 storage system. With either switch type the number of storage systems can be increased if the number of servers is reduced to provide switch ports for storage. A minimum of 1 ISL is required between the 2 switches, more may be required based on the specific server and storage configuration and count, and the applications utilized. The ISL(s) can be up to 10 km allowing connectivity between two remote sites. 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 4 Two Switches, Fabric/ QuickLoop Configuration Example

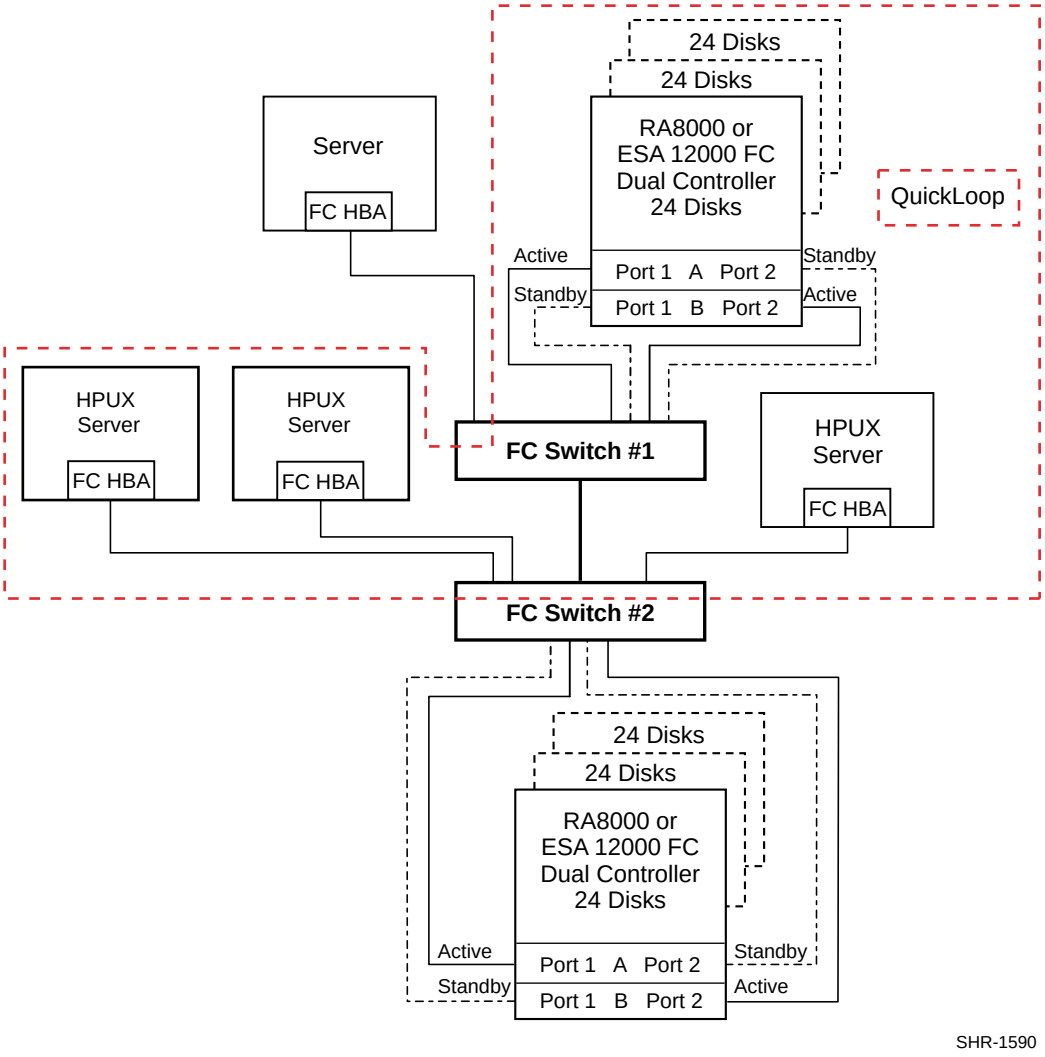

SAN Configuration 3 Server/Storage Connectivity Rules

- An HP-UX server connects to the switch as an Arbitrated Loop (FC-AL) Private Loop device. All storage systems it accesses must also be Arbitrated Loop devices. The server and storage systems must be in the QuickLoop on the switches. **For more information on QuickLoop, please see Appendix A.**

- Any mix of HP-UX Servers - standalone or clustered, and storage systems using these maximums. Assumes 1 FC HBA per server, an even storage port count, and no EBS ports.
- 16-port FC Switches, total number of switch ports available for Servers and Storage is 30
  Maximum Server connectivity,
  Up to 22 Servers and 2 Storage Systems (8 storage ports/4 Active, 4 Standby)
  [32 – 2 ISL ports – 8 storage ports = 22 server ports]
  Maximum of 16 servers configured on one controller pair
  Maximum of 8 servers configured on any one active controller host port
  Or Maximum Storage capacity,
  Up to 7 Storage Systems (26 Storage Ports/13 Active, 13 Standby) and 4 Servers
  [32 – 2 ISL ports – 4 server ports = 26 storage ports]
  Maximum of 4 active controller host ports per FC HBA
- 8-port FC Switches total number of switch ports available for Servers and Storage is 14
  Maximum Server connectivity,
  Up to 10 Servers and 1 Storage System (4 storage ports/2 Active, 2 Standby)
  [16 – 2 ISL ports – 4 storage ports = 10 server ports]
  Maximum of 8 servers configured on any one active controller host port
  Or Maximum Storage capacity,
  Up to 3 Storage Systems (12 storage ports/6 Active, 6 Standby) and 2 Servers
  [16 – 2 ISL ports – 2 server ports = 12 storage ports]
  Maximum of 4 active controller host ports per FC HBA
- All HSG80 array controllers configured in Transparent Failover Mode

SAN Configuration 3 Fabric Rules

- 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 = 15
- 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. Refer to general Fabric rules for other cable types
  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 3 Notes

- Each FC Switch must have a unique domain number (Domain_ID)
- For servers accessing controller host port 1 use LUN number and offset values in the range of 0 – 99, for controller host port 2 use number and offset values in the range 100 -199
- Performance considerations, assuming relatively equal server load:

  If using two storage systems, for balanced performance across all servers configure half of the servers on each of the storage systems and equally on the controller port pairs.
  The recommended maximum number of controller host ports configured per FC HBA is 2 for high bandwidth applications and 4 for high throughput applications.
  Use these general performance rules to determine the optimum number of ISLs required between both switches.
  For balanced SAN performance configure 1 storage system on each FC switch.
  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 High Availability Fault Tolerant Fabric, MC/ServiceGuard/ Controller Multiple-Bus Failover Mode, Quick Loop

SAN configuration 4 (Figure 5) is a high availability storage cascaded switch configuration that provides increased server, disk, and tape storage connectivity capabilities. This configuration uses two separate Fabrics/Quick Loops to provide two data paths between servers and disk storage configured in Multiple-Bus failover mode. Up to 26 separate (non-cooperating) servers or up to 13 pairs of clustered (cooperating) servers can be connected through 4 16-port FC switches to 4 storage systems. With 8-port FC switches up to 6 servers can be connected to 1 storage system. With either switch type the number of storage systems can be increased if the number of servers is reduced to provide switch ports for storage. This SAN configuration utilizes the high availability features of controller Multiple-Bus failover and LVM PV Links. With 2 FC HBAs in each server and 4 FC switches in 2 separate isolated Fabrics, a second separate path to the storage is provided to each server. The 2 switch pairs form isolated Fabrics providing the highest level of storage path fault tolerance should a failure occur with any path component - FC HBA, FC Switch, Fabric path interconnect, or storage controller. Non-HA configured servers and storage configured in Transparent Failover Mode are supported in this configuration, however access may be limited to one or the other of the two Fabrics unless 2 FC HBAs are utilized, each connected to a different Fabric.

Figure 5 Four Switches, Cascaded HA Fault Tolerant Fabric Configuration Example

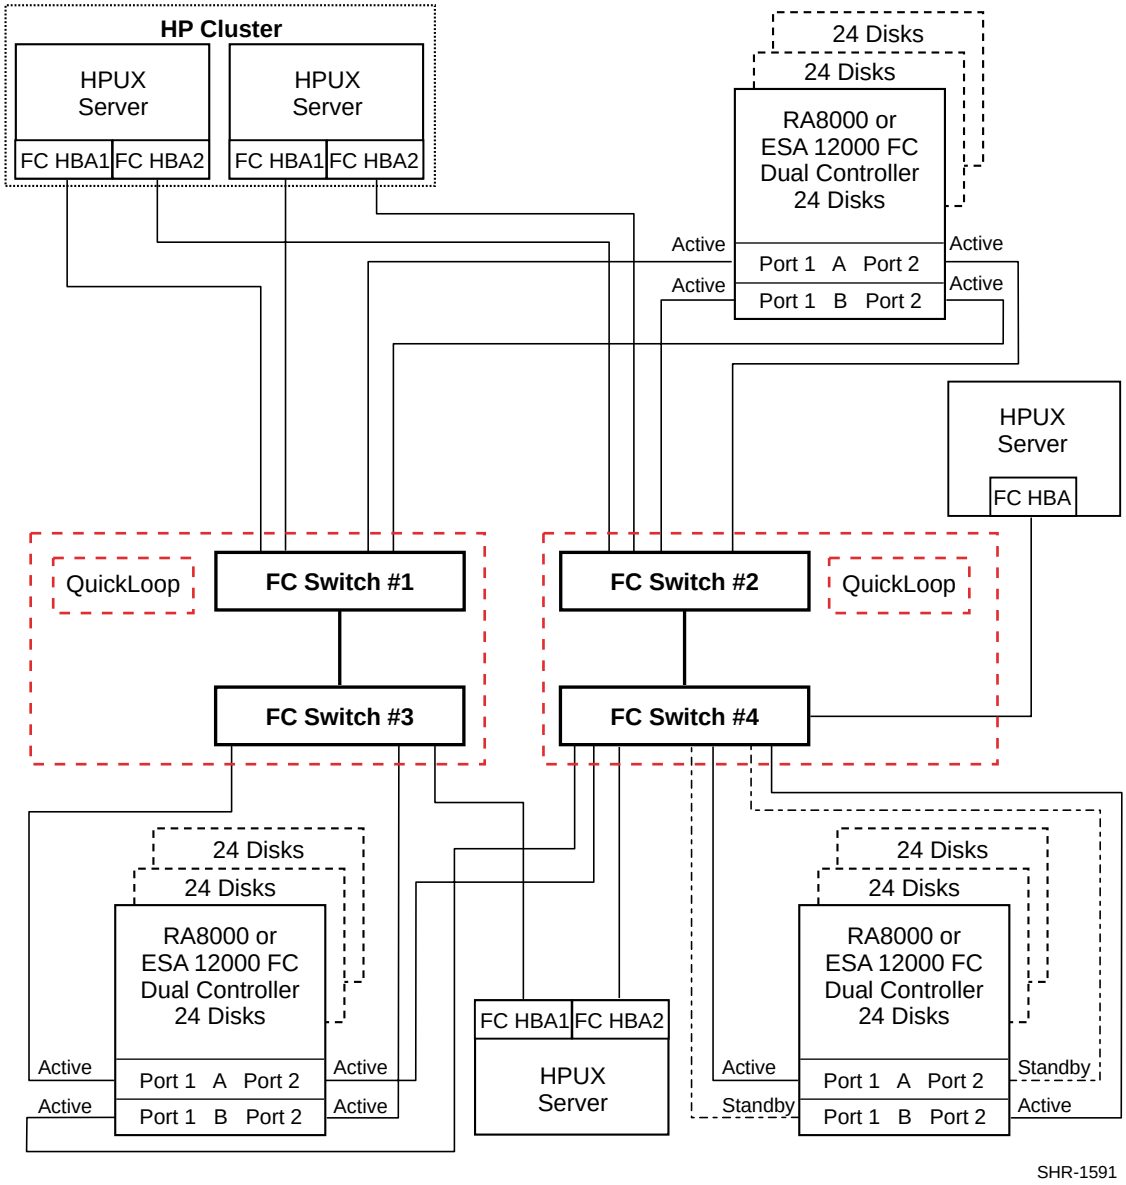

SAN Configuration 4 Server/Storage Connectivity Rules

- An HP-UX server connects to the switch as an Arbitrated Loop (FC-AL) Private Loop device. All storage systems it accesses must also be Arbitrated Loop devices. The server and storage systems must be in the QuickLoop (QuickLoop enabled) on the switches. For more information on QuickLoop, please see Appendix A.
- Any mix of HP-UX Servers - standalone or clustered, configured for HA or non-HA, and any mix of Multiple-bus Failover and Transparent Failover storage systems using these maximums. Assumes all servers configured for HA, 2 FC switches and 2 FC HBA's per server for HA, an even storage port count, and no EBS ports.
- 16-port FC Switches, total number of switch ports available for Servers and Storage is 60
  Maximum Server connectivity,
  Up to 26 Servers and 2 Storage Systems (8 Active storage ports)
  [(64 – 4 ISL ports – 8 storage ports)/2 = 26 servers]
  Maximum of 16 servers configured on any one controller pair

Maximum of 8 servers configured on any one active controller host port

Or Maximum Storage capacity,

Up to 12 Storage Systems (48 Storage Ports) and 6 HA Servers
[64 – 4 ISL ports – 12 server ports = 48 storage ports]
Maximum of 4 active controller host ports per FC HBA

- 8-port FC Switches, total number of switch ports available for Servers and Storage is 28

  Maximum Server connectivity,

  Up to 12 Servers and 1 Storage System (4 storage ports)
  [(32 – 4 ISL ports – 4 storage ports)/2 = 12 server ports]
  Maximum of 8 servers configured on any active controller host port

  Or Maximum Storage capacity,

  Up to 6 Storage Systems (22 storage ports) and 3 HA Servers
  [32 – 4 ISL ports – 6 server ports = 22 storage ports]
  Maximum of 4 active controller host ports per FC HBA

- For dual path high availability, HSG80 controller pairs are configured in Multiple-Bus Failover Mode, for single path servers (non-HA) controllers are configured in Transparent Failover Mode

SAN Configuration 4 Fabric Rules

- Total Switches = 4, 16-port or 8-port
- Maximum Switch Hops (nominal) = 1 in each path
- Maximum Switch Hops (worst case) = 1 in each path
- Maximum number of ISLs between the two switch in each path = 15
- 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. Refer to general Fabric rules for other cable types
  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 4 Notes

- Each FC Switch within a path must have a unique domain number (Domain_ID)
- For servers accessing controller host port 1 use LUN number and offset values in the range of 0 – 99, for controller host port 2 use number and offset values in the range 100 -199
- Performance considerations, assuming relatively equal server load:

  If using two storage systems, for balanced performance across all servers configure half of the servers on each of the storage systems and equally on the controller port pairs.
  The recommended maximum number of controller host ports configured per FC HBA is 2 for high bandwidth applications and 4 for high throughput applications.

Parts List

- **HPUX**

Compaq Part #	Description
166315-001	RA8000/ESA12000 FC Solution Software V8.5 for HP-UX
TBS	ACS V8.5F Controller SW
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)	FC Switch 16 port (no GBICs) FC Switch 16 port (no GBICs) FC Switch 8 port (no GBICs) FC Switch 8 port (no GBICs)
380561-B21	FC Optical GBIC
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	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

Product Details

The RA8000/ESA12000 storage systems utilize the Compaq HSG80 RAID controller running Array Controller Software (ACS) V8.5. 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 implementation, disk mirroring, and disk partitioning capabilities. In addition, ACS manages host 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 unlimited 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 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 512 MB of cache in each controller. ESA12000 FC controllers come with 256 MB of cache. Cache options allow for 512 MB 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 remaining 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 5 power supplies providing N+1 redundancy. Power can be upgraded to full redundancy with the addition of 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 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 host platform.

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.

Configuration References

The following documents provide further information on the configuration of RA8000 and ESA12000 Fibre Channel Storage Systems:

Table 4 Configuration Reference Material

Topic	Document Title	Order Number
Heterogeneous SAN Application Note	*Heterogeneous Storage Area Networks*	*EK-SMA30-AN*
Sun SAN Application Note	*Storage Area Network Configurations for RA8000/ESA12000 on Sun Solaris*	*EK-SMA37-AN*
Windows NT intel SAN Application Note	*Storage Area Network Configurations for RA8000/ESA12000 on Windows NT*	*AA-RHH6B-TE*
OpenVMS SAN Application Note	*Storage Area Network Configurations for RA8000/ESA12000 on OpenVMS*	*EK-SMA34-AN*
Tru64 UNIX SAN Application Note	*RA8000/ESA12000 FC-Switch Configurations for Tru64 UNIX*	*EK-SMA33-AN*
Windows NT – Intel Standard FC-AL Configurations Application Note	*RA8000/ESA12000 FC-AL Configurations for Windows NT – Intel*	** AA-RH0RA-TE*
Windows NT – Intel High Availability FC-AL Configurations Application Note	*RA8000/ESA12000 FC-AL High Availability Configurations for Windows NT – Intel*	**AA–RH0SA–TE*
Windows NT – Intel HSG80 Array Controller	*Release Notes, HSG80 Array Controller (ACS V8.5F) for Windows NT Server – Intel*	*AA–RH0QB–TE or 387393-002*
Windows NT – Intel RA8000/ESA12000 Installation	*RA8000/ESA12000 Fibre Channel Solution Software for Windows NT – Intel Installation Reference Guide*	*AA–RFA9C–TE*
Secure Path Installation	*StorageWorks Secure Path for Windows NT Installation Guide*	*EK-WNTMP-MH*
HSG80 Array Controller Configuration	*HSG80 Array Controller ACS Version 8.5 Configuration Guide*	*EK-HSG85-CG or 165144-001*
HSG80 Array Controller Command Line Interface	*HSG80 Array Controller ACS Version 8.5 CLI Reference Guide*	*EK-HSG85-RG or 165145-001*
HSG80 Array Controller Maintenance	*HSG80 Array Controller ACS Version 8.5 Maintenance and Service Guide*	*EK-HSG84-SV or 118620-002*
StorageWorks Command Console	*Command Console V2.2 (HSG80) for RAID Array 8000/ESA12000, User's Guide*	*AA–RFA2D–TE*

* Refer to the website for the latest information.

Appendix A

QuickLoop on Fibre Channel Switches

- HP-UX servers can only be configured for Arbitrated Loop (FC-AL) and must be used inside a QuickLoop. They only access devices in the same QuickLoop. All HP-UX host connections must be set for **HP**.
- HSG80 Controllers using ACS V8.5F only can participate in a QuickLoop.
- To add an HP-UX Arbitrated Loop (FC-AL) Private Loop devices to your SAN, use the switch QuickLoop feature.
- HSG80 controllers used in a QuickLoop, configured for Arbitrated Loop, must use a host port topology of LOOP_HARD. QuickLoop is supported only on specific fibre channel switch models.

- Compaq fibre channel switches provide an Arbitrated Loop connectivity feature called QuickLoop. It provides the ability to connect Arbitrated Loop devices to the switch. QuickLoop is used in Private Loop mode only.
- Each QuickLoop consists of fabric loop devices connected to specified switch ports on one switch or two switches partnered together. The partnered switches may have intermediate switches between them that do not directly participate in the QuickLoop.
- In a configuration where a SAN has both fabric and loop devices connected to it, the switches provide a translative mode which assigns a public loop address to fabric devices communicating with loop devices in the QuickLoop. The fabric devices can access the loop devices in the QuickLoop, but not vice versa.
- The maximum number of switches per QuickLoop is two switches. The maximum number of QuickLoops per switch is one QuickLoop.
- In a SAN where multiple QuickLoops exist, devices in different QuickLoops cannot communicate with each other.
- QuickLoop is supported only on the following models of fibre channel switches:

Compaq 16 port Fibre Channel SAN Switch, part # 158223-B21

Compaq 8 port Fibre Channel SAN Switch, part # 158222-B21

NOTE

Currently, fibre channel hubs are not supported connected to a QuickLoop. For more information regarding QuickLoop, please consult the *SAN Switch QuickLoop Management Guide* and other supporting documentation supplied with your fibre channel SAN switch.

QuickLoop Commands

QuickLoop is enabled on a fibre channel switch through commands entered in a telnet session logged into a switch account with administrative privileges. If two fibre channel switches are partnered in one QuickLoop each switch is configured separately.

QuickLoops are managed with the following commands:

qlEnable, qlDisable
: These are used for configuring an entire switch for QuickLoop.

qlPortEnable, qlPortDisable
: These are used for configuring individual switch ports for QuickLoop.

qlPartner
: This is used for partnering two switches into one QuickLoop.

qlShow, qlHelp
: These are used for the management of QuickLoops.

5.2 QuickLoop Example

This example shows how to set up a QuickLoop on two switches.

Figure 6 QuickLoop Example

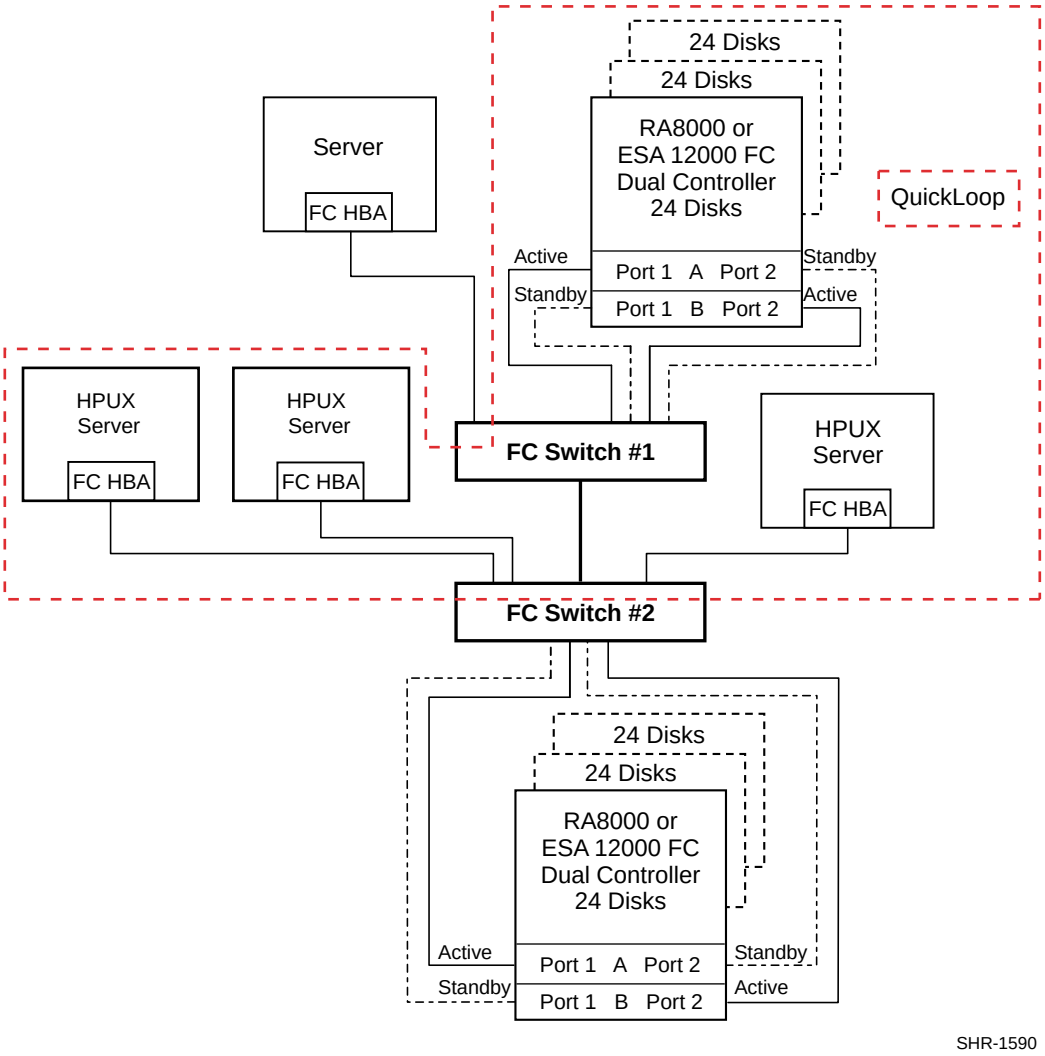

In Figure 6:

- FC Switch #1 and FC Switch #2 are partnered in a QuickLoop
- The cluster, the standalone server on the right, and the top storage system are in the QuickLoop
- The top standalone server and the bottom storage system are outside of the QuickLoop
- The top standalone server is connected to Port 0 of FC Switch #1
- The top storage system is connected to FC Switch #1
  - Host Port 1 of the top controller is connected to Switch Port 4

  - Host Port 1 of the bottom controller is connected to Switch Port 5
  - Host Port 2 of the top controller is connected to Switch Port 6
  - Host Port 2 of the bottom controller is connected to Switch Port 7
- The cluster is connected to Port 0 and Port 1 of FC Switch #2
- The standalone server on the right is connected to Port 2 of FC Switch #2
- The bottom storage system is connected to FC Switch #2
  - Host Port 1 of the top controller is connected to Switch Port 4
  - Host Port 1 of the bottom controller is connected to Switch Port 5
  - Host Port 2 of the top controller is connected to Switch Port 6
  - Host Port 2 of the bottom controller is connected to Switch Port 7
- Port 3 of FC Switch # 1 is connected to Port 3 of FC Switch #2
- FC Switch #1 has a World Wide Name of **1000 0060 a706 9c4e**
- FC Switch #2 has a World Wide Name of **1000 0060 444d 8302**

The following telnet session log shows the switch commands used to set up the QuickLoop on FC Switch #1 in this example. The user input is in **bold** and explanations are in boxes:

switch1:admin> **qlPortEnable 4**

Setting port 4 to QuickLoop mode,

Committing configuration...done.

Activate looplet 4

switch1:admin> **qlPortEnable 5**

Setting port 5 to QuickLoop mode,

Committing configuration...done.

Activate looplet 5

switch1:admin> **qlPortEnable 6**

Setting port 6 to QuickLoop mode,

Committing configuration...done.

Activate looplet 6

switch1:admin> **qlPortEnable 7**

Setting port 7 to QuickLoop mode,

Committing configuration...done.

Activate looplet 7

These enable QuickLoop for Ports 4, 5, 6 and 7 on the switch, for the top storage system.

switch1:admin> **qlPartner "10:00:00:60:44:4d:83:02"**

Setting QuickLoop to dual-switch mode,

Committing configuration...done.

This indicates the switch with the WWN 1000 0060 444d 8302 (FC Switch #2) is to be the QuickLoop partner for this switch (FC Switch #1).

The following telnet session log shows the switch commands used to set up the QuickLoop on FC Switch #2:

switch2:admin> **qlPortEnable 0**

Setting port 0 to QuickLoop mode,

Committing configuration...done.

Activate looplet 0

switch2:admin> **qlPortEnable 1**

Setting port 1 to QuickLoop mode,

Committing configuration...done.

Activate looplet 1

> These enable QuickLoop for Port 0 and 1 on the switch, for the cluster.

switch2:admin> **qlPortEnable 2**

Setting port 2 to QuickLoop mode,

Committing configuration...done.

Activate looplet 2

> This enables QuickLoop for Port 2 on the switch, for the standalone server on the right.

switch2:admin> **qlPartner "10:00:00:60:a7:06:9c:4e"**

Setting QuickLoop to dual-switch mode,

Committing configuration...done.

> This indicates the switch that has the WWN 1000 0060 a706 9c4e (FC Switch #1) is to be the QuickLoop partner for this switch (FC Switch #2).

The following telnet session log shows the QuickLoop configuration on FC Switch #1:

switch1:admin> **qlShow**

Self: 10:00:00:60:a7:06:9c:4e domain 1

Peer: 10:00:00:60:44:4d:83:02 domain 2

State: Master

Scope: dual

AL_PA bitmap: 03000000 00000011 00000000 00000000

Remote AL_PAs

[021000]: 04

[021100]: 08

[021200]: 10

Local AL_PAs

[011400]: 71

[011700]: 72

Local looplet states

Member: 4 5 6 7

Online: 4 - - 7

Looplet 4: online

Looplet 5: offline

Looplet 6: offline

Looplet 7: online

> This shows the QuickLoop configuration for the switch FC Switch #1.

The following telnet session log shows the QuickLoop configuration on FC Switch #2:

```
switch2:admin> qlShow
Self:   10:00:00:60:44:4d:83:02 domain 2
Peer:   10:00:00:60:a7:06:9c:4e domain 1
State:  Non-Master
Scope:  dual
AL_PA bitmap:   03000000 00000011 00000000 00000000
Remote AL_PAs
  [011400]:   71
  [011700]:   72
Local AL_PAs
  [021000]:   04
  [021100]:   08
  [021200]:   10
Local looplet states
Member:    0 1 2
  Online:    0 1 2
  Looplet 0:  online
  Looplet 1:  online
  Looplet 2:  online
```

This shows the QuickLoop configuration for the switch FC Switch #2.

The QuickLoop allows the top standalone server to access both the top and bottom storage systems. The cluster and the standalone server on the right can only access the top storage controller.

NOTE

For more information regarding QuickLoop, please consult the *SAN Switch QuickLoop Management Guide* and other supporting documentation supplied with your fibre channel SAN switch.

Appendix B

The "mkvglinks" Utility

The "mkvglinks" utility will create volume groups with physical volume on the RA8000 / ESA12000 Storage Subsystem and link all the paths to that volume. It will also link all the paths to the physical volume on existing volume groups that use the RA8000 / ESA12000 subsystem. Mkvglinks will also launch hsg_vgmon to monitor the volume group. Logical volumes must then be created according to user requirements.

The "hsg_vgmon" Daemon

The "hsg_vgmon" will check that there is a volume group that is attached to a RA8000 / ESA12000 RAID Array and that the RAID Array is in Multibus Mode. If there is none then it will exit. If there is a RAID Array in Multibus Mode, it will check that the volume group has access to its physical volume. If the current path to the physical volume is not "ready", hsg_vgmon will send a start command to that physical volume. To ensure hsg_vgmon starts when rebooting the system, an entry should be in the /etc/inittab file as follows:

hsvg::once:/opt/steam/bin/hsg_vgmon >> /var/adm/syslog/steamd.log 2>&1

Configuring PV Links using the "mkvglinks" Utility

Place the controller in Multibus_Mode, as described in the *HSG80 Array Controller ACS Version 8.3 Configurations and CLI Reference Guide (EK-HSG80-RG*). Make the preferred controller the one that is connected to the primary path, using the following command to determine the primary path of an existing volume group:

vgdisplay –v /dev/<volume-group>

Set the preferred path using the following HSG80 Command Line Interface command:
HSG80> **SET unit-number PREFERRED_PATH = THIS_CONTROLLER (or OTHER_CONTROLLER)**

Adding Links to the Physical Volume

The mkvglinks utility will create a new volume group or will take an existing volume group that uses physical volume units on the RA8000/ESA12000 storage system. It will add all the links to the physical volume. Both methods are described below.

NOTE
It is important that the volume group associated with the physical volume has no mounted file systems and is deactivated on all other attached hosts.

Using the "mkvglinks" Utility on Existing Volume Groups

On an existing volume group, the mkvglinks utility will link all the paths to the physical volumes. It will also allow you to provide primary paths to other physical volumes that need to be added. The command is:

/opt/steam/bin/mkvglinks <volume-group> [/dev/dsk/<physical-volume>]

If the physical volume already has volume group information on it, the utility will ask whether to overwrite the volume. Enter **"ye**s" to continue adding the volume.

Using the mkvglinks Utility to Create a New Volume Group

Both the volume group name and the primary path to the physical volume must be included. The command is:

/opt/steam/bin/mkvglinks <volume-group> /dev/dsk/<physical-volume>

If the physical volume already has volume group information on it, the utility will ask whether to overwrite the volume. Enter **"ye**s" to continue adding the volume.

Logical Volumes

Next, create your logical volumes on the volume group. It is recommended that you create a volume group map file in /etc/lvmconf with the volume group name and map extension as follows:

/usr/sbin/vgchange –a n /dev/<volume-group>
/usr/sbin/vgexport –p [-s] –m /etc/lvmconf/<volume-group>.map /dev/<volume-group>
/usr/sbin/vgchange –a y /dev/<volume-group>

The *mkvglinks* utility must be run on each existing volume groups that use RA8000/ESA12000 LUNs as physical volumes. If this is a MC/Service Guard cluster then *mkvglinks* must be run on the other hosts that share volume groups connected to the RA8000/ESA12000 storage system. If the volume group does not exist on the other hosts then copy the map file, created in the above procedure, over to the other hosts that will be sharing those volumes. Import the volume group as follows:

import –s –m /etc/lvmconf/<volume-group>.map /dev/<volume-group>

If for any reason a deactivated volume group cannot be activated, either reboot the machine or use the following commands to recover:

1. # **ll /dev/<volume-group>/group** (and note the minor number, the field before the date field)
2. # **mkdir /dev/<volume-group>**
3. # **mknod /dev/<volume-group>/group c 64 <minor_number>**
4. # **/usr/sbin/vgexport /dev/<volume-group>**
5. # **/usr/sbin/vgimport [-s] –m /etc/lvmconf/<volume-group>.map /dev/<volume-group>**
6. # **vgchange –a y /dev/<volume-group>**

For shared volume groups between hosts, copy the map file to the other host and perform steps 2 through 6, above. Note that when using LVM maintenance commands that reference devices directly, (i.e., pvmove), you must first verify that the path is active using the following command:

 # /opt/steam/bin/ctrlinfo . <device_special_file>
Where: device_special_file does not include the path.

If the message **"0 unk unk unk 0 un**k" is returned, then the path is not active, and you must use a different path. To produce a list of alternate paths, type:

 # vgdisplay –v /dev/<volume-group>