

# ***Storage Area Network Configurations for RA8000/ESA12000, MA8000/EMA1200 and MA6000 on HP-UX***

# ***Compaq StorageWorks***

**Application Note**  
EK-SMA32-AN. B01

---

## **Introduction**

This application note describes Storage Area Network (SAN) configurations for RA8000/ESA12000, MA8000/EMA12000, and MA6000 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.

## **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 about 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](http://www.compaq.com/storageworks)

## **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, MA8000/EMA12000, and MA6000 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/MA8000/MA6000 and ESA12000/EMA12000 FC storage systems are based on a common architecture and offer customers centralized management, high availability (HA)<sup>1</sup>, and exceptional performance and scalability in open systems environments. These products provide a Fibre Channel storage solution with industry leading SAN support using both FC-Fabric (switch) and FC-AL (arbitrated loop) technology.

The RA8000/MA8000 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/MA8000 supports up to 2.6 TB of storage. The ESA12000/EMA12000 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/EMA12000 FC offers unprecedented levels of storage, scaling in performance and capacity as user requirements increase.

The MA6000/MA8000 is a modular array RAID Storage system. The MA6000/MA8000 provides an easy-to-deploy, flexible, scalable solution for open systems. The MA6000 is a lower-priced member of the MA8000/EMA12000 product family. The MA6000 supports up to 864GB of storage capacity. The MA6000 is designed to provide a single architecture from the desktop to the data center. The MA6000 is also an excellent choice for an initial SAN installation and/or remote office.

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 effect on performance while preserving data availability and supporting high availability.<sup>1</sup>

The RA8000/MA8000/MA6000 and ESA12000/EMA12000 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.<sup>2</sup> PV Links detects the failure of I/O operations to complete on a failed path and automatically reroutes all traffic to the surviving path.

---

<sup>1</sup> The configurations in this App Note are supported under HP-UX 10.20 and 11.00 configurations. Only the HP-UX 10.20 systems support high availability (HA); for 11.00 systems, contact your Compaq storage consultant.

<sup>2</sup> Multiple-Bus failover mode is supported under HP-UX 10.20 only.

---

## SAN Configuration Guidelines

This application note describes homogeneous HP-UX SAN configurations. It is expected that customers may want 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

- If you are configuring a homogeneous HP-UX SAN, refer to the HP-UX Server/Storage Rules configuration rules in this document to determine how best to configure servers and storage.
- 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.

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 four 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 reroute to a 2-hop path on a single link failure. A hop is defined as one or more connections between two FC switches, for example; two 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 two 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 reroute, 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 reroute, 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 reroute, 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**



SHR-1585

### HP-UX Server/Storage Rules

Table 1 describes the HP-UX platform maximums when in a homogenous HP-UX SAN.

**Table 1 Guidelines for HP-UX Platform Supported Maximums**

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

*Table Notes:*

- The maximum number of host bus adapters per server depends on the specific server model.
  - The recommended maximum number of active HSG controller host ports that a single HBA can simultaneously access; the actual number of Targets available is 16.
  - The maximum number of LUNs per SCSI target. The maximum number of LUNs per target that an HBA can address in a SAN may be less.
  - The maximum number of HBAs per server that can be connected to the same switch zone.
  - The maximum number of active hosts for one storage system in a Homogeneous SAN using the specific operating system. This assumes one FC HBA per host for non-HA Transparent Failover configurations and two FC HBAs per host for HA Multiple-bus Failover configurations.
- The configurations in this App Note are supported under HP-UX 10.20 and 11.00. HP-UX configurations. HP-UX 10.20 systems only support high availability (HA); for 11.00 systems, contact your Compaq storage consultant.
  - 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 fewer 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, MA8000/EMA12000, and MA6000 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/MA8000/MA6000 or ESA12000/EMA12000 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 HSG60/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 HSG controller pair is 128.
- Dual redundant controllers can be configured for Transparent Failover Mode or Multiple-Bus Failover Mode. Multiple-bus Failover mode requires Volume 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 controller maintenance serial port or a TCP/IP interface
- All configurations use ACS V8.5F, V8.5b, and 8.5L (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 one FC HBA in a server or more than one server accessing through the same HSG60/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 MCSG cluster enables two servers to share an RA8000/MA8000/MA6000 or ESA12000/EMA12000 FC storage system through a Fibre Channel Fabric Attached Arbitrated Loop (Quick Loop). If a server failure should occur, the storage I/O is rerouted 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/MA8000/ MA6000 or ESA12000/EMA12000 in Multiple bus\_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.

### Multiple Bus/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 Multiple Bus-Failover Mode.<sup>2</sup>
- 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 use ACS V8.5 (Array Controller Software) configured for FC-AL topology
- All configurations shown use 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 *installation and configuration documentation for your storage system*). To support Multiple Bus\_Failover, PV Links **must** be configured for full redundancy in the manner described in the *Configuring PV Links using the mkvglinks Utility* in Appendix B.

- When using Multiple Bus\_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 2 lists six SAN Fabric/QuickLoop configuration examples for RA8000/MA8000/MA6000 and ESA12000/EMA12000 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 2 RA8000 / ESA12000, MA8000/EMA12000, and MA6000 SAN Configurations for HP-UX - Index**

SAN Configuration #	Description	Pages
1	One FC Switch, Controller Transparent Failover Mode, QuickLoop	8 - 9
2	Two FC Switches, High Availability, PV Links, Controller Multiple-Bus Failover Mode, Quick Loop	10-11
3	Two FC Switches, Cascaded, Controller Transparent Failover Mode, Quick Loop	12-12

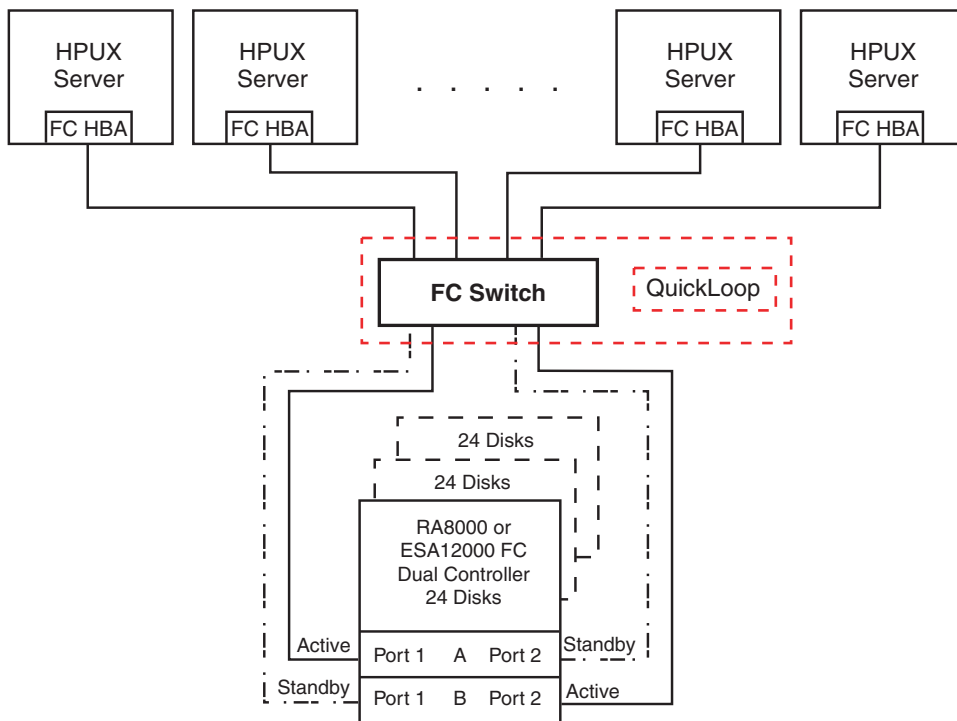
**NOTE:**

HP-UX 10.20 systems only support high availability (HA); for 11:00 systems, contact your Compaq storage consultant

### 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 a single 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 HSG60/HSG80 array controllers within the RA8000/MA8000/MA6000 or ESA12000/EMA12000 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



SHR-1589

**Note: For the MA6000, this figure applies to a dual-controller (12 disks) only.**



### **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 about QuickLoop, 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 HSG60/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:

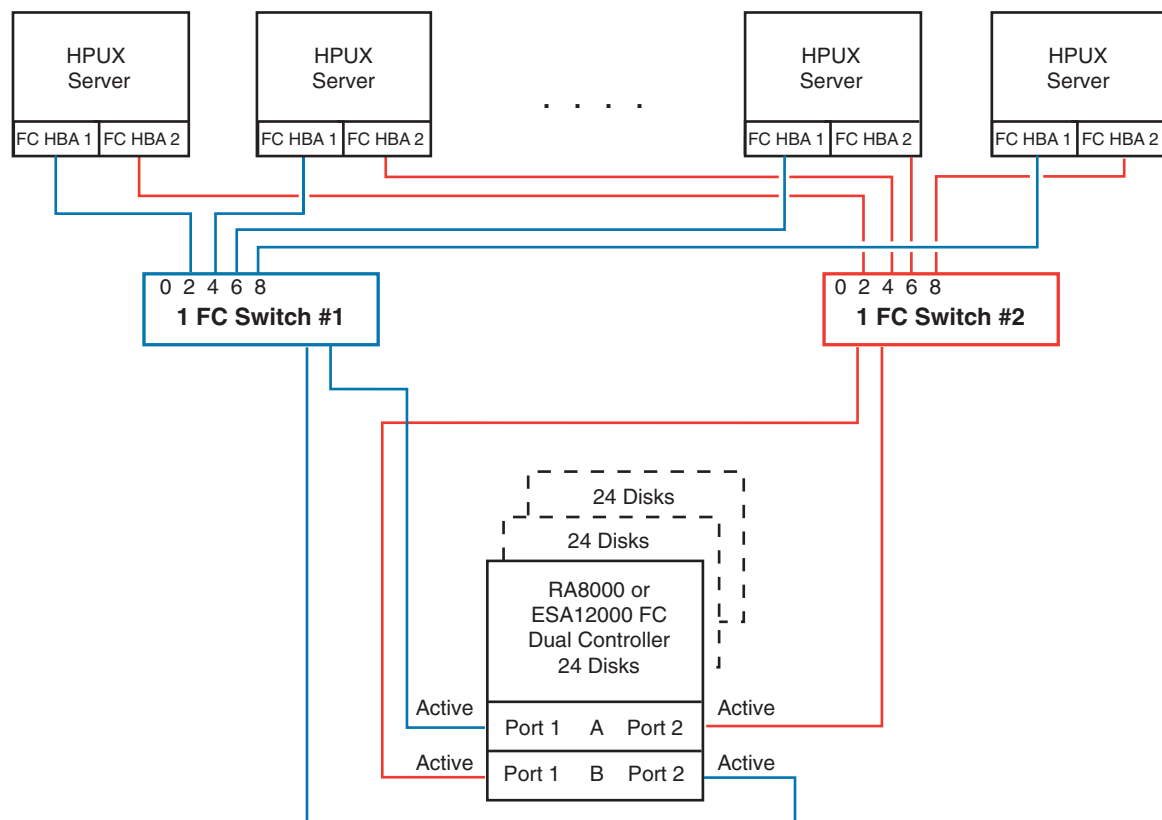
If using a single storage system, for balanced performance across all servers configure half of the servers on each of the controller port pairs. If using multiple storage systems configure an equal number of servers on each storage system.

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 Mode1

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 one single storage system. With 8-port FC switches up to six servers can be connected to one single 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 uses the high availability features of controller Multiple-Bus failover and LVM PV Links. With two FC HBAs in each server, and two FC switches, a second separate path to the storage is provided to each server. The two switches form isolated Quick Loops 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 two FC HBAs are used, each connected to a different Quick Loop.

Figure 3 Two Switch HA Quick Loop Configuration Example



SHR-1588

## 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 that it accesses must also be Arbitrated Loop devices. The server and storage systems must be in the QuickLoop (QuickLoop enabled) on the switches.

NOTE: For more information about QuickLoop, 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, two FC switches and two FC HBAs 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)
    - $[(32 - 4 \text{ storage ports})/2 = 14 \text{ servers}]$
    - Maximum of 8 FC HBAs 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
    - $[32 - 8 \text{ server ports} = 24 \text{ storage ports}]$
    - 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)
    - $[(16 - 4 \text{ storage ports})/2 = 6 \text{ servers}]$
  - Or Maximum Storage capacity,
    - Up to 3 Storage Systems (12 Active storage ports) and 2 HA Servers
    - $[16 - 4 \text{ server ports} = 12 \text{ storage ports}]$
- For dual path high availability, HSG 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:

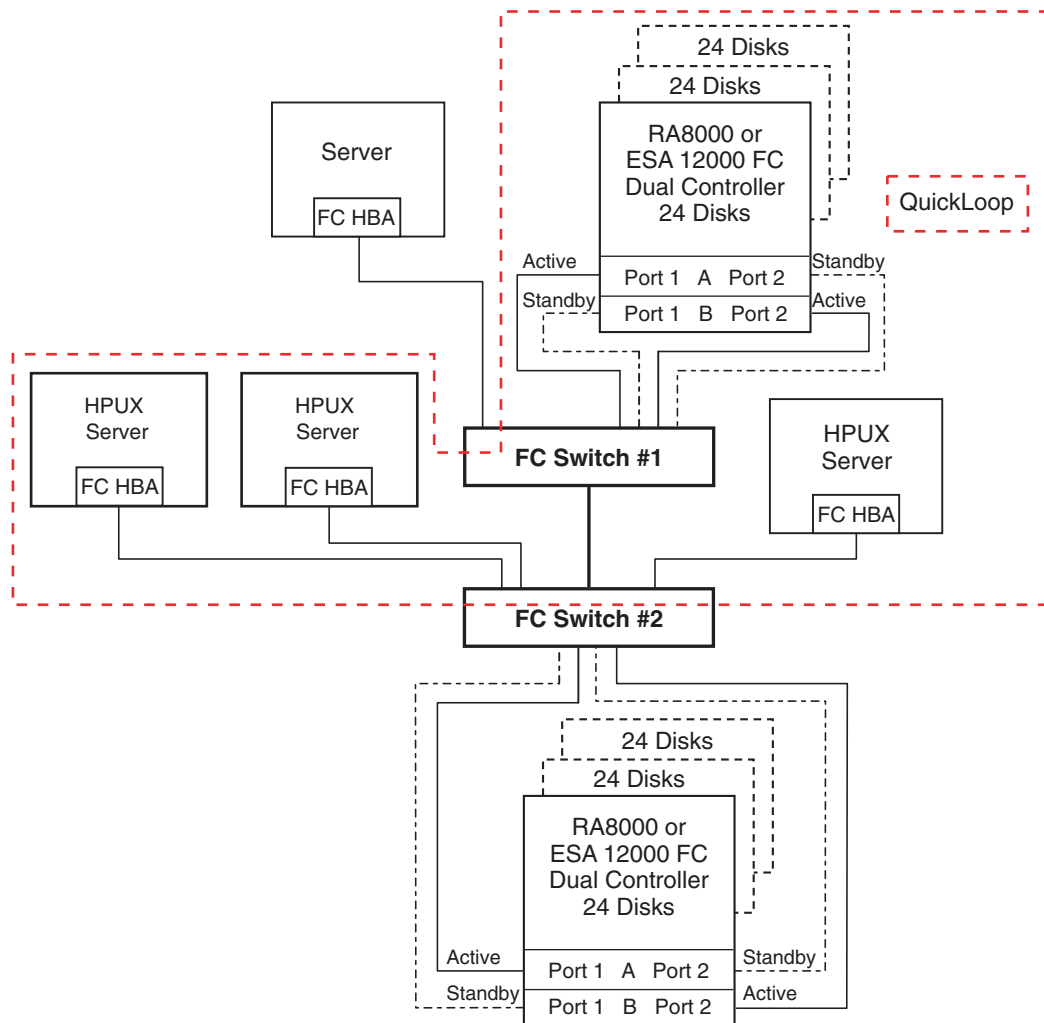
If using a single storage system, for balanced performance across all servers configure half of the servers on each of the controller port pairs. If using multiple storage systems configure an equal number of servers on each storage system.

The recommended maximum number of controller host ports configured per FC HBA is two for high bandwidth applications and four 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 two storage systems. With an 8-port FC switch up to 10 servers can be connected to one single 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 one ISL is required between the two switches, more may be required based on the specific server and storage configuration and count, and the applications used. The ISL(s) can be up to 10 km allowing connectivity between two remote sites. The HSG array controllers within the SotrageWorks 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



SHR-1590

### 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.

NOTE: For more information about QuickLoop, 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 two 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 two 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 six storage controller ports on one switch being accessed by a server on the other switch.

## Appendix A

### QuickLoop on Fibre Channel Switches

- HP-UX servers can be configured only 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 and HSG60 Controllers using ACS V8.5L 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.
- HSG60/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 that are partnered. 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 that 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; the maximum number of QuickLoops per switch is one.
- 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.

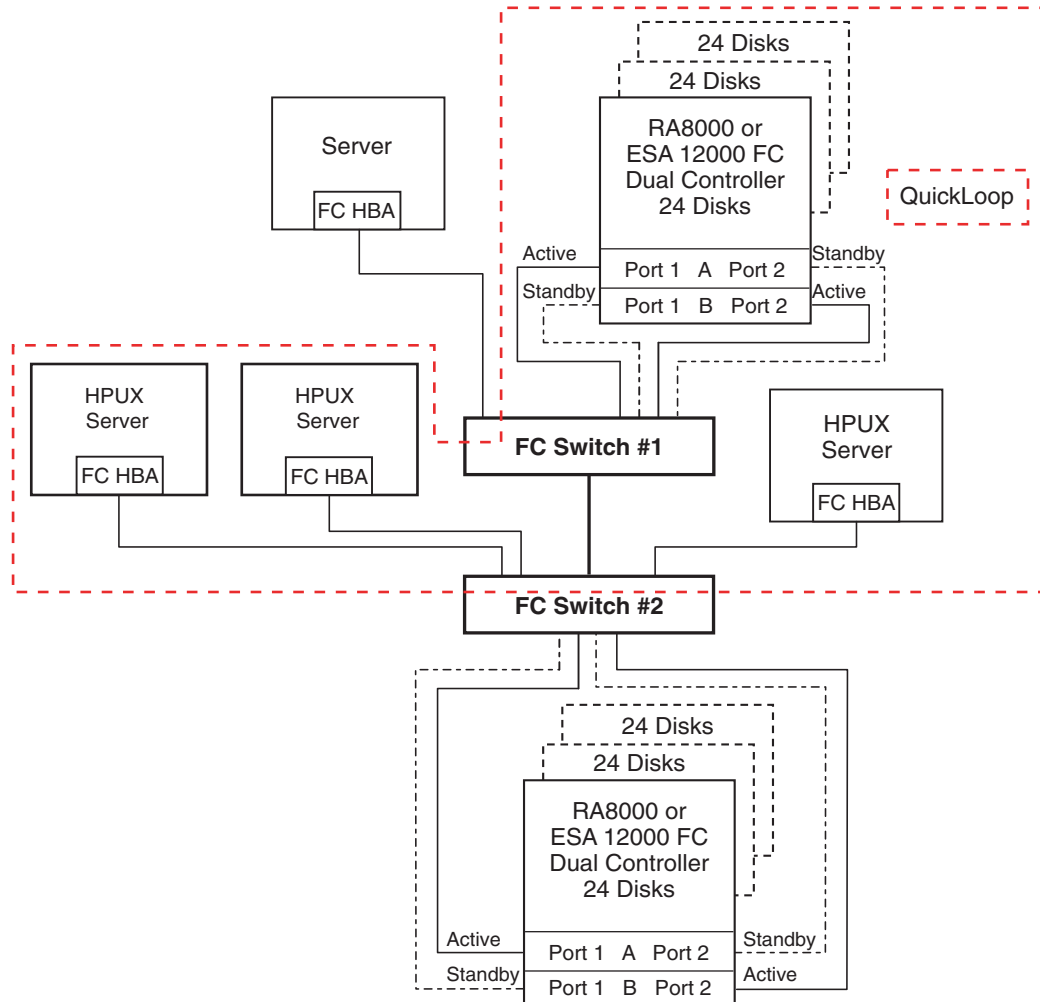
`qlShow`, `qlHelp`

These are used for the management of QuickLoops.

## QuickLoop Example

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

Figure 5 QuickLoop Example



SHR-1590

**Note: For the MA6000, this figure applies to a dual-controller (12 disks) only.**



### About Figure 5

---

- 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**
-

***Telnet Session Log Showing Switch Commands for QuickLoop Setup on FC Switch #1***

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).

---

**Telnet Session Log Showing Switch Commands for QuickLoop Setup on FC Switch #2**

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).

***Telnet Session Log Showing QuickLoop Configuration on FC Switch #1***

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.
---

---

**Telnet Session Log Showing QuickLoop Configuration on FC Switch #2**

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 creates volume groups with physical volume on the RA8000 /ESA12000, MA8000/EMA12000, and MA6000 Storage Subsystems 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, MA8000/EMA12000, and MA6000 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, MA8000/EMA12000, and MA6000 RAID Array and that the RAID Array is in Multiple-bus mode. If there is none then it will exit. If there is a RAID Array in Multiple-bus bus 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 Multiple Bus\_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 HSG60/HSG80 Command Line Interface command:

```
HSG60/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, MA8000/EMA12000, and MA6000 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 “yes” 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 “yes” to continue adding the volume.

*RA8000 / ESA12000 FC-AL High Availability Configurations for Hewlett Packard HP-UX*  
Page 10 EK-FCALM-AA. A011

## 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 existing volume groups that use RA8000/ESA12000 MA8000/EMA12000, and MA6000 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 MA8000/EMA12000, and MA6000 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 share 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 2-6 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 unk**” is returned, then the path is not active, and you must use a different path. To produce a list of alternate paths, type:

```
# vdisplay -v /dev/<volume-group>
```