

# *Storage Area Network Configurations for RA8000/ESA12000 on LINUX*

# *Compaq StorageWorks*

## **Application Note**

EK-SMA42-AN. A01

---

### **Introduction**

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

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

### **Visit Our Web Sites for the Latest Information**

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

<http://www.compaq.com/storage>

<http://www.compaq.com/support>

#### **NOTE**

This application note describes RA8000/ESA12000 switched fabric configurations for Linux.

### **Enterprise Network Storage Architecture (ENSA)**

Compaq StorageWorks RA8000 FC and ESA12000 FC products are the latest in fully integrated Fibre Channel RAID storage solutions.

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

## **Product Overview**

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

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

The ESA12000 FC, designed for the data center, is available in easy-to-order building blocks. Each block is 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 performance while preserving data and high availability support.

Platform kits are components of software and documentation created for each specific operating system that can be used with the Compaq StorageWorks RAID arrays. Platform kits exist for the SUN (Solaris) operating system, Hewlett-Packard (HP-UX), Linux Alpha, and Linux x86 operating system, and so on.

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

## Configuration References

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

**Table 1 Configuration Reference Material**

<b>Topic</b>	<b>Document Title</b>	<b>Order Number</b>
<i>Linux HSG80 Array Controller</i>	<i>Release Notes, HSG80 Array Controller (ACS V8.5) for Linux (Intel)</i>	<i>AA-RKJPB-TE</i>
	<i>Release Notes, HSG80 Array Controller (ACS V8.5) for Linux (Alpha)</i>	<i>AA-RKJRB-TE</i>
<i>Linux RA8000/ESA12000 Installation</i>	<i>RA8000/ESA12000 and MA8000/EMA12000 Solution Software V8.5b for Linux Installation Reference Guide</i>	<i>AA-RKJQB-TE</i>
<i>HSG80 Array Controller Configuration</i>	<i>HSG80 Array Controller ACS Version 8.5 Configuration Guide</i>	<i>EK-HSG85-CG</i>
	<i>HSG80 Array Controller ACS Version 8.5 CLI Reference Guide</i>	<i>EK-HSG85-RG</i>
<i>HSG80 Array Controller Maintenance</i>	<i>HSG80 Array Controller ACS Version 8.5 Maintenance and Service Guide</i>	<i>EK-HSG84-SV</i>
<i>StorageWorks Command Console</i>	<i>Compaq StorageWorks Command Console V2.3 User Guide</i>	<i>AA-RFA2F-TE</i>
<i>Warranty Information</i>	<i>DIGITAL StorageWorks Warranty Terms and Conditions</i>	<i>EK-HSXSW-WC</i>
<i>HSG80 Array Controller Parts</i>	<i>HSG80 Array Controller Illustrated Parts Map</i>	<i>EK-HSG80-MP or 355217-001</i>

Intel is a registered trademark of Intel Corporation.

Windows NT and MSCS are trademarks of Microsoft Corporation.

StorageWorks and the Compaq Logo are trademarks of Compaq Computer Corporation.

All other trademarks and registered trademarks are the property of their respective owners.

## **General Configuration Guidelines**

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

## **Configuration Maximums**

### **The HSG80 Controller**

1. Maximum 128 visible LUNs
2. Maximum 200 assignable unit numbers
3. Maximum 72 physical devices
4. Maximum 20 RAID-5 StorageSets
5. Maximum 30 (RAID-5 and RAID-1) StorageSets
6. Maximum 45 (RAID-5 and RAID-1 and RAID-0) StorageSets
7. Maximum 8 partitions of a StorageSet or individual disk
8. Maximum 6 members of a RAID-1 (MirrorSet)
9. Maximum 14 members per RAID-5 StorageSet
10. Maximum 24 members per RAID-0 (StripeSet)
11. Maximum 48 physical devices per Striped-MirrorSet
12. Maximum 64 Connections
13. Maximum 9 Characters per Connection Name
14. Maximum 512 GB per StorageSet

## The Fibre Channel SAN

1. Maximum 500 meters per optical cable segment, 50 micron, multi-mode, short wave GBICs.  
175 meters per optical cable segment, 62.5 micron, multi-mode, short wave GBICs.  
10 km per Inter-Switch Link (ISL) cable segment, 9 micron, single mode, long wave GBICs.
2. Maximum 11 kilometer nominal distance between Server and Storage.(21 km. worst case single fault reroute).
3. Maximum 4 Servers per RA8000 or ESA12000
4. Maximum 2 HBAs per Server
5. Maximum 4 Fibre Channel switches in a single fabric, i.e., SAN with all switches interconnected.
6. Maximum Number of ISLs between switches depends on the number of available ports and the required redundancy. A reasonable minimum is two links for redundancy.

## Assumptions

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

### Document References

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

## General Configuration Notes

1. All single controller storage configurations can be upgraded to dual controller configurations.
2. Controller pairs can be configured for either Dual Redundant Failover Mode *or* Multiple-bus Failover Mode. If the controllers are in Dual Redundant Failover Mode, it is referred to as "Transparent Failover Mode".

### **This application note describes Transparent Failover Mode configurations only.**

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

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

7. Configurations use both short and long wavelength GBICs and cables.

Inter-Switch Links (ISLs) are supported with long wavelength GBICs and 9-micron single mode cables.
8. All configurations **require** the Connection Name *Operating System* parameter set to "SUN".  
(CLI: "SET *connection-name* OPERATING\_SYSTEM = SUN")
9. All Configurations require the controller to be set to SCSI-3 mode.  
(CLI: "SET THIS\_CONTROLLER SCSI = SCSI-3")

As a result, the HSG80 Command Console LUN (CCL) is enabled on both unit number D0 and D100.
10. Controller units: The HSG80 controllers support a range of 200 unit numbers. In Transparent Failover Mode, the units D0-D99 are assigned to Port 1, and units D100-D199 are assigned to Port 2.

#### **NOTE**

However, there are 128 visible LUNs with D0 and D100 reserved for the CCL in SCSI-3 mode.

11. Record the FC HBA Worldwide Identification (WWID) address. Some HBAs display the WWID on the HBA.
11. Use the recorded WWID to identify the connection name displayed from the CLI at the controller using the CLI: “SHOW CONNECTIONS” command.
12. Rename connection names. By default, new connection names are automatically added to the controller connection name table as “!NEWCONnn” (where nn is a number from 1 to 64). The controller adds a connection when a path to an adapter from each active controller host port is detected. After you identify each connection using the FC HBA WWID, rename the connection names to a name meaningful to the specific configuration. (CLI: “RENAME *old-name new-name*”. Connection names may contain nine (9) characters). In some cases, host name and Port value may suffice (i.e., WESTP1, WESTP2, EASTP1, EASTP2, etc.).
13. To define access control from Server to Storage: it is possible to define a specific combination of Storage Unit and Host by enabling an ACCESS\_PATH. The steps are:
  - Disable/Clear all current access\_paths  
(CLI: “SET *unit-number* DISABLE\_ACCESS\_PATH = ALL”)
  - Set a specific Host access (through the connection name )  
(CLI: “SET *unit-number* ENABLE\_ACCESS\_PATH = *connection-name*”)
14. See the next section, “Connections.”

## **Connections**

For our purposes, "Connections" are defined as the end points of the physical path between a server and a storage system within a SAN. Obviously, a FC-Switch may serve to connect, direct or protect a path. Yet, our focus is to treat the connection as the end points of the path. There are two points of view when considering fibre channel SAN connections. Managing connections from these two points of view is the key to the definition, support, and maintenance of the servers and systems in a SAN.

### **The First View**

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

### **The Second View**

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

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

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

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

### ***Managing Connections - Server Strategies***

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

If the HBA is already installed or if the WWNN is not affixed to the adapter, the WWNN of the HBA is displayed on the system console when booting. The WWNN is a hexadecimal value of the form, nnnn-nnnn-nnnn-nnnn.

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

2. The server needs to know to which storage system(s) it will connect. The identification of the storage system is facilitated with the WWNN of each controller in the SAN.

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



---

### *Managing Connections - Controller Strategies*

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

If the controller reports that the controller is not in FABRIC MODE ( CLI: SHOW\_THIS ), do the following:

```
CLI: SET THIS_CONTROLLER PORT_1_TOPOLOGY = OFF_LINE
```

```
CLI: SET THIS_CONTROLLER PORT_2_TOPOLOGY = OFF_LINE
```

```
CLI: SET THIS_CONTROLLER PORT_1_TOPOLOGY = FABRIC
```

```
CLI: SET THIS_CONTROLLER PORT_2_TOPOLOGY = FABRIC
```

```
Perform the same commands for OTHER_CONTROLLER
```

and then

```
CLI: RESTART OTHER_CONTROLLER
```

```
CLI: RESTART THIS_CONTROLLER
```

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

Connections that are not required may be deleted.

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

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

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

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

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

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

Offsets are the designation of the mapping of the LUN 0 from a specific server to a unit on the storage system. Our problem started with two servers accessing the same units from the same

port. Thus, from server A, its LUN 0 maps to the lowest unit, for example D3. From server B its LUN 3 maps to D3. To solve this conflict problem as we set the base (LUN 0) to be all units above D40 for server A, by doing the following:

CLI: "SET *connection-name* UNIT\_OFFSET = *n*"

From our example,

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

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

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

The steps are:

- Disable/Clear all current access\_paths  
(CLI: "SET *unit-number* DISABLE\_ACCESS\_PATH = ALL")
- Set a specific server access (through the connection name )  
(CLI: "SET *unit-number* ENABLE\_ACCESS\_PATH = *connection-name*")

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

## Adjusting LUNs per SCSI Target ID

If you wish to connect many Storage Systems to a particular Linux host, it may be necessary to adjust the number of LUNs presented per SCSI Target ID on the host. During installation of the Solution Software, /usr/src/linux/drivers/scsi/scsi.c is patched to extend the SCSI Multi-LUN capability from 8 to 64 LUNs per target ID. Depending upon your SAN and Storage configuration, it may be desirable to adjust the number of LUNs per target ID to 8, 16 or 32. This is a three-step process:

1. Adjust the "Luns per Target" value of the QLA2200 HBA using the Fast!UTIL utility (invoked by pressing ALT-Q during QLA2200 initialization at system boot time). This option can be found in the Advanced Adapter Settings area.

### NOTE

The AlphaServer SRM Console does not support invoking the Fast!UTIL utility. Alpha users must make these changes by plugging the HBA into a x86 system.

2. Edit /usr/src/linux/drivers/scsi/scsi.c and change "static int max\_scsi\_luns" to the desired value. For example: static int max\_scsi\_luns = 16; (Remember to include the terminating ";")
3. Rebuild your kernel and re-boot.

---

## Configuration Descriptions

Table 2 lists four configuration examples for RA8000/ESA12000 storage systems in an FC-Fabric SAN using Linux platforms.

**Table 2 Examples of RA8000/ESA12000 FC-Switch Linux Configurations**

Configuration #	Description	Pages
1	One Server, Single Controller, One FC Switch	12
2	One Server, Dual Controller Transparent Failover, One FC Switch	13
3	Two Servers, Dual Controller Transparent Failover, One FC Switch	14
4	Four Servers, Two Dual Controller Pairs Transparent Failover, Two FC Switches	15

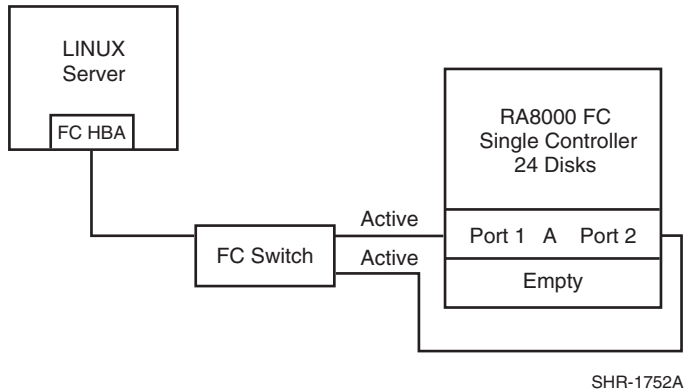
The examples follow a logical progression of an initial, entry level SAN that uses a single FC-Switch to a more complex configuration that uses multiple switches.

8-port and 16-port FC switches are supported. Server and storage system port counts can be adjusted and optimized based on FC-Switch port availability, and server and storage connectivity needs.

### Configuration 1 - One Server, Single Controller, One FC Switch

The single controller configuration (**Figure 1**) is an entry-level RA8000 FC or ESA12000 that can be expanded to accommodate additional storage capacity or performance needs as required. The configuration consists of an RA8000 or ESA12000 storage enclosure with a single HSG80 array controller on a single fibre channel switch. In this configuration, both controller host ports and the server HBA are connected to the switch. Additional servers can be added to this configuration, as required, using the existing switch.

**Figure 1 Configuration 1**



### Configuration 1 Notes

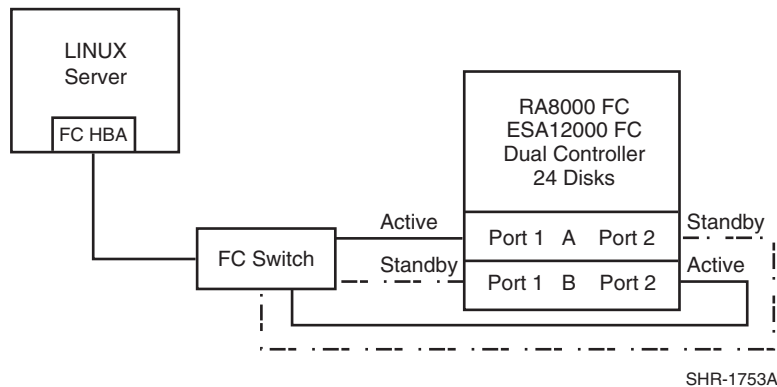
- Single Fibre Channel Switch
- Single HSG80 controller (controller failover is not available)
- Up to 500 meters per 50-micron cable segment
- Up to 72 disks with 2 optional expansion enclosures
- Controller in SCSI-3 mode
- Port topology set to FABRIC
- Connection O/S set to SUN

## Configuration 2 - One Server, Dual Controller Transparent Failover, One FC Switch

This two-controller configuration (**Figure 2**) provides dual redundant controllers configured in Transparent Failover mode. The configuration consists of an RA8000 FC or ESA12000 storage enclosure with two HSG80 array controllers on a single switch.

This configuration provides full controller redundancy if either controller fails. During failover, the surviving controller transitions the standby port to the active state providing access to all units for the server.

**Figure 2 Transparent Failover**



### Configuration 2 Notes


- Single Fibre Channel Switch
- Dual HSG80 controllers configured in Transparent Failover Mode
- Up to 500 meters per 50-micron cable segment
- Up to 72 disks with 2 optional expansion enclosures
- By default, setting the controller pair in Transparent Failover Mode will automatically configure Port 1 and Port 2 of both controllers to the Active and Standby state indicated, and set the unit offset values to 0 for Port 1 and 100 for Port 2
- Controller in SCSI-3 mode
- Port topology set to FABRIC
- Connection O/S set to SUN
- 2 logical connections are available with up to 126 LUNs. Suggested controller unit numbers: Port 1, units D1 – D99; Port 2, units D101 – D199 (default)

### Configuration 3 – Two Servers, Dual Controller Transparent Failover, One FC Switch

Configuration 3 (**Figure 3**) uses a dual controller pair to provide storage and controller redundancy to a second, separate (non-cooperating) server. The two servers share the storage enclosure in that each server has exclusive access to its own LUNs through the use of mutually exclusive access path management. This configuration uses one fibre channel switch. If a controller failover occurs, the surviving controller continues to provide full access to units of both servers.

In normal operation Controller A provides the active path for all storage sets assigned to the servers through Port 1. Controller B provides the active path for all storage sets assigned to the Servers through Port 2. Each controller has a dedicated standby port if a controller failover occurs. Additional servers can be added to this configuration as required based on the available ports at the FC Switch.

**Figure 3 Configuration 3**



### Configuration 3 Notes

- One Fibre Channel Switch
- Shared storage enclosure
- Dual HSG80 controllers configured in Transparent Failover Mode
- Up to 500 meters per 50-micron cable segment
- Up to 72 disks with 2 optional expansion enclosures
- Up to 126 LUNs total, with 2 active controller host ports.
- By default, setting the controller pair in Transparent Failover Mode will automatically configure Port 1 and Port 2 of both controllers to the Active and Standby state indicated, and set the unit offset values to 0 for Port 1 and 100 for Port 2
- Controllers in SCSI-3 mode
- Port topology set to FABRIC
- Connection O/S set to SUN
- Rename connections to meaningful names (CLI:RENAME)
- Each server has exclusive access to its own LUNs through path management.

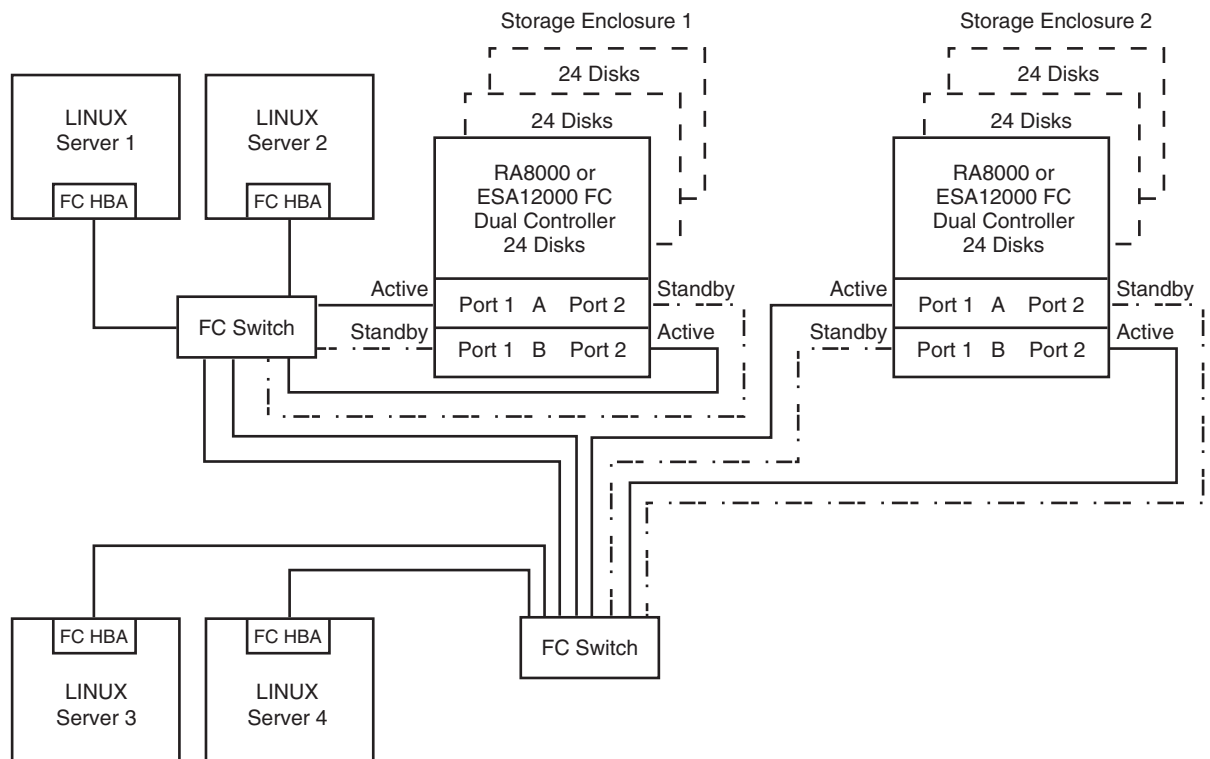
CLI: "SET *unit-number* DISABLE\_ACCESS\_PATH = ALL"

CLI: "SET *unit-number* ENABLE\_ACCESS\_PATH = *connection-name*"

## Configuration 4 – Up to Four Servers, Two Dual Controller Pairs, Transparent Failover, Two FC Switches

Configuration 4 (**Figure 4**) provides storage and controller redundancy for up to four, separate (non-cooperating) servers. The four servers share the storage enclosures in that each server has exclusive access to its own set of LUNs within each enclosure. This is made possible by the setting of connection-name access to give each server exclusive access to its LUNs in the SAN.

**Figure 4 Configuration 4**



SHR-1755A

### Configuration 4 Notes

- Two Fibre Channel Switches
- Shared storage enclosures
- Two Dual HSG80 controller pairs configured in Transparent Failover Mode
- Up to 500 meters per 50-micron cable segment
- Up to 144 disks with 4 optional expansion enclosures
- In this configuration, 16 connections exist. Each server has 4 connections.
- By default, setting the controller pairs in Transparent Failover Mode will automatically configure Port 1 and Port 2 of both controllers to the Active and Standby state indicated, and set unit offset values to 0 for Port 1 and 100 for Port 2.
- Port topology set to FABRIC
- Controllers in SCSI-3 mode
- Connection O/S set to SUN

- Rename connections to meaningful names (CLI:RENAME)
- Set exclusive access for each LUN using access path management.  
CLI: SET unit-number DISABLE\_ACCESS\_PATH = ALL  
CLI: SET unit-number ENABLE\_ACCESS\_PATH = connection-name