

***RA8000/ESA12000
FC-SW High Availability
Configurations
for Sun Solaris***

COMPAQ
STORAGeworks

Application Note

AA-RLMGA-TE

April, 2000

Introduction

This application note serves as a guide to some High-Availability Fibre Channel Fabric (HA FC-SW) Configurations for Compaq StorageWorks RAID Array 8000 (RA8000 FC) and Enterprise Storage Array 12000 (ESA12000 FC) storage system(s) in conjunction with the Sun family of servers running the Solaris Operating System.

Clearly, there are many factors of complexity in dealing with servers. They include performance, usage, load, the configuration of Fibre Channel fabric and the configuration of the storage sets on the RAID subsystem. Each installation or site may have needs that are beyond the scope of this application note, and to that end, we recommend contacting your local Compaq Services representative for on-site consultation.

Visit Our Web Site for the Latest Information

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

www.compaq.com/products/storageworks

NOTE

This application note describes RA8000/ESA12000 High Availability Fibre Channel Fabric configurations.

For High Availability Fibre Channel Arbitrated Loop configurations, refer to "RA8000/ESA12000 FC-AL High Availability Configurations for Sun Solaris".

For standard configurations refer to "RA8000/ESA12000 Configurations for Sun Solaris" (EK-FCALS-AA) and "Storage Area Network Configurations for RA8000/ESA12000 on Sun Solaris(EK-SMA37-AN)"

Enterprise Network Storage Architecture (ENSA)

The Compaq StorageWorks RAID Array 8000 and Enterprise Storage Array 12000 products are the latest in fully integrated Fibre Channel RAID storage solutions. These products provide the key initial steps in delivering on 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.

Table of Contents

Section Title	Page
1 StorageWorks RAID Product Overview	3
2 StorageWorks RAID Product Details	3
3 High Availability Application Options	4
3.1 SANworks Secure Path Version 2.1 for Sun Solaris	4
VERITAS Cluster Server V1.1.1 for Solaris V2.6/V2.7	5
VERITAS FirstWatch V2.5.5 for Solaris V2.6 (only)	5
4 High Availability Configuration Information	6
4.1 Reference Documentation	6
4.2 General Configuration Guidelines	7
Configuration Maximums	7
4.2.1.1 SAN Concepts	8
4.2.1.2 SAN Maximums	9
4.2.2 Assumptions	10
4.2.3 Configuration Notes	11
5 High Availability Configuration Topologies and Modes	14
5.1 One Server, Secure Path, Multiple-bus Failover	15
5.2 Two Servers, VERITAS Cluster Server, Transparent Failover	16
5.3 Four Servers, VERITAS Cluster Server, Transparent Failover, One Shared Storage Enclosure	18
5.5 Four Servers, VERITAS Cluster Server, Secure Path, Multiple-Bus Failover, Two Independent Server Pairs, One Shared Storage Enclosure	19
5.6 Up to Four Servers, Secure Path, Multiple-Bus Failover, One Shared Storage Enclosure	20
5.7 Up to Four Servers, Secure Path, Multiple-Bus Failover, Two Shared Storage Enclosures	22
5.8 Up to Eight Servers, Secure Path, Multiple-Bus Failover, Two Shared Storage Enclosures	24
	26
6 High Availability Parts List	28
Appendix A Controller Failover Transitions	30

1 StorageWorks RAID Product Overview

The RA8000 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 SAN (Storage Area Network) support utilizing both fibre channel switch (FC-SW) and fibre channel Arbitrated Loop (FC-AL) topology and technology.

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

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

These storage systems incorporate the latest in RAID technology providing RAID levels 0, 1, 0+1, adaptive 3/5, and non-RAID (JBOD) 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 high availability support.

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

2 StorageWorks RAID Product Details

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

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

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

RA8000 FC controllers come with 64 MB of cache. Cache options allow for 128 MB, 256 MB and 512MB of cache in each controller. ESA12000 FC controllers come with 256 MB of cache. Cache options allow for 512MB of cache in each controller. Cache features include read ahead and mirrored write back cache for use in optimizing performance where redundant capabilities are required for high availability system configurations. In the rare event that a storage controller fails the mirrored data is immediately available to the 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. 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 supplies and a 2nd power distribution unit for AC input redundancy. StorageWorks components are “hot” swappable including storage controllers when operating as a redundant pair.

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

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

3 High Availability Application Options

This section discusses the following High Availability Application Options:

- SANworks Secure Path Version 2.1 for Sun Solaris
- VERITAS Cluster Server V1.1.1 for Solaris V2.6 and V2.7
- VERITAS FirstWatch V2.5.5 for Solaris V2.6 (**only**)

3.1 SANworks Secure Path Version 2.1 for Sun Solaris

SANworks Secure Path is a high availability software product providing continuous data access for RA8000 and ESA12000 FC storage products configured on Sun Solaris. Redundant hardware, advanced RAID technology and automated failover capability are used to enhance data availability. Secure Path effectively eliminates controllers, interconnect hardware and host bus adapters as single points of failure in the storage system.

The key to Secure Path's functionality is the capability of RA8000/ESA12000 FC dual HSG80 controllers to operate in the Multiple-bus failover mode. This failover mode allows each controller to be configured on its own bus and to process I/O independently under normal operation. Available storage units are “preferred” to one or the other of the two controllers which determines which controller path is used for access at system boot time. During runtime, storage units may be moved between paths at anytime using the Secure Path Management Tool.

Controllers in Multiple-bus failover mode monitor each other and automatically move storage units from the failed member of a controller pair. The Secure Path software detects the failure of I/O operations to complete on a failed path and automatically re-routes all traffic to the surviving path. Controller and path failovers are completed seamlessly, without process disruption or data loss.

The Secure Path management utility provides continuous monitoring capability and identifies failed paths and failed-over storage units. To facilitate static load balancing, storage units can be moved between paths using the Secure Path Management Tool.

The integration of Secure Path and StorageWorks RAID Array Fibre Channel technology provides the maximum level of fault-tolerance, data availability, and performance required for mission-critical environments.

3.2 VERITAS Cluster Server V1.1.2 for Solaris V2.6 and 2.7

VERITAS Cluster Server is a high-availability cluster solution that provides proactive management of application services. The VERITAS Cluster Server Java-based management GUI provides a single point of management for the entire cluster configuration. VERITAS Cluster Server:

- Is architecture-independent supporting up to 32 servers and compatible with single server, parallel and distributed applications
- Supports cascading and multi-directional application failover and application services can be manually migrated to alternate servers for maintenance purposes
- Provides a comprehensive availability management solution designed to minimize both planned and unplanned down time
- Can accommodate new and existing Sun and third party hardware and can be easily adapted to meet site-specific requirements
- Minimizes deployment time with features such as a self-replicating single configuration file per configuration, installation wizards and templates, auto-discovery of new cluster resources, off-the-shelf agents to fault manage the industry's most popular applications
- Provides a unique global broadcast mechanism, ClusterStat, to enable robust and high performance interserver communication for fault management and heartbeat monitoring
- Provides scalability with a fully multi-threaded availability engine that leverages a multi-level fault detection scheme. This scalability allows support of the industry's largest commercial cluster configurations today, with planned support in the future for hundreds of cluster servers

3.3 VERITAS FirstWatch V2.2.5.1 for Solaris V2.6 (only)

VERITAS FirstWatch for SPARC/Solaris maximizes the overall availability of business-critical servers, applications, and data. Unrivaled in its ease of configuration and management, it can be successfully deployed within hours. When it is combined with other VERITAS products, such as VERITAS Volume Manager and VERITAS File System, the products work together to deliver a comprehensive availability solution that provides greater flexibility, superior performance, high availability, and easier manageability for the enterprise.

VERITAS FirstWatch:

- is simple to install and configure, and can be easily customized to meet the needs of any environment.
- automates the recovery of application services, configured in service groups, from failures and provides for the manual migration of service groups to other nodes for maintenance purposes.
- automates application service failover between multiple active peer systems or one or more production servers with a hot standby system. Up to 8 nodes are supported. Support for service level failover (SLF) eases data center consolidation projects. Identical hardware is not required.
- simplifies failover administration and management, reducing the likelihood of operator error.
- can be implemented using off-the-shelf hardware, software, and network components, supporting the use of existing components for investment preservation or the latest products for new functionality.

4 High Availability Configuration Information

This section provides information about the following topics:

- Reference Documentation
- Configuration Guidelines (Maximums, Assumptions, and General Notes)

4.1 Reference Documentation

Table 1 lists related documents supporting RA8000 and ESA12000 Fibre Channel Storage Systems.

Table 1 Reference Documentation

Topic	Document Title	Order Number
<i>Solution Kit Overview</i>	<i>Compaq StorageWorks Solution Software Kit Overview</i>	<i>EK-SOLSR-AA</i>
<i>RA8000/ESA12000 Configurations Application Note</i>	<i>RA8000/ESA12000 Configurations for Sun Solaris</i>	<i>EK-FCALS-AA</i>
<i>RA8000/ESA12000 High Availability Configurations for Sun Solaris Application Note - Loop</i>	<i>RA8000/ESA12000 FC-AL High Availability Configurations for Sun Solaris</i>	<i>EK-FCALH-AA</i>
<i>Storage Area Network Configurations for RA8000/ESA12000 on Sun Solaris</i>	<i>Storage Area Network Configurations for RA8000 / ESA12000 on Sun Solaris</i>	<i>EK-SMA37-AN</i>
<i>Release Notes</i>	<i>Compaq StorageWorks Release Notes, RA8000/ESA12000 HSG80 Solution Software V8.5 for Sun Solaris</i>	<i>AA-RFBPE-TE</i>
<i>Quick Setup Guide</i>	<i>RAID Array 8000/ESA12000 Fibre Channel Storage Subsystem for Sun Solaris</i>	<i>AA-RFBNA-TE</i>
<i>Sun Solaris RA8000/ESA12000 Installation</i>	<i>RA8000/ESA12000 HSG80 Fibre Channel Solution Software V8.5 for Solaris - Installation Reference Guide</i>	<i>AA-RFBQE-TE</i>
<i>Secure Path User Guide</i>	<i>Compaq SANworks Secure Path V2.1 for Sun Solaris - Installation and Reference Guide</i>	<i>AA-RKYDC-TE</i>
<i>Secure Path Release Notes</i>	<i>Compaq SANworks Release Notes for Secure Path V2.1 for Sun Solaris</i>	<i>AA-RKYEB-TE</i>
<i>HSG80 Array Controller CLI Reference Guide</i>	<i>HSG80 Array Controller ACS Version 8.5 CLI Reference Guide</i>	<i>EK-HSG85-RG</i>
<i>HSG80 Array Controller Maintenance</i>	<i>HSG80 Array Controller ACS Version 8.5 Configuration Guide</i>	<i>EK-HSG85-CG</i>
<i>StorageWorks Command Console</i>	<i>Command Console V2.2 (HSG80) for RAID Array 8000/ESA12000, User's Guide</i>	<i>AA-RFA2D-TE</i>
<i>Warranty Information</i>	<i>COMPAQ StorageWorks Warranty</i>	<i>EK-HSXSU-WC</i>
<i>HSG80 Array Controller Parts</i>	<i>HSG80 Array Controller Illustrated Parts Map</i>	<i>EK-HSG80-MP</i>

Solaris is a registered trademark of Sun Microsystems Inc.

Intel is a registered trademark of Intel Corporation.

Windows NT and MSCS are trademarks of Microsoft Corporation.

SANWORKS, STORAGEWORKS and the Compaq Logo are trademarks of Compaq Computer Corporation.

VERITAS Cluster Server, VERITAS FirstWatch are registered trademarks of the VERITAS Software Corporation.

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

4.2 General Configuration Guidelines

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

4.2.1 Configuration Maximums

The HSG80 Controller	
1. Maximum	128 visible LUNS (64 per port)
2. Maximum	200 assignable unit numbers
3. Maximum	72 physical devices
4. Maximum	20 RAID-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 MirrorSet
9. Maximum	14 members per RAID-5 StorageSet
10. Maximum	24 members per StripeSet
11. Maximum	48 physical devices per Striped-MirrorSet
12. Maximum	64 Connections for a dual-controller pair; Refer to the HSG80 Array Controller ACS Version 8.5 - Configuration Guide (Chapter 1) for connection information.
13. Maximum	9 Characters per Connection Name

The Solaris Operating System	
1. Maximum	256GB file system per LUN (Solaris Versions: V2.6, or V2.7 (32-bit mode only))
2. Maximum	64 LUNs per FC Target
3. Patches	Refer to <i>RA8000/ESA12000 HSG80 Solution Software V8.5F for Sun Solaris Reference Guide (AA-RFBQE-TE)</i> for required patch levels.

4.2.1.1 SAN Concepts

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

SAN Management

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

- SAN Storage and FC-Switch Management

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

- SAN Access Management

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

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

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

Additional SAN information is available in the documentation listed in Table 1.

4.2.1.2 SAN (Storage Area Networks) Maximums	
1. Minimum	2 meters per optical cable segment.
2. Maximum	500 meters per optical cable segment, 50 micron, multi-mode using short-wave length laser GBICs; 1.5 km total distance nominal, 2.0 km worst case (single fault re-route) between server and storage.
	200 meters per optical cable segment, 62.5 micron, multi-mode using short-wave length laser GBICs; 600 meters total distance nominal, 800 meters worst case (single fault re-route) between server and storage.
	10 km per Inter-Switch Link (ISL) cable segment, 9 micron, single-mode, using long-wave length laser GBICs ; 11 km total distance nominal, 21 km worst case (single fault re-route) between server and storage. 10 km links are only supported in switch to switch configurations.
3. Maximum	Number of ISLs between switches depend on the number of available ports, required redundancy protection and data loads from servers to storage systems. A reasonable minimum is 2 links for redundancy.
4. Maximum	1 Unique Domain Number for each switch in a single fabric.
5. Maximum	16 fibre channel Host Based Adapters (HBA) per server. Currently 2 pairs of adapters per server under Secure Path V2.1
6. Maximum	4 fibre channel switches in a single fabric. ie., SAN with all switches interconnected.
7. Maximum	A hop is defined as one (1) or more connections between two FC switches. One (1) switch hop between servers and storage, maximum of two (2) switch hops (worst case) with a single fault. A fabric can re-route to a 2-hop path on a single link failure. 2 switches cascaded is one hop. Server to FC switch segments and FC switch to storage segments are not counted as hops.
8. Maximum	8 Solaris servers (assumes two FC HBA per server) per RA8000 or ESA12000 dual controller storage system, with a maximum of 8 servers configured on each port pair.
9. Maximum	2.6 TB of storage, assuming 72 disks, 36GB each, in a cabinet with two (2) expansion cabinets.
10. Maximum	Transparent or Multiple-Bus Failover Mode...
Maximum	4 Servers on a single Fabric with ...
Maximum	2 Storage Systems with ...
Maximum	4 Controllers with 2 Active controller ports/controller pair.
OR	
Maximum	8 Servers on two Fabrics with ...
Maximum	1 Storage System per Fabric with...
Maximum	2 Controllers with 2 Active controller ports.

4.2.2 Assumptions

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

Document References

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

4.2.3 Configuration Notes

1. All single controller storage configurations can be upgraded to dual controller configurations. All High Availability configurations discussed in this document assume dual controller configurations.
2. Controller pairs must be configured for either Dual Redundant Failover Mode *or* Multiple-bus Failover Mode. When controllers are in Dual Redundant Failover Mode, this is referred to as "Transparent Failover Mode".

NOTE: This Application Note describes configurations with either Transparent Failover or Multiple-Bus Failover. Converting from Transparent Failover Mode to Multiple-bus Transparent Failover mode (and vice-versa) is documented in Appendix A.

NOTE: For configurations that use the Secure Path application, the RAID storage controllers must be configured in Multiple-Bus Failover.

3. All configuration diagrams show a front 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 providing Fibre Channel loop and fabric topology support at the controller level.
6. All configurations **require** RA8000/ESA12000 FC Solution Software V8.5 for the Sun Solaris operating system. This software supports both the 32-bit and 64-bit FC host bus adapters.
 - Secure Path V2.1 supports the 64-bit S-bus and 32-bit PCI adapters only.
7. All configurations utilize short wavelength lasers and multi-mode fibre channel optical cables.
8. All configurations **require** the World Wide Port Names. The WWPN is derived from the World Wide Node Name of the RAID storage. The World Wide Node Name is located on the controller pair and is a 16 character value in the format of NNNN-NNNN-NNNN-NNNN. This hexadecimal value is the name of the RAID storage. By using the CLI command
 HSG80> SHOW THIS_CONTROLLER
 and locating the Reported PORT_ID section, the World Wide Port name(s) can be located as shown below:

Host PORT_1:

```
Reported PORT_ID = 5000-1FE1-0000-0D43
PORT_1_TOPOLOGY = FABRIC (fabric up)
Address          = 021200
```

9. All configurations **require** the Connection Name *Operating System* parameter set to "SUN". (CLI: "SET *connection-name* OPERATING_SYSTEM = SUN")

10. By default the HSG80 Command Console LUN (CCL) is enabled. This may be disabled.

(CLI: “SET *controller* NOCOMMAND_CONSOLE_LUN”)

UNIT NUMBERS

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. In Multiple-Bus Failover Mode, all 200 unit numbers may be visible to all ports. Connection offsets are used to restrict access/visibility.

Solaris supports up to 64 LUNs per target. The RAID platform kit installation defines 16 LUNs per target. This is considered more than sufficient for general use. These LUNs are defined in the file: `/kernel/drv/sd.conf`

11. For configurations with more than one server on the same fabric:

Use the SET *connection-name Unit_Offset* value to set the range of unit numbers to be visible from each Server (CLI: “SET *connection-name* UNIT_OFFSET = n”) This will provide a first level of access control from Server to Storage.

It is recommended that a second level of access control be defined that creates a specific **access path** from Server to storage unit. 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*”)

12. Record the FC HBA Worldwide Identification (WWID) address. Some HBAs display the WWID on the HBA itself. For Solaris, the WWID is displayed when a host system boots and the Fibre Channel driver attaches to the adapter. These HBA addresses will also be found in the `/var/adm/messages` file. The format is `nnnn-nnnn-nnnn-nnnn`. This WWID should be recorded.

13. Use the recorded WWID to identify the connection name displayed from the CLI at the controller using the CLI: “SHOW CONNECTIONS” command.

14. Rename connection names. When the RAID controller detects a path from a HBA to a controller port, it creates a connection and connection names are automatically added to the controller connection name table as “!NEWCONnn” (where nn is a number from 0 to 99). Currently, the maximum number of connections is 64. Once a connection is identified using the FC HBA WWID, rename the connection 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.).

15. Each active controller host port presents one SCSI Target ID with up to sixty-four (64) LUNs to each FC HBA on the same fabric. LUNs (logical units) can consist of single disks (JBOD), a storageset of multiple disks or a partition (up to 64), configured for a specific RAID level of 0, 0 + 1, or 3/5.

16. **Thinking About Connections.** There are two points of view when speaking of fibre channel connections.

- The first view of connections is taken from the server side. The server has a configuration file that allows the operating system to identify targets and logical units used for I/O. 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 6 targets (WWPNs) on a fibre channel fabric, the configuration file may/may not list all 6 targets. In short, a specific server may not look for or try to access a specific target.

- The second view of connections is from the controller side. In a fabric, the controller knows about each host-adapter combination (HAC) that it communicates with because they were defined either at or after installation. The controller creates one connection on each active controller port for each HAC.

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

In the configurations of this application note, we use the server's point of view to limit access to certain targets. This means we have unused connections. This also means that when implementing a configuration, we only need to identify the connections that we are using, not all of them. Thus, the naming of connections is done for only those connections that are active in the specific configuration. We simply disregard the unused connections.

NOTE

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

5 High Availability Configuration Topologies and Modes

This section provides a detailed description of the eight most common RA8000/ESA12000 FC Sun Solaris high availability configurations, as follows:

- 5.1 One Server, Secure Path V2.1, Multiple-Bus Failover
- 5.2 Two Servers, VERITAS Cluster Server, Transparent Failover
- 5.3 Four Servers, VERITAS Cluster Server, Transparent Failover, One Shared Storage Enclosure
- 5.4 Two Servers, Secure Path V2.1, VERITAS Cluster Server, Multiple-Bus Failover
- 5.5 Four Servers, VERITAS Cluster Server, Secure Path, Multiple-Bus Failover
- 5.6 Up to Four Servers, Secure Path V2.1, Multiple-Bus Failover, One Shared Storage Enclosure
- 5.7 Up to Four Servers, Secure Path V2.1, Multiple-Bus Failover, Two Shared Storage Enclosures
- 5.8 Up to Eight Servers, Secure Path V2.1, Multiple-Bus Failover, Two Shared Storage Enclosures

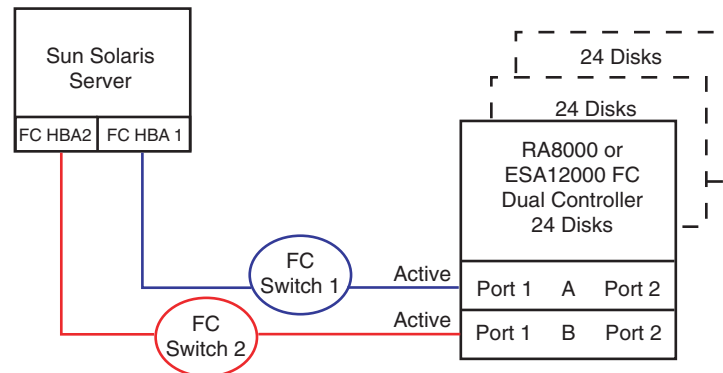
NOTE:

All configurations have one or more HSG80 Storage Systems with dual-controllers bound in either Transparent Failover or Multiple_bus Failover Mode.

5.1 HA Configuration 1 – One Server, Secure Path, Multiple-Bus Failover

HA Configuration 1 (Figure 1) represents an entry-level RA8000 or ESA12000 high availability configuration that can be expanded to support additional storage capacity or performance needs as required. The configuration consists of a single RA8000 or ESA12000 storage enclosure with two HSG80 array controllers configured in Multiple-Bus Failover mode. The configuration utilizes two separate paths (two fabrics) to the storage. In the event of a failure of any component within the access path – HBA, cables, switch, or controller host port interface, all LUNs on the failing path will become accessible through the alternate path. The initial configuration supports up to 24 disks. Support for an additional 48 disks can be added using two expansion enclosures. Additional servers can be added to this configuration as required using the existing switches. (See Configuration 6)

Figure 1 HA Configuration 1



SHR-1668

HA Configuration 1 Notes

- Two Fibre Channel Fabrics
- RA8000/ESA12000 HSG80 Solution Software V8.5 for Sun Solaris installed
- Secure Path V2.1 for Sun Solaris application software installed
- Dual HSG80 controllers configured in Multiple-Bus Failover Mode
- Up to 500 meters per cable segment
- Up to 72 disks with 2 optional expansion cabinets
- Up to 64 LUNs with 2 active controller host ports. Suggested controller unit number assignment: Port 1 units D0 – D63; Port 2 is not used in this configuration.
- Set the Unit Preferred Path. Use the Preferred Path unit parameter to assign units to specific controllers. Units should be distributed equally across controllers for load balancing.

Suggested Controller Unit Number Assignment:

Prefer units D0 – D31 to Controller A

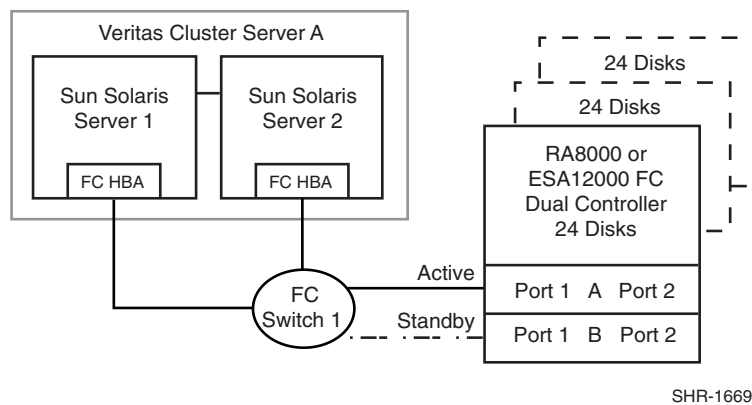
Prefer units D32 – D63 to Controller B

- In this configuration 2 logical connections are available
- Rename the connection names, i.e., SERV1P1A, SERV1P1B
- Update `/kernel/drv/sd.conf` as required to account for the required LUNs

5.2 HA Configuration 2 - Two Servers, VERITAS Cluster Server, Transparent Failover

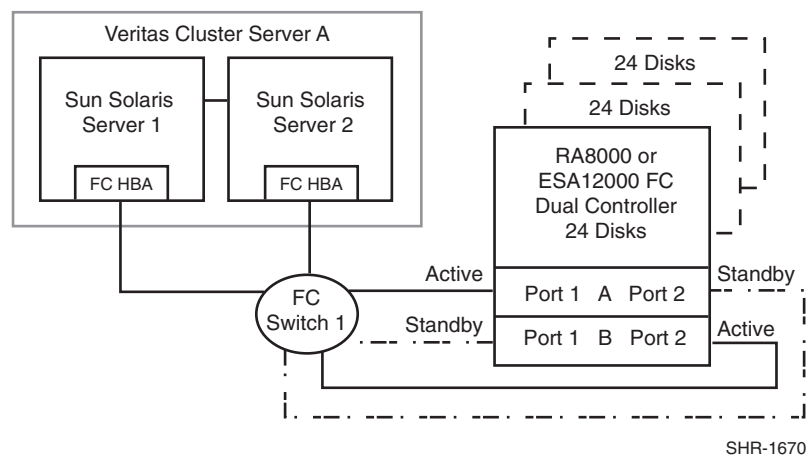
HA Configuration 2 (Figure 2) provides dual redundant controllers configured in Transparent Failover mode. The configuration consists of a single RA8000 or ESA12000 storage enclosure with two HSG80 array controllers connected to two servers running VERITAS Cluster Server Application software. In normal operation Controller A provides the active path for all storagesets to the server pair through Controller A, Port 1. A second controller, Controller B, is a hot standby. In the event of a failure in Controller A, Controller B will automatically become active. All units being accessed through Port 1 of Controller A will become accessible through Port 1 of Controller B. In the event of a server failure, VERITAS Cluster Server will transfer access to units through the surviving server. In the initial configuration a single controller host port (Port 1) is active providing up to 64 LUNs across up to 24 disks. An additional 48 disks can be added using two expansion enclosures.

Figure 2 HA Configuration 2



With the addition of two fibre channel cables, Controller B, can become an active controller to increase performance and double LUN capacity. Two additional fibre channel cables, one connected to Controller A Port 2, and one connected to Controller B Port 2, result in the configuration shown in Figure 3. This configuration provides full controller redundancy in the event that Controller A or Controller B fails. In either failover scenario the surviving controller will transition a standby port to active providing access to all units in the VERITAS Cluster Server pair.

Figure 3 HA Configuration 2 with Increased Performance and Doubled LUN Capacity



HA Configuration 2 Notes

Figure 2

- Single Fibre Channel Fabric
- Dual HSG80 controllers configured in Transparent Failover Mode
- RA8000/ESA12000 HSG80 Solution Software V8.5 for Sun Solaris installed
- VERITAS Cluster Server application software installed
- Up to 500 meters per cable segment
- Up to 72 disks with 2 optional expansion cabinets
- Up to 64 LUNs with 1 active controller host port.
- Suggested controller unit number assignment: Port 1 units D0 – D63.
Two (2) logical connections are available
- Rename the connection names, i.e., SERVER1P1, SERVER1P2
- Update */kernel/drv/sd.conf* as required to account for the required LUNs

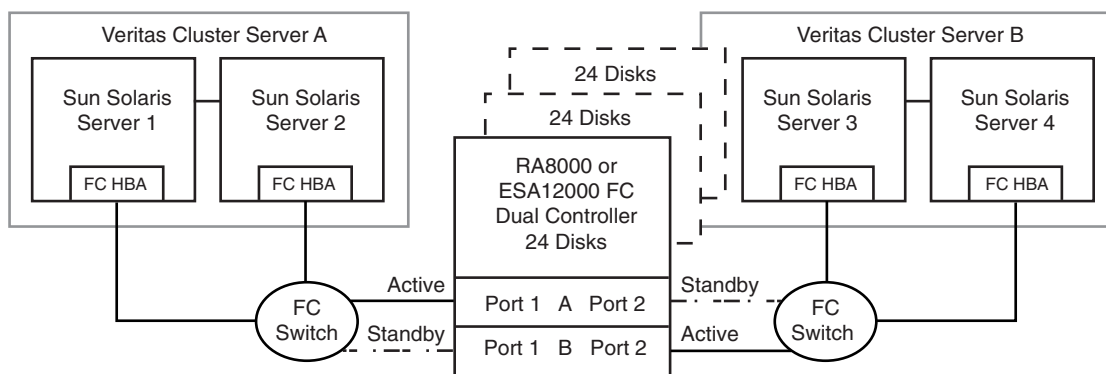
Figure 3

- Single Fibre Channel Fabric
- Dual HSG80 controllers configured in Transparent Failover Mode
- RA8000/ESA12000 HSG80 Solution Software V8.5 for Sun Solaris installed
- VERITAS Cluster Server application software installed
- Up to 500 meters per cable segment
- Up to 72 disks with 2 optional expansion cabinets
- Up to 128 LUNs with 2 active controller host ports. Suggested controller unit number assignment:
Port 1 units D0 – D63, Port 2 units D100 – D163. (Four (4) logical connections are available).
- Utilization of controller host Port 2 requires two additional fibre channel cables
- Setting the controller pair in Transparent Failover mode will, by default 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
- Rename the connection names, i.e., SERVER1P1, SERVER1P2, SERVER2P1, SERVER2P2
- Update */kernel/drv/sd.conf* as required to account for the required LUNs

5.3 HA Configuration 3 - Four Servers, VERITAS Cluster Server, Transparent Failover, Shared Storage Enclosure

The dual-controller configuration (Figure 4) utilizes two pairs of controller host ports to provide storage for a second server pair under VERITAS Cluster Server control. Each Server pair can be configured with up to 64 LUNs. The two VERITAS Cluster Servers share the storage enclosure in that each cluster has exclusive access to its own LUNs through the use of mutually exclusive LUN numbering. The controllers are configured in Transparent Failover mode providing controller redundancy for both clusters. In the event of a controller failover, the surviving controller will continue to provide full access to units in both VERITAS clusters. In the event of a server failure, VERITAS Cluster Server failover will be invoked to provide access to units through the surviving server. In this configuration two active controller ports provide up to 128 LUNs (64 per server pair) across up to 24 disks. An additional 48 disks can be added using two expansion enclosures to increase storage capacity.

Figure 4 HA Configuration 3



SHR-1671

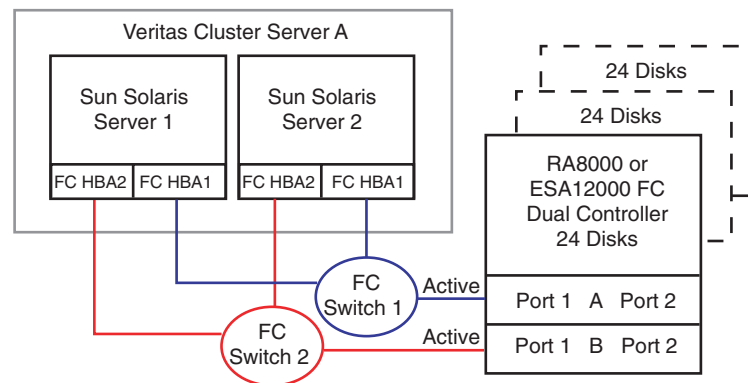
HA Configuration 3 Notes

- RA8000/ESA12000 HSG80 Solution Software V8.5 for Sun Solaris installed
- VERITAS Cluster Server application software for two server pairs installed
- Shared storage enclosure
- Two Fibre Channel Fabrics
- Dual HSG80 controllers configured in Transparent Failover Mode
- Up to 500 meters per cable segment
- Up to 72 disks with 2 optional expansion enclosures
- Up to 128 LUNs total/64 per server pair with 2 active controller host ports.
- Suggested controller unit number assignment:
Server pair A, Controller A Port 1 units D0 – D63
Server pair B, Controller B Port 2 units D100 – D163
- 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
- In this configuration 4 logical connections are available
- Rename the connection names, i.e., SERVER1P1, SERVER2P1, SERVER3P2, SERVER4P2
- Update `/kernel/drv/sd.conf` as required to account for the required LUNs

5.4 HA Configuration 4 - Two Servers, Secure Path, VERITAS Cluster Server, Multiple-Bus Failover

HA Configuration 4 (Figure 5) combines the benefits of VERITAS Cluster Server pairing, Secure Path for Solaris software and controller redundancy to provide the highest level of availability. The configuration utilizes two separate physical paths to the storage. In the event of a failure of any component within the access path – FC HBA, cables, switch, or controller host port interface, all units on the failing path will become accessible through the alternate path. In the event of a server failure, VERITAS Cluster Server failover will be invoked to provide access to units through the surviving server.

Figure 5 HA Configuration 4



HA Configuration 4 Notes

- Two Fibre Channel Fabrics
 - RA8000/ESA12000 HSG80 Solution Software V8.5 for Sun Solaris installed
 - Secure Path V2.1 for Solaris application software installed
 - Dual HSG80 controllers configured in Multiple-Bus Failover Mode
 - Up to 500 meters per cable segment
 - Up to 72 disks with 2 optional expansion cabinets
 - Up to 64 LUNs with 2 active controller host ports.
 - Suggested controller unit number assignment: Port 1 units D0 – D63
 - Set the Unit Preferred Path. Use the Preferred Path unit parameter to assign units to specific controllers. Units should be distributed equally across controllers for load balancing.
- For example:

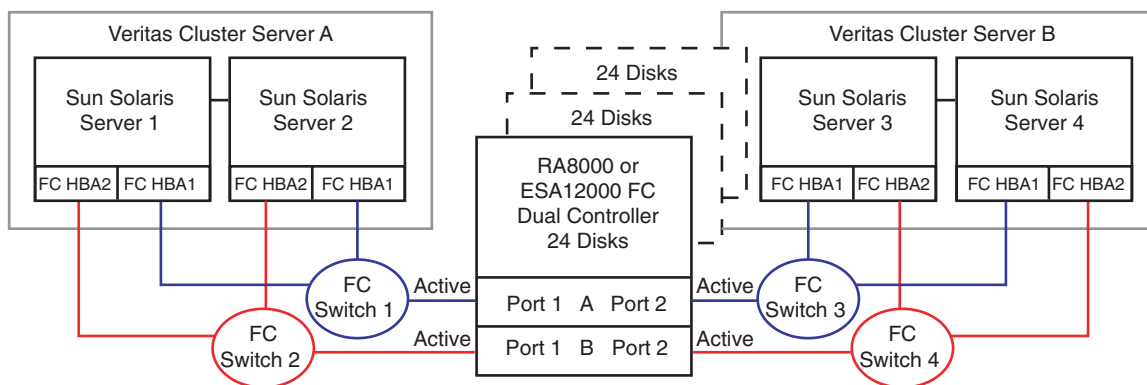
Prefer units D0 – D31 to Controller A
Prefer units D32 – D63 to Controller B

- In this configuration 4 logical connections are available
- Rename the connection names, i.e., SERV1P1A, SERV1P1B, SERV2P1A, SERV2P1B
- Update `/kernel/drv/sd.conf` as required to account for the required LUNs

5.5 HA Configuration 5 - Four Servers, Secure Path, VERITAS Cluster Server, Multiple-Bus Failover, Two Independent Server Pairs, Shared Storage Enclosure

HA Configuration 5 (Figure 6) builds on configurations 3 and 4, combining the benefits of VERITAS Cluster Server software, Secure Path for Solaris software, controller redundancy, and sharing of the storage enclosure. The configuration provides two independent server pairs with two separate physical paths to the storage through two pairs of controller host ports. In the event of a failure of any component within the access path – FC HBA, cables, switch, or controller host port interface, all units on the failing path will become accessible through the alternate path. In the event of a server failure, cluster failover will be invoked to provide access to units through the surviving server. Each VERITAS Cluster Server pair can be configured with up to 64 LUNs. The two server pairs share the storage enclosure as each server pair has exclusive access to its own LUNs through the use of mutually exclusive LUN numbering. An additional 48 disks can be added using two expansion enclosures to increase storage capacity.

Figure 6 HA Configuration 5



SHR-1673

HA Configuration 5 Notes

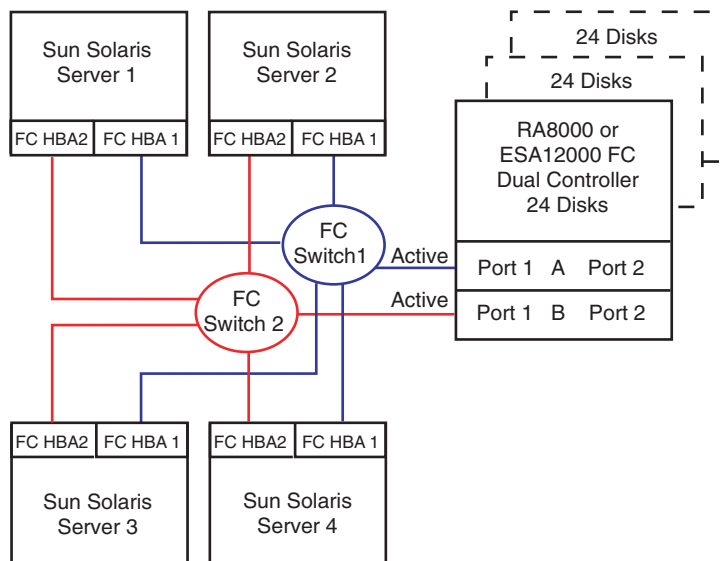
- Four Fibre Channel Fabrics
- RA8000/ESA12000 HSG80 Solution Software V8.5 for Sun Solaris installed
- Secure Path V2.1 for Sun Solaris and VERITAS Cluster Server application software installed
- Dual HSG80 controllers configured in Multiple-Bus Failover Mode
- Up to 500 meters per cable segment
- Up to 72 disks with 2 optional expansion cabinets
- Up to 128 LUNs total/64 server pair with 4 active controller host ports.
- Suggested controller unit number assignment:
Port 1 units D0 – D63
Port 2 units D64 – D127
- Set the Unit Preferred Path. Use the Preferred Path unit parameter to distribute units equally across controllers for load balancing, as follows:
Prefer units D0 – D31 to Controller A
Prefer units D32 – D63 to Controller B
Prefer units D64 – D95 to Controller A
Prefer units D96 – D127 to Controller B
(In this configuration 8 logical connections are available)
- Rename the connection names, i.e., SERV1P1A, SERV1P1B, SERV2P1A, SERV2P1B, SERV3P2A, SERV3P2B, SERV4P2A, SERV4P2B

- Set exclusive unit access for each unit to specific server connection names, i.e., for unit D0 ENABLE = SERV1P1A, SERV1P1B, SERV2P1A, SERV2P1B, for unit D32 ENABLE = SERV3P2A, SERV3P2B, SERV4P2A, SERV5P2B.
- With SWCC, new units are enabled on all known connection names. The *connection_names* must be explicitly disabled to prevent access.
- Set the connection name unit offset values for each connection as shown below
Note: The suggested unit numbering and connection parameters results in the units being distributed equally across clusters/controller host ports as follows:
 - Server 1: Port 1 units D0 – D63
Connection names, unit offset values: SERV1P1A, 0. SERV1P1B, 0
 - Server 2: Port 1 units D0 – D63
Connection names, unit offset values: SERV2P1A, 0. SERV2P1B, 0
 - Server 3: Port 2 units D64 – D127
Connection names, unit offset values: SERV3P2A, 64. SERV3P2B, 64
 - Server 4: Port 2 units D64 – D127
Connection names, unit offset values: SERV4P2A, 64. SERV4P2B, 64
- Update */kernel/drv/sd.conf* as required to account for the required LUNs

5.6 HA Configuration 6 - Up to Four Servers, Secure Path, Multiple-Bus Failover, Shared Storage Enclosure

HA Configuration 6 (Figure 7) builds on Configuration 1 providing storage for up to four separate (non-cooperating) servers on two fibre channel fabrics. The four servers share the storage enclosure in that each server has exclusive access to its own set of LUNs. This is made possible by the setting of explicit “connection name access” at the unit level and using connection name “unit offsets” to give each server an accessible range of unit numbers. Additional levels of access control are necessary for this configuration because all servers are sharing the two fibre channel fabrics. The configuration utilizes two separate paths to the storage from each server. In the event of a failure of any component within the access path – HBA, cables, switch, or controller host port interface, all LUNs on the failing path will become accessible through the alternate path.

Figure 7 HA Configuration 6



SHR-1674

HA Configuration 6 Notes

- Two Fibre Channel Fabrics
- Shared storage enclosure
- RA8000/ESA12000 HSG80 Solution Software V8.5 for Sun Solaris installed
- Secure Path V2.1 for Solaris application software installed
- Dual HSG80 controllers configured in Multiple-Bus Failover Mode
- Up to 500 meters per cable segment
- Up to 72 disks with 2 optional expansion cabinets
- Up to 128 LUNs, 32 available to each server with 2 active controller host ports.
- Set the Unit Preferred Path. Use the Preferred Path unit parameter to distribute units equally across controllers/ports for load balancing
 - Prefer units D0 – D15 to Controller A
 - Prefer units D16 – D31 to Controller B
 - Prefer units D32 – D47 to Controller A
 - Prefer units D48 – D63 to Controller B
 - Prefer units D64 – D79 to Controller A
 - Prefer units D80 – D95 to Controller B
 - Prefer units D96 – D111 to Controller A
 - Prefer units D112 – D127 to Controller B
- In this configuration 8 logical connections are available

- Rename the connection names, i.e., SRV1A1P1A, SRV1A2P1B, SRV2A1P1A, SRV2A2P1B, SRV3A1P1A, SRV3A2P1B, SRV4A1P1A, SRV4A2P1B
- Set exclusive unit access for each unit to specific server connection names, i.e., for unit D0 ENABLE = SRV1A1P1A, SRV1A2P1B, for unit D32 ENABLE = SRV2A1P1A, SRV2A2P1B.
- If you are using SWCC to create units, by default units are enabled on all known connection names. You must explicitly disable connection names for units to prevent access.
- Set the connection name unit offset values for each connection as shown below.

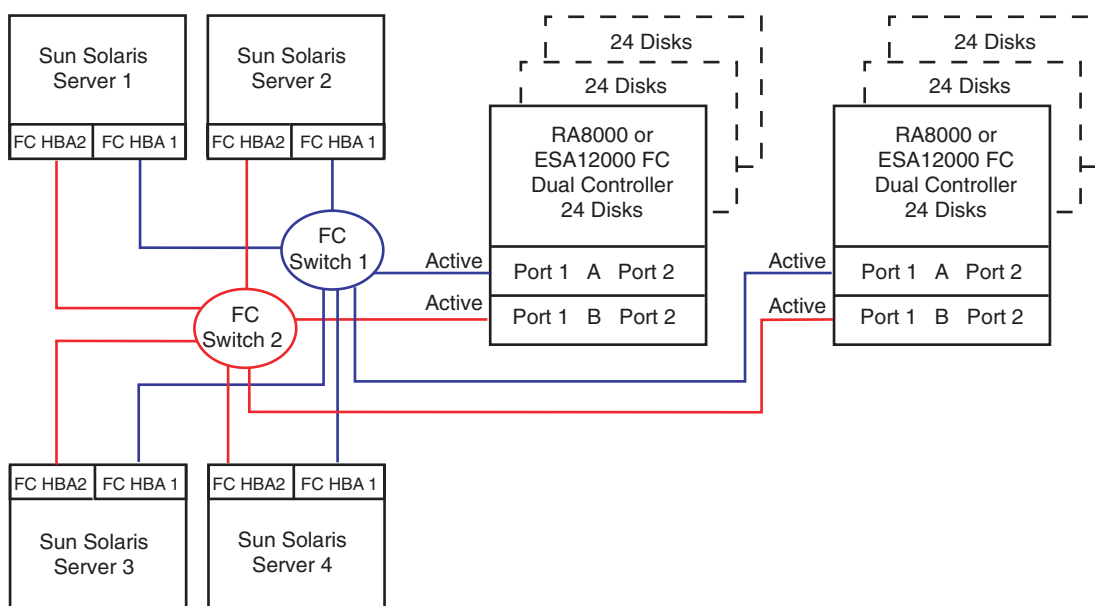
Note: The suggested unit numbering and connection parameters results in the units being distributed equally across servers and controller host ports, as follows:

- Server 1: Port 1 units D0 – D31
Connection names, unit offset values: SRV1A1P1A, 0. SRV1A2P1B, 0
 - Server 2: Port 1 units D32 – D63
Connection names, unit offset values: SRV2A1P1A, 32. SRV2A2P1B, 32
 - Server 3: Port 1 units D64 – D95
Connection names, unit offset values: SRV3A1P1A, 64. SRV3A2P1B, 64
 - Server 4: Port 1 units D96 – D127
Connection names, unit offset values: SRV4A1P1A, 96. SRV4A2P1B, 96
- Update */kernel/drv/sd.conf* as required to account for the required LUNs

5.7 HA Configuration 7 - Up to Four Servers, Secure Path, Multiple-Bus Failover, Two Shared Storage Enclosures

HA Configuration 7 (Figure 8) builds on Configuration 6 by adding storage capacity using a second enclosure and second pair of controllers. The second pair of controllers are connected to two fibre channel fabrics providing redundant paths for up to four separate (non-cooperating) servers. The four servers share both storage enclosures in that each server has exclusive access to its own LUNs in each enclosure through the use of mutually exclusive LUN numbering, explicit unit host access, and the use of connection unit offsets. Additional levels of access control are necessary for this configuration because all servers are sharing the two fibre channel fabrics. The configuration utilizes two separate paths to the storage from each server. In the event of a failure of any component within the access path – HBA, cables, switch, or controller host port interface, all LUNs on the failing path will become accessible through the alternate path.

Figure 8 HA Configuration 7



SHR-1675

HA Configuration 7 Notes

- Two Fibre Channel Fabrics
- Shared storage enclosure
- RA8000/ESA12000 HSG80 Solution Software V8.5 for Sun Solaris installed
- Secure Path V2.1 for Solaris application software installed
- Two HSG80 controller pairs configured in Multiple-Bus Failover Mode
- Up to 500 meters per cable segment
- Up to 144 disks with 4 optional expansion cabinets
- Up to 256 LUNs, 64 available to each server with 4 active controller host ports. Refer to Configuration 6 for a suggested way to define the units that distributes them equally across servers/host ports. Unit configurations and numbering can be identical for each storage enclosure, as each server will see each active controller host port as a unique target ID with a unique set of LUNs.

- For each storage enclosure (controller pair), set the Unit Preferred Path. Use the Preferred Path unit parameter to distribute units equally across controllers/ports for load balancing.

For Example:

- Prefer units D0 – D15 to Controller A
 - Prefer units D16 – D31 to Controller B
 - Prefer units D32 – D47 to Controller A
 - Prefer units D48 – D63 to Controller B
 - Prefer units D64 – D79 to Controller A
 - Prefer units D80 – D95 to Controller B
 - Prefer units D96 – D111 to Controller A
 - Prefer units D112 – D127 to Controller B
- In this configuration 16 logical connections are available, 8 in each enclosure
 - For each storage enclosure (controller pair), rename the connection names, i.e., SRV1A1P1A, SRV1A2P1B, SRV2A1P1A, SRV2A2P1B, SRV3A1P1A, SRV3A2P1B, SRV4A1P1A, SRV4A2P1B
 - Set exclusive unit access for each unit to specific server connection names, i.e., for unit D0 ENABLE = SRV1A1P1A, SRV1A2P1B, for unit D32 ENABLE = SRV2A1P1A, SRV2A2P1B. When using SWCC to create units, by default, units are enabled on all known connection names. You must explicitly disable connection names for units to prevent access.
 - Set the connection name unit offset values for each connection as shown below
 - Suggested unit numbering and connection parameters would result in the units being distributed equally across enclosures and controller host ports as follows:

Storage Enclosure 1

Server 1: Port 1 units D0 – D31

Connection names, unit offset values: SRV1A1P1A, 0. SRV1A2P1B, 0

Server 2: Port 1 units D32 – D63

Connection names, unit offset values: SRV2A1P1A, 32. SRV2A2P1B, 32

Server 3: Port 1 units D64 – D95

Connection names, unit offset values: SRV3A1P1A, 64. SRV3A2P1B, 64

Server 4: Port 1 units D96 – D127

Connection names, unit offset values: SRV4A1P1A, 96. SRV4A2P1B, 96

Storage Enclosure 2

Server 1: Port 1 units D0 – D31

Connection names, unit offset values: SRV1A1P1A, 0. SRV1A2P1B, 0

Server 2: Port 1 units D32 – D63

Connection names, unit offset values: SRV2A1P1A, 32. SRV2A2P1B, 32

Server 3: Port 1 units D64 – D95

Connection names, unit offset values: SRV3A1P1A, 64. SRV3A2P1B, 64

Server 4: Port 1 units D96 – D127

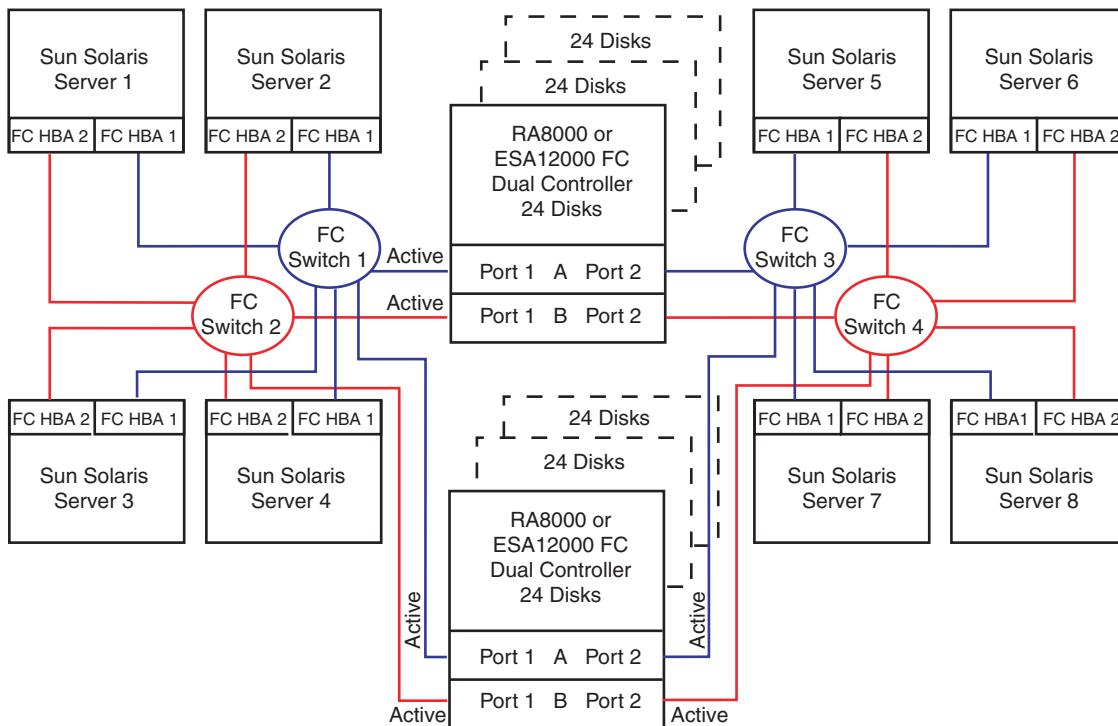
Connection names, unit offset values: SRV4A1P1A, 96. SRV4A2P1B, 96

- Update `/kernel/drv/sd.conf` as required to account for the required LUNs

5.8 HA Configuration 8 - Up to Eight Servers, Secure Path, Multiple-Bus Failover, Two Shared Storage Enclosures

Configuration 8 (Figure 9) builds on Configurations 6 and 7 utilizing two dual controller pairs on four fabrics to provide storage path and controller redundancy for up to eight separate (non-cooperating) servers. The eight servers share both storage enclosures in that each server has exclusive access to its own LUNs in each enclosure through the use of mutually exclusive LUN numbering, explicit unit host access, and the use of connection unit offsets. Additional levels of access control are necessary for this configuration because each group of four servers is sharing two fibre channel fabrics. The configuration utilizes two separate paths to the storage from each server. In the event of a failure of any component within the access path – HBA, cables, switch, or controller host port interface, all LUNs on the failing path will become accessible through the alternate path.

Figure 9 HA Configuration 8



SHR-1676

HA Configuration 9 Notes

- Four Fibre Channel Fabrics
- Shared Storage Enclosure
- RA8000/ESA12000 HSG80 Solution Software V8.5 for Sun Solaris installed
- Secure Path V2.1 for Solaris application software installed
- Two HSG80 controller pairs configured in Multiple-Bus Failover Mode
- Up to 500 meters per cable segment
- Up to 144 disks with 4 optional expansion cabinets
- Up to 128 LUNs, 32 available to each server with 8 active controller host ports. Refer to Configuration 6 for a suggested way to define the units that distributes them equally across servers/host ports. Unit configurations and numbering can be identical for each storage enclosure, as each server will see each active controller host port as a unique target ID with a unique set of LUNs.

- For each storage enclosure (controller pair), set the Unit Preferred Path. Use the Preferred Path unit parameter to distribute units equally across controllers/ports for load balancing.

Example:

Prefer units D0 – D7 to Controller A
 Prefer units D8 – D15 to Controller B
 Prefer units D16 – D23 to Controller A
 Prefer units D24 – D31 to Controller B
 Prefer units D32 – D39 to Controller A
 Prefer units D40 – D47 to Controller B
 Prefer units D48 – D55 to Controller A
 Prefer units D56 – D63 to Controller B
 Prefer units D64 – D71 to Controller A
 Prefer units D72 – D79 to Controller B
 Prefer units D80 – D87 to Controller A
 Prefer units D88 – D95 to Controller B
 Prefer units D96 – D103 to Controller A
 Prefer units D104 – D111 to Controller B
 Prefer units D112 – D119 to Controller A
 Prefer units D120 – D127 to Controller B

(In this configuration 32 logical connections are available, 16 in each enclosure).

- For each enclosure (controller pair), rename the connection names, i.e., SRV1A1P1A, SRV1A2P1B, SRV2A1P1A, SRV2A2P1B, SRV3A1P1A, SRV3A2P1B, SRV4A1P1A, SRV4A2P1B
- Set exclusive unit access for each unit to specific server connection names, i.e., for unit D0 ENABLE = SRV1A1P1A, SRV1A2P1B, for unit D16 ENABLE = SRV2A1P1A, SRV2A2P1B. When using SWCC to create units, by default, units are enabled on all known connection names. You must explicitly disable connection names for units to prevent access.
- Set the connection name unit offset values **at each enclosure** for each connection as shown below. Suggested unit numbering and connection parameters would result in the units being distributed equally across enclosures and controller host ports as follows:
 - Server 1: Port 1 units D0 – D15
 Connection names, unit offset values: SRV1A1P1A, 0. SRV1A2P1B, 0
 - Server 2: Port 1 units D16 – D31
 Connection names, unit offset values: SRV2A1P1A, 16. SRV2A2P1B, 16
 - Server 3: Port 1 units D32 – D47
 Connection names, unit offset values: SRV3A1P1A, 32. SRV3A2P1B, 32
 - Server 4: Port 1 units D48 – D63
 Connection names, unit offset values: SRV4A1P1A, 48. SRV4A2P1B, 48
 - Server 5: Port 2 units D64 – D79
 Connection names, unit offset values: SRV5A1P2A, 64. SRV5A2P2B, 64
 - Server 6: Port 2 units D80 – D95
 Connection names, unit offset values: SRV6A1P2A, 80. SRV6A2P2B, 80
 - Server 7: Port 2 units D96 – D111
 Connection names, unit offset values: SRV7A1P2A, 96. SRV7A2P2B, 96
 - Server 8: Port 2 units D112 – D127
 Connection names, unit offset values: SRV8A1P2A, 112. SRV8A2P2B, 112
- Update */kernel/drv/sd.conf* as required to account for the required LUNs

6 High Availability Parts List

- Sun Servers with Solaris Version 2.6 and 2.7 (**32-bit mode**)
- Quantities assumed as needed for a specific configuration.

Compaq Part #	Description
380554-001	RA8000/ESA12000 FC Solution SoftwareV8.5 for Sun Solaris (QB-65RAG-SA)
128697-B21	ACS V8.5F Controller Software (QB-6BUAA-SA)
165991-B21	Secure Path for Sun Solaris V2.1 (QB-669AE-SA)
380576-001 123503-001	PCI FC HBA (32-bit) (SWSA4-PC) SBUS FC HBA (64-bit) (DS-SWSA4-SC)
380560-B21 (Blue) 380560-B22 (Opal) 380660-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) 380630-001 (Blue) 380630-002 (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 ESA1200 w/2 pairs/dual HSG80 48 Slot 60 HZ ESA1200 w/2 pairs/dual HSG80 48 Slot 60 HZ ESA1200 w/2 pairs/dual HSG80 48 Slot 50 HZ ESA1200 w/2 pairs/dual HSG80 48 Slot 50 HZ
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
380591-B21 380591-B22 158222-B21 158223-B21 380578-B21 380578-B22 158224-B21 158225-B21	8 port switch (DS-DSGGA-AA) 8 port switch (DS-DSGGA-AC) 8 port switch (DS-DSGGB-AA) 8 port switch (DS-DSGGB-AB) 16 port switch (DS-DSGGA-AB) 16 port switch (DS-DSGGA-AD) 16 port switch (DS-DSGGB-BA) 16 port switch (DS-DSGGB-BB)

Compaq Part #	Description
380561-B21	FC Optical GBIC (DS-DXGGA-SA)
234457-B21	FC 2 Meter Optical Cable (BNGBX-02)
234457-B22	FC 5 Meter Optical Cable (BNGBX-05)
234457-B23	FC 15 Meter Optical Cable (BNGBX-15)
234457-B24	FC 30 Meter Optical Cable (BNGBX-30)
234457-B25	FC 50 Meter Optical Cable (BNGBX-50)
380691-B21	4GB UW 7200 RPM Disk (DS-RZ1CF-VW)
380595-B21	9GB UW 7200 RPM Disk (DS-RZ1DF-VW)
380694-B21	18GB UW 7200 RPM Disk (DS-RZ1EF-VW)
380588-B21	9GB UW 10000 RPM Disk (DS-RZ1DD-VW)
380589-B21	18GB UW 10000 RPM Disk (DS-RZ1ED-VW)
147599-001	36GB UW 7200RPM Disk (DS-RZ1FB-VW)
159138-001	36GB UW 1000 RPM Disk (DS-RZ1FC-VW)

Appendix A: Controller Failover Transitions

Introduction

This appendix provides the procedures for setting dual-redundant HSG80 controllers from one failover state to another. Controllers can exist in one of the following failover states: Transparent Failover (TF), Multiple-Bus Failover (MBF) and Nofailover (NF).

The controller failover state changes, discussed in the following sections, are:

- A. From Transparent Failover to NoFailover Mode
- B. From Transparent Failover to Multiple-bus Failover Mode
- C. From Multiple-bus Failover to Nofailover and then to Transparent Failover Mode

To change the failover state of a controller, connect a serial line to the top controller. This will be referenced as THIS_CONTROLLER. The second controller will be referenced as the OTHER_CONTROLLER. All commands directed to the RAID storage system (i.e., those preceded by HSG80> prompt in the following sections), must be issued through this serial connection to the top controller.

Verifying the current state of the controllers is achieved by entering:

```
HSG80> SHOW THIS_CONTROLLER
```

The display from the SHOW command has a number of sections. The information required for changing a controller failover state is located in the first section, under the “Controller:” heading. A sample display for Transparent Failover is shown below, and the failover state is identified with an arrow, “->” preceding the text of note.

Example:

```
Controller:
  HSG80 ZG83502145 Software V85F-0, Hardware E03
  NODE ID          = 5000-1FE1-0000-3350
  ALLOCATION_CLASS  = 0
  SCSI VERSION     = SCSI-2
  -> Configured for dual-redundancy with ZG80200290
  ->      In dual-redundant configuration
```

When a controller state is changed using the procedures described in one of the next sections, the display will be refreshed to verify that the change has completed successfully.

A. Changing from Transparent Failover to No Failover Mode

1. At the HSG80> prompt, enter:

```
HSG80> SET NOFAILOVER
```

This action will cause the OTHER_CONTROLLER to shutdown.

2. At the HSG80> prompt, enter:

```
HSG80> SHOW THIS_CONTROLLER
```

From the resulting display, verify that the controller state has changed to no failover.

Example:

```
Controller:
HSG80 ZG83502145 Software V85F-0, Hardware E03
NODE_ID          = 5000-1FE1-0000-3350
ALLOCATION_CLASS  = 0
SCSI_VERSION     = SCSI-2
-> Not Configured for dual-redundancy
```

3. Restart the OTHER_CONTROLLER by pressing the RESET button on the OTHER_CONTROLLER.

The OTHER_CONTROLLER will sound an alarm as it discovers the second controller but detects that it is not bound in a failover mode. The alarm (may be silenced) and the message about the controllers being misconfigured may be disregarded.

4. To verify the change in controller state, enter:

```
HSG80> SHOW THIS_CONTROLLER
```

Example:

```
Controller:
HSG80 ZG83502145 Software V85F-0, Hardware E03
NODE_ID          = 5000-1FE1-0000-3350
ALLOCATION_CLASS  = 0
SCSI_VERSION     = SCSI-2
-> Not Configured for dual-redundancy
-> Controller misconfigured -- other controller present
```

B. Changing from Transparent Failover to Multi-bus Failover Mode

Whether there are defined UNITS or not for the RAID system, the following steps will implement Transparent Failover to Multi-bus Failover:

1. At the HSG80> prompt, enter:

```
HSG80> SET NOFAILOVER
```

This action will cause the OTHER_CONTROLLER to shutdown.

2. At the HSG80> prompt, enter:

```
HSG80> SHOW THIS_CONTROLLER
```

From the resulting display, verify that the controller state has changed to no failover.

Example:

```
Controller:
HSG80 ZG83502145 Software V85F-0, Hardware E03
NODE_ID           = 5000-1FE1-0000-3350
ALLOCATION_CLASS   = 0
SCSI_VERSION      = SCSI-2
-> Not Configured for dual-redundancy
```

3. Restart the OTHER_CONTROLLER by pressing the RESET button on the OTHER_CONTROLLER.

The OTHER_CONTROLLER will sound an alarm as it discovers the second controller but knows that it is not bound in a failover mode. The alarm (may be silenced) and the message about the controllers being misconfigured may be disregarded.

Example:

```
Controller:
HSG80 ZG83502145 Software V85F-0, Hardware E03
NODE_ID           = 5000-1FE1-0000-3350
ALLOCATION_CLASS   = 0
SCSI_VERSION      = SCSI-2
-> Not Configured for dual-redundancy
Controller misconfigured -- other controller present
```

4. When the OTHER_CONTROLLER is on-line, enter the following command to put the controllers into Multi-bus Failover mode:

```
HSG80> SET MULTIBUS_FAILOVER COPY=THIS
```

This action will copy all unit and connection information to the OTHER_CONTROLLER and restart both controllers.

After both controllers have restarted, the controller pair will be bound in Multi-bus Failover mode with consistent views of all the RAID system information.

5. Verify that the controllers are now in Multi-bus failover. Enter:

```
HSG80> SHOW THIS_CONTROLLER
```

Example:

```
Controller:
HSG80 ZG83502145 Software V85F-0, Hardware E03
NODE_ID           = 5000-1FE1-0000-3350
ALLOCATION_CLASS   = 0
SCSI_VERSION      = SCSI-2
-> Configured for MULTIBUS_FAILOVER with ZG80200290
-> In dual-redundant configuration
```


6. If the RAID system had connections prior to making this transition, display the connections examine their “offset value” (last column), using this command:

HSG80> **SHOW CONNECTIONS**

Note: In Transparent Failover mode, the controller, by default, assigns an offset value of 0 to the left hand port and an offset value of 100 to the right hand port. In Multi-bus Failover mode, the controller assigns an offset value of 0 to all ports, unless existing connections have non-zero offset values.

Secure Path Note: Secure Path software requires connections to be reset to zero (0) so that both paths will be able to access the entire set of UNITS on the RAID system. To reset the connections to zero, enter:

HSG80> **SET CONNECTION *connection_name* UNIT_OFFSET=0**

Repeat for each connection.

C. Changing from Multibus Failover Mode to Nofailover and then to Transparent Failover Mode

1. Inspect the “Connections:” display for connections on the storage system. Enter:

```
HSG80> SHOW CONNECTIONS
```

Delete all connections by entering the following command for each connection found:

```
HSG80> DELETE connection_name (Repeat for each connection).
```

Note: The connections will be regenerated later.

2. If there are units (Dn) on the storage system, they must be deleted. This is due to the inconsistencies incorporated in the volumes' WWID in different failover modes.

```
HSG80> SHOW UNITS
```

```
HSG80> DELETE Dn (for each Dn)
```

Note: The UNITS will be restored after the controller state is changed. It is advised that the Dn values and associated information as well as the storage set information be recorded for later use. The controller state change will not affect the data on the storage sets.

3. If the controllers are currently in a failover mode, enter:

```
HSG80> SET NOFAILOVER
```

This action will cause the OTHER_CONTROLLER to shutdown.

Verify the current state of the controller by entering:

```
HSG80> SHOW THIS_CONTROLLER
```

Example:

```
Controller:
HSG80 ZG83502145 Software V85F-0, Hardware E03
NODE ID           = 5000-1FE1-0000-3350
ALLOCATION CLASS  = 0
SCSI VERSION     = SCSI-2
-> Not Configured for dual-redundancy
```

4. Restart the OTHER_CONTROLLER by pressing the RESET button on the OTHER_CONTROLLER.

The OTHER_CONTROLLER will sound an alarm as it discovers the second controller but knows that it is not bound in a failover mode. The alarm (may be silenced) and the message about the controllers being misconfigured may be disregarded.

Verify the current state of the controller by entering:

```
HSG80> SHOW THIS_CONTROLLER
```

Example:

```
Controller:
HSG80 ZG83502145 Software V85F-0, Hardware E03
NODE ID           = 5000-1FE1-0000-3350
ALLOCATION CLASS  = 0
SCSI VERSION     = SCSI-2
-> Not Configured for dual-redundancy
-> Controller misconfigured -- other controller present
```

5. When the OTHER_CONTROLLER is available, enter:

```
HSG80> SET FAILOVER COPY=THIS
```

This action will copy all unit and configuration information to the OTHER_CONTROLLER restart it. When restarted, the controller pair will be bound in Transparent Failover mode.

6. Verify the controller state change by entering:

HSG80> **SHOW THIS_CONTROLLER**

Example:

```
Controller:
  HSG80 ZG83502145 Software V85F-0, Hardware E03
  NODE ID           = 5000-1FE1-0000-3350
  ALLOCATION_CLASS   = 0
  SCSI VERSION      = SCSI-2
  -> Configured for dual-redundancy with ZG80200290
  ->   In dual-redundant configuration
```

7. Restore the UNIT to storage set mapping that was recorded earlier. Enter:

HSG80> **ADD UNIT Dn storage_set_name**

Warning! Do not initialize the units. This action will destroy data on the storageset(s).

8. Restart both controllers to reacquire connections. Enter:

HSG80> **RESTART OTHER_CONTROLLER**

HSG80> **RESTART THIS_CONTROLLER**

An alternative method to reestablish the connections is to reboot the server(s).