Shared Storage Solutions Using Fibre Channel Hubs

COMPAQ STORAGEWORKS

Application Note

EK–SMA26–AN. A01

Visit Our Web Site for the Latest Information

At Compaq, we are continually making additions to our storage solution product line. Please check our web site for more information on our fibre channel product line as well as the latest drivers, technical tips, and documentation. Visit our local web site at:

http://www.compaq.com/storageworks/

Introduction

The *Compaq StorageWorks* Enterprise Storage Array 12000 Fibre Channel (ESA12000 FC) and RAID Array 8000 Fibre Channel (RA8000 FC) storage systems, using the HSG80 storage controllers, are designed to support shared storage configurations. Shared storage configurations accommodate multiple host environments that demand any combination of the following features: storage data path security, flexible storage system accessibility, and/or storage compatibility with one or more host server operating systems.

A shared storage configuration consists of multiple host servers connected to a storage system, through fibre channel arbitrated loops, using fibre channel hubs. The hardware can be configured in such a way that it offers No-Single-Point-Of-Failure (NSPOF) storage data path security. In this environment, using special software, an alternate data path from the storage system to the server is available should the hardware components of a primary data path fail. In contrast, a Non-NSPOF configuration enables shared storage, without requiring the alternate data path.

Whether the shared storage environment is configured for NSPOF or Non-NSPOF, it can be configured to support more than one host server operating system. This is called a **Heterogeneous** operating system configuration. A shared storage configuration in which the host servers run the same operating system is called a **Homogeneous** operating system configuration.

This application note describes the concept of shared storage, the function of a Fibre Channel Arbitrated Loop (FC-AL), and how HSG80 controllers support Heterogeneous and Homogeneous operating system configurations. It also provides information for configuring NSPOF and Non-NSPOF shared storage environments, including hardware and software requirements. The presentation of this material is as follows:

1 Shared Storage Concepts

In a multiple host environment, where each host server or cluster requires exclusive access to its own data, a separate data storage system for each host server or cluster is typically employed. In this scenario, each storage system must be individually configured and maintained, each has its own set of hardware components and each requires its own physical location.

Shared storage, however, provides a better alternative to the complex storage demands of a multiple host environment. It enables more effective management of storage for a lower total cost. Multiple host storage needs are consolidated into a few storage systems. There are fewer hardware components and fewer physical locations. Each host server or cluster maintains exclusive access to its own data.

2 Fibre Channel Arbitrated Loop (FC-AL)

A fibre channel arbitrated loop consists of two or more fibre channel devices connected serially in a loop. Information is routed around the loop, being repeated by the intermediate devices, until it reaches its destination device. A fibre channel hub provides loop reliability by allowing any device to be added or removed from the loop without disrupting the flow of information around the loop.

Figure 1 Logical View of FC-AL

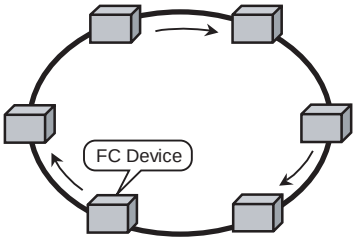

Figure 2 Physical View of FC-AL with an FC Hub

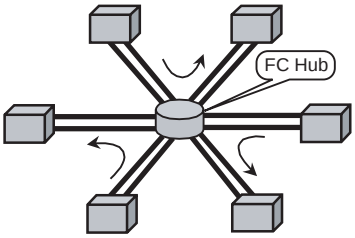

SHR-1502

3 Heterogeneous Operating System Configuration using HSG80 Technology

Heterogeneous configurations take advantage of the HSG80 controller's dual host ports. Host servers running one operating system are connected by an arbitrated loop to the first host port. Host servers, running a second operating system, are connected by a second arbitrated loop to the second host port. The host ports are isolated from each other at the hardware level. This isolation enables all operating systems supported using the HSG80 controller to be automatically supported in Heterogeneous shared storage configurations, two at a time (one operating system per loop). As of the writing of this document, the qualified operating systems are:

- Windows NT (Intel)
- Sun Solaris
- HP-UX
- SGI IRIX
- Novell NetWare
- SCO UnixWare (planned for Q3 1999)

The host servers can be standalone servers or a single cluster. A cluster cannot be mixed with standalone servers. However, the isolation of the HSG80 controller's host ports simultaneously allows the host servers on one host port to be standalone servers and the host servers on the other host port to be a cluster. Figure 3 and Figure 4 reveal how configuring separate arbitrated loops into isolated host ports accommodates multiple operating systems in a shared storage environment.

4 No-Single-Point-Of-Failure (NSPOF) Shared Storage Configurations

No-Single-Point-Of-Failure (NSPOF) configurations provide the maximum protection from data path failures. Host servers use two independent data paths (two loops) to access the storage system. A single data path consists of one fibre channel host bus adapter connected to one fibre channel hub connected to one storage controller. Should any part of a single data path fail, the host servers still have access to storage through the other data path. Up to two pairs of arbitrated loops are supported.

Figure 3 Multiple Host No-Single-Point-Of-Failure Shared Storage Configuration Layout

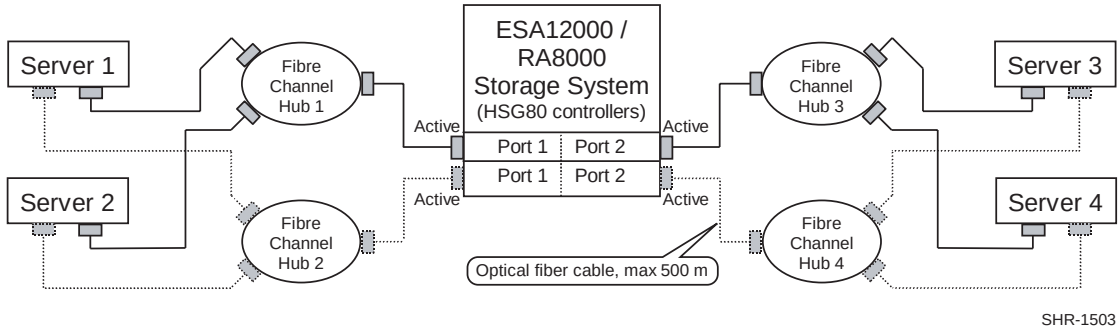

4.1 Requirements for NSPOF Shared Storage Configurations

This section describes the minimum requirements of a supported StorageWorks NSPOF shared storage configuration, which include:

- Storage Systems
- StorageWorks Command Console (SWCC)
- Host Server Operating System(s) (Heterogeneous/Homogeneous)
- Fibre Channel Hubs

4.1.1 Storage Systems (up to two)

The following Compaq StorageWorks Storage Systems are supported in a NSPOF shared storage configuration:

- Compaq StorageWorks Enterprise Storage Array 12000 Fibre Channel (ESA12000 FC)
- Compaq StorageWorks RAID Array 8000 Fibre Channel (RA8000 FC)

Controller Configuration

Storage systems must have dual-redundant HSG80 controllers with ACS V8.3 firmware. The controllers must be configured for multiple bus failover mode using the following HSG80 command line interface (CLI) command:

```
SET MULTIBUS_FAILOVER
```

Loop Configuration

One pair of loops (using Hubs 1 and 2, in Figure 3) is connected to host port 1 of both controllers. The other pair of loops (using Hubs 3 and 4) is connected to host port 2 of both controllers. One controller host port per loop.

NOTE

For more information regarding ESA12000 or RA8000 configuration requirements, please consult the *HSG80 Array Controller, ACS Version 8.3, Configuration and CLI Reference Guide* and other supporting documentation supplied with your storage system.

4.1.2 StorageWorks Command Console (SWCC)

The StorageWorks Command Console is recommended for setting up and managing these storage configurations. SWCC is a centralized graphical storage management console that allows real-time configuration of the storage environment and delivers reliable, real-time monitoring and notification of storage events.

NOTE

In a multiple host environment, each host may have an SWCC agent installed, but only **one** SWCC agent in the configuration may be running at a time. A single SWCC agent gives the Command Console client complete ability to monitor and control the RAID system. For more information about SWCC, please consult the *Command Console Version 2.1, User's Guide*.

4.1.3 Host Servers with Heterogeneous or Homogeneous Operating Systems

There are three main requirements for host servers supported in an NSPOF shared storage configuration, which are:

1. All host servers on one pair of loops must use the same operating system
   - Any operating system that is qualified with the HSG80 controller in an arbitrated loop and has multiple path support can be used. As of the writing of this document, the qualified operating systems are:

     Windows NT (Intel), HP-UX
   - The host servers on one pair of loops can use a different operating system from host servers on the other pair of loops
   - The operating system must provide multiple data access path support and the host servers must be configured to use it
   - Each server must have the *RA8000/ESA12000 FC Solution Software Kit* for its operating system installed
   - The number of host bus adapters within a single host server that can connect to the same loop is dependent on the operating system and the server platform
2. On one pair of loops, the host servers are all standalone or all part of a single cluster; cluster servers and standalone servers cannot be mixed on the same pair of loops
3. Maximum of two host server connections (host bus adapters) per loop recommended, up to four host server connections per loop supported

NOTE

For more information regarding host server configuration requirements, including operating system version, please consult the *RA8000/ESA12000 HSG80 Solution Software V8.3, Installation Reference Guide* for the host server's operating system and other supporting documentation supplied with your storage system.

4.1.4 Fibre Channel Hubs

Four hubs (two per operating system) and up to twenty Gigabit Interface Converters (GBIC) and fibre cables are required for a NSPOF shared storage configuration. The equipment supported is as follows:

- Compaq Fibre Channel Storage Hub 12 (12 ports), part # 295573-B22
- Compaq Fibre Channel Storage Hub 7 (7 ports), part # 234453-001
- Optical Short Wave GBIC, part # 380561-B21
- Fibre Channel Optical Cable, 50 Micron Short Wave Multi-mode, part # 234457-B21/B22/B23/B24/B25 (2/5/15/30/50m)

5 Non-NSPOF Shared Storage Configurations

Non-NSPOF configurations provide a lower cost solution where full protection from data path failures is not required.

Figure 4 Multiple Host Non-NSPOF Shared Storage Configuration Layout

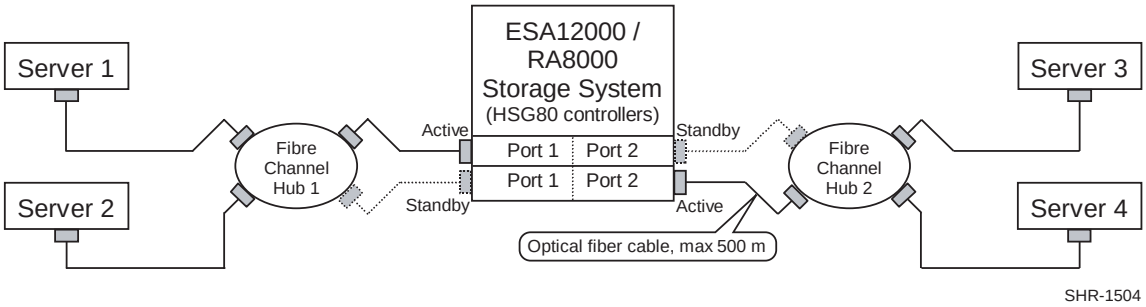

5.1 Requirements for Non-NSPOF Configurations

This section describes the minimum requirements of a supported StorageWorks Non-NSPOF shared storage configuration, which include:

- Storage Systems
- StorageWorks Command Console (SWCC)
- Host Server Operating Systems (Heterogeneous/Homogeneous)
- Fibre Channel Hubs

5.1.1 Storage Systems (up to two)

The following Compaq StorageWorks Storage Systems are supported in a Non-NSPOF shared storage configuration:

- Compaq StorageWorks Enterprise Storage Array 12000 Fibre Channel (ESA12000 FC)
- Compaq StorageWorks RAID Array 8000 Fibre Channel (RA8000 FC)

Controller Configuration

Storage systems must have HSG80 controllers with ACS V8.3 firmware. A single controller or dual redundant controllers may be used. If dual-redundant controllers are used, they must be configured for transparent failover mode using the following HSG80 command line interface (CLI) command:

```
SET FAILOVER
```

Dual-redundant controllers provide high availability and better performance than a single controller does. In the unlikely event that one controller should fail, the load moves immediately to the other controller transparently; the host servers see no interruption in service. A single controller must handle the load from both loops. Dual-redundant controllers split the load by each being dedicated to one loop.

Loop Configuration

One loop is connected to host port 1 of both controllers. The other loop is connected to host port 2 of both controllers.

NOTE

For more information regarding ESA12000 or RA8000 configuration requirements, please consult the *HSG80 Array Controller, ACS Version 8.3, Configuration and CLI Reference Guide* and other supporting documentation supplied with your storage system.

5.1.2 StorageWorks Command Console (SWCC)

The StorageWorks Command Console is recommended for setting up and managing these storage configurations. SWCC is a centralized graphical storage management console that allows real-time configuration of the storage environment and delivers reliable, real-time monitoring and notification of storage events.

NOTE

In a multiple host environment, each host may have an SWCC agent installed, but only **one** SWCC agent in the configuration is to be running. A single SWCC agent gives the Command Console client complete ability to monitor and control the RAID system. For more information about SWCC, please consult the *Command Console Version 2.1, User's Guide*.

5.1.3 Host Servers with Heterogeneous or Homogeneous Operating Systems

There are three main requirements for host servers supported in a Non-NSPOF shared storage configuration, which are:

1. All host servers on one loop must use the same operating system
   - Any operating system that is qualified with the HSG80 controller in an arbitrated loop can be used. As of the writing of this document, the qualified operating systems are:

     Windows NT (Intel), Sun Solaris, HP-UX, SGI IRIX, Novell NetWare,
     (SCO UnixWare planned for Q3 1999)
   - The host servers on one loop can use a different operating system from host servers on the other loop
   - Each server must have the *RA8000/ESA12000 FC Solution Software Kit* for its operating system installed
   - The number of host bus adapters within a single host server that can connect to the same loop is dependent on the operating system and the server platform
2. On one loop, the host servers are all standalone or all part of a single cluster; cluster servers and standalone servers cannot be mixed on the same loop
3. Maximum of two host server connections (host bus adapters) per loop recommended, up to four host server connections per loop supported

NOTE

For more information regarding host server configuration requirements, including operating system version, please consult the *RA8000/ESA12000 HSG80 Solution Software V8.3, Installation Reference Guide* for the host server's operating system and other supporting documentation supplied with your storage system.

5.1.4 Fibre Channel Hubs

Two hubs (one per operating system) and up to twelve Gigabit Interface Converters (GBIC) and fibre channel optical cables are required for a Non-NSPOF shared storage configuration. The equipment supported is as follows:

- Compaq Fibre Channel Storage Hub 12 (12 ports), part # 295573-B22
- Compaq Fibre Channel Storage Hub 7 (7 ports), part # 234453-001
- Optical Short Wave GBIC, part # 380561-B21
- Fibre Channel Optical Cable, 50 Micron Short Wave Multi-mode, part # 234457-B21/B22/B23/B24/B25 (2/5/15/30/50m)

6 NSPOF and Non-NSPOF Shared Storage Configuration Notes

6.1 Selective Storage Presentation

Server access control for NSPOF and Non-NSPOF shared storage configurations is provided by the HSG80 controller's selective storage presentation feature. Each storage unit in an ESA12000FC or RA8000FC can be selectively presented to specific host servers. This provides each host server or set of host servers with exclusive access to their own data. Selective storage presentation is accomplished by specifying for each storage unit a list of host connections allowed access and/or by using unit offsets.

6.1.1 List of Host Connections Allowed Access

A host connection is created when an HSG80 controller finds a host bus adapter connected to the fibre channel loop. Each HSG80 controller assigns a host connection name (e.g., !NEWCON01) to each host bus adapter. These host connection names can be changed to more meaningful names; e.g., SERVER1.

Each storage unit has a list of host connections that are allowed access to it. By default, a storage unit allows access by all host connections. Selective storage presentation occurs by listing a subset of host connections allowed access.

For example, in a configuration there is a storage unit, D24, and three host connections, SERVER1, SERVER2 and SERVER3. If SERVER1 is the only host connection name on the D24's list of host connections allowed access, then D24 is presented only to SERVER1. If SERVER2 is added to D24's list, then D24 is presented to both SERVER1 and SERVER2, but not to SERVER3.

NOTES

Use care when implementing a storage configuration that does not match either of the configuration layouts described earlier. If a storage system has an active host port 1 and an active host port 2 connected to the same loop, each host port will create its own host connection for the same host bus adapter. In other words, each host bus adapter will have associated with it two different host connection names; one assigned by host port 1 and the other by host port 2. Although this is not a problem, it does require careful management of the selective storage presentation.

The assignment of two host connection names does not occur when the NSPOF and non-NSPOF shared storage configurations are used.

6.1.2 Unit Offsets

The Compaq StorageWorks fibre channel storage systems use the SCSI protocol. The arbitrated loop physical address (ALPA) determines the SCSI ID for each node on an arbitrated loop. Each HSG80 controller has one ALPA which determines its target SCSI ID.

A host server uses this target ALPA to address the controller and uses the Logical Unit Number (LUN) to identify the storage unit presented by the HSG80 controller. For example, a host server addresses a storage unit using ALPA 72, LUN 4. The ALPA 72 identifies which HSG80 controller is the target of the SCSI message and LUN 4 identifies the storage unit presented by

that HSG80 controller. The range of LUNs a host server can address is determined by the capability of the host bus adapter it uses.

Unit offsets are a base for addressing storage unit names. The HSG80 controller uses unit offsets to allow multiple host servers to be able to address different storage units using the same ALPA and LUN combination. This allows the host servers to have exclusive access to storage and be able to use the full range of LUNs their host bus adapters can address.

Each host connection is assigned a unit offset. By default, the HSG80 controller assigns a unit offset of 0 (zero) to host connections on host port 1 and 100 to host connections on host port 2.

Adding the host connection's unit offset to the LUN being addressed gives the name of the storage unit that is accessed. For example, a host connection, SERVER1, has a unit offset of 100 and it addresses a storage unit by using LUN 3. Adding SERVER1's unit offset, 100, to LUN 3 gives the storage unit name D103; D103 is seen by SERVER1 as LUN 3.

By assigning different unit offsets to each host connection, they all have a different LUN for the same storage unit. A storage unit is addressable through the host connection that has the correct unit offset. The host connection with the correct unit offset will be presented a LUN that can be addressed by the host adapter. All other host connections will be presented a LUN outside of the range of LUNs that can addressed their host adapters.

For example, a host connection, SERVER1, and a second host connection, SERVER2, each represent a host bus adapter in different host servers. Both adapters can only address LUN 0 through LUN 7. SERVER1 is given a unit offset of 20 and SERVER2 is given a unit offset of 10. There exists a storage unit named D24. It is addressable by SERVER1, but not by SERVER2.

D24 is presented to SERVER1 as LUN 4:

D24 minus unit offset 20 equals LUN 4

D24 is presented to SERVER2 as LUN 14, outside the host adapter's addressable range:

D24 minus unit offset 10 equals LUN 14

For D24 to be addressable by both SERVER1 and SERVER2 as LUN4, both host connections must have a unit offset of 20. An example application would be where SERVER1 and SERVER2 are host connections to two host servers in a cluster.

The LUN addressing capability of the host bus adapter determines what storage units it can address. For example, if the host bus adapter is only capable of addressing LUN 0 through LUN 7 and its host connection has a unit offset of 50, then only storage units D50 through D57 are addressable.

Non-NSPOF configurations: An HSG80 controller configured for transparent failover mode has an additional selective storage presentation feature. Host port 1 only presents storage units D0 through D99 and host port 2 only presents storage units D100 through D199.

NSPOF configurations: For an HSG80 controller configured for multiple bus failover mode, both host port 1 and host port 2 present storage units D0 through D199.

NOTE

For more information regarding unit offsets and lists of host connections allowed access, please consult the *HSG80 ACS Version 8.3, Configuration and CLI Reference Guide* and other supporting documentation supplied with your storage system. Lists of host connections allowed access is found under the topic "Assigning Access Paths".

6.1.3 Configuring Server Access Control for a Cluster

To selectively present a storage unit to the host servers in a cluster one or both of the following can be used:

- The list of host connections allowed access contains only the host connection names for the host servers in the cluster
- The host connections for the cluster servers are given the same unique unit offset that puts the storage unit inside the range of LUNs addressable only by the cluster servers

6.1.4 Configuring Server Access Control for Standalone Servers

To selectively present a storage unit to a single host server one or both of the following can be used:

- The list of host connections allowed access contains only the host connection name for the standalone host server
- The host connection is given a unique unit offset that puts the storage unit inside the range of LUNs addressable only by the standalone host server

6.1.5 Heterogeneous Operating Systems Support

Heterogeneous operating systems sharing the same storage system is supported by the HSG80 controller's hardware isolation of its host ports. The hardware isolation prevents the loop connected to host port 1 from interacting with the loop connected to host port 2. By connecting host servers using one operating system to one host port and host servers using another operating system to the other host port, any possible unwanted interactions between operating systems are prevented.

Supported operating systems, as of the writing of this document:

Windows NT (Intel), Sun Solaris, HP-UX, SGI IRIX, Novell NetWare,
(SCO UnixWare planned for Q3 1999)

7 Reference Material

Table 1 lists the documents to reference for further information on the configuration of the RA8000 and ESA12000 Fibre Channel Storage Systems.

Table 1 Related Configuration Documentation

Topic	Document Title	Order Number
Windows NT (Intel)	*RA8000/ESA12000 FC-AL High Availability Configurations for Windows NT – Intel*	*AA-RH0SA-TE*
Windows NT (Intel)	*RA8000/ESA12000 FC-AL Configurations for Windows NT – Intel*	*AA-RH0RA-TE*
Windows NT (Intel)	*RA8000/ESA12000 HSG80 Solution Software V8.3 for Windows NT - Intel, Installation Reference Guide*	*387387-001 or AA-RFA9A-TE*
Sun Solaris	*RA8000/ESA12000 HSG80 Solution Software V8.3b for Sun Solaris 2.x, Installation Reference Guide*	*387384-003 or AA-RFBQC-TE*
HP-UX	*RA8000/ESA12000 HSG80 Solution Software V8.3 for HP-UX, Installation Reference Guide*	*387374-003 or AA-RFBEC-TE*
SGI IRIX	*RA8000/ESA12000 HSG80 Solution Software V8.3b for SGI IRIX, Installation Reference Guide*	*387399-002 or AA-RFBKB-TE*
Novell NetWare	*RA8000/ESA12000 Fibre Channel Solutions Software V8.3 for Novell NetWare, Installation Reference Guide*	*387376-001 or AA-RFB9A-TE*
HSG80 Controller	*HSG80 Array Controller, ACS Version 8.3, Configuration and CLI Reference Guide*	*387402-001 or EK–HSG80–RG*
StorageWorks Command Console	*Command Console V2.1 (HSG80) for Raid Array 8000/ESA12000, User's Guide*	*387405-001 or AA–RFA2A–TE*