

***RA8000/ESA12000***  
***FC-AL High Availability***  
***Configurations for***  
***Hewlett-Packard HP-UX***

**COMPAQ**  
STORAGeworks

**Application Note**  
EK-FCALM-AA. A01

---

**Introduction**

This application note serves as a guide to two common Fibre Channel Arbitrated Loop (FC-AL) configurations available for Compaq StorageWorks RAID Array 8000 (RA8000 FC) and Enterprise Storage Array 12000 (ESA12000 FC) storage system(s) in conjunction with the Hewlett-Packard family of servers running the HP-UX Operating System. It shows how to obtain RAID storage protection, ranging from RAID controller failover, to server failover, to complete independent path failover by configuring hardware and software, including MC/ServiceGuard and the Logical Volume Manager.

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

**Visit Our Web Site for the Latest Information**

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

[www.compaq.com/storageworks](http://www.compaq.com/storageworks)

**Enterprise Network Storage Architecture (ENSA)**

Compaq StorageWorks RAID Array 8000 (RA8000 FC) and Enterprise Storage Array 12000 (ESA12000 FC) products are the latest in fully integrated Fibre Channel RAID storage solutions. These products provide the key initial step in delivering Compaq's ENSA vision. ENSA addresses the issues that customers expect to face now and in the future. Both products address today's issues including: economical capacity growth, data protection, high availability, increased distance, multi-vendor platform support, and investment protection by being the base for building the ENSA vision.

## Product Overview

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

The RA8000 FC is a mid-range storage system available in a pedestal enclosure for departmental use in the office. It can also be installed in server racks for small data centers. An easy-to-deploy, flexible solution for open systems, the RA8000 supports up to 1.3 TB of storage.

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

These storage systems incorporate the latest in RAID technology providing RAID levels 0, 1, 0+1, adaptive 3/5, and non-RAID disks. Compaq's implementation of RAID capability assures that data availability will be maintained despite hardware failure. Features such as read ahead cache and mirrored write back cache improve performance while preserving data availability and supporting high availability.

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

## MC/Service Guard (MCSG)

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

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

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

## General Configuration Guidelines

Fibre Channel storage technology offers more configuration flexibility than parallel SCSI storage. This application note describes typical RA8000/ESA12000 FC-AL HP-UX High Availability configurations based on the configuration maximums listed. This document is not intended to describe all possible configurations. For configurations not shown here, the configuration maximums should be used to determine the viability of the proposed configuration. Exceeding the maximums listed will result in an unsupported configuration.

The maximums and notes listed apply to all configurations, additional information specific to particular configurations is provided in each configuration section.

## Configuration Maximums

- 500 meters per optical cable segment, 1 kilometer distance between server and storage controller
- Multibus\_Failover Mode supports up to 32 host-adapter connections
- Maximum storage capacity stated assumes 18GB disks, increases in disk drive capacity will increase overall storage capacity for the same enclosure footprint

## General Configuration Notes

- Dual redundant controllers can be configured for Failover Mode (referred to as Transparent Failover Mode) or Multibus\_Failover Mode.
- All configuration diagrams show a front view of the storage system, Controller A is the top controller; Controller B is the bottom controller. Controller Host Port 1 is the left port; Host Port 2 is the right port
- Each storage system can be configured using SWCC through a TCP/IP interface on the host
- All configurations utilize ACS V8.3 (Array Controller Software) configured for FC-AL topology
- All configurations shown utilize short wavelength laser/50 micron multi-mode fibre channel optical cables in all cable segments
- All configurations require that the *Arbitrated Loop Physical Address* (ALPA), for each loop device, be unique (refer to the *RA8000/ESA12000 FC Solution Software V8.3 for HP-UX - Installation Reference Guide*, Appendix D, for a list of valid addresses), and is set to *LOOP\_HARD*.
- Set the *Connection Name Operating System* parameter to “WINNT”, as HP mode is not supported in V8.3.
- To support Multibus\_Failover, PV Links **must** be configured for full redundancy in the manner described in the *Configuring PV Links using the mkvglinks Utility* paragraph contained in this document.
- Offsets  
For full redundancy, in Multibus\_Failover Mode, Offsets should *not* be used.
- For two enclosure (4 controller) configurations with more than one controller host port on the same loop, set the *ALPA* for each controller host port in the second enclosure to a different hard address (using the CLI, “SET controller PORT\_1\_AL\_PA=73”, “SET controller PORT\_2\_AL\_PA=74”)
- Record the worldwide ID (WWID) for the FC Host Bus Adapters by using the following command:  

```
# /opt/fc/bin/fcmsutil /dev/fcmsn
```

 (where *n* is the adapter number)
- Use the recorded WWID to identify the connection name displayed from the CLI at the controller using the “SHOW CONNECTIONS” command
- Rename connection names. By default new connection names are automatically added to the controller connection name table by the controllers when they detect a path to an adapter from each active controller host port. The default connection name assigned by the controllers will be “!NEWCONnn”, where nn is a number from 1 to 32. After the controllers detect all paths, rename each connection name to be more meaningful to the configuration i.e., SERVER1P1, SERVER1P2, etc. (connection names can be a maximum of 9 characters) Each active controller host port on a loop presents one SCSI Target ID with up to eight LUNS to each server (FC HBA) on the same loop. LUNs (logical units) can consist of single disks (JBOD), or a storageset of multiple disks or multiple partitions configured for a specific RAID level of 0, 0 + 1, or 3/5
- Partitioned storage sets are not permissible with controllers set to Multibus\_Failover Mode

## Configuration Descriptions

Table 2 lists the most common RA8000 / ESA12000 FC HP-UX high availability configurations and the pages of this document where a description of capacities, configuration notes, performance considerations, are provided.

**Table 2 Common RA8000/ESA12000 FC-AL High Availability Configurations for HP-UX - Index**

High Availability Configuration #	Description	Pages
1	One Server, PV Links, Dual Controller Multibus_Failover Mode	5 - 6
2	Two Servers, HP-UX, MC/ServiceGuard, Dual Controller Multibus_Failover Mode	7 - 8

## Configuration Parts Lists

A list of parts for each specific configuration is provided at the end of this document following the configuration descriptions.


## Additional Information/References

Table 3, at the end of this document, lists documentation to reference for further details on product installation and configuration.

## HA Configuration 1 – One Server, PV Links, Dual Controller Multibus\_Failover Mode

HA Configuration 1 (Figure 1) represents a basic 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 Multibus\_Failover mode. The configuration utilizes four paths (utilizing two separate fibre channel loops) to the storage. In the event of a failure of any component within the access path – HBA, cables, hub, or controller host port interface, all LUNs on the failing path will be accessible through an alternate path. The initial configuration provides up to 24 disks, an additional 48 disks can be added using two expansion enclosures. Additional servers can be added to this configuration as required using the existing hubs. HP-UX supports 8 LUNs per target. Each path is considered a target.

**Figure 1 HA Configuration 1**



SHR-1444

### HA Configuration 1 Notes

- Two Fibre Channel Loops
- Uses HP-UX's PV Links.
- Dual HSG80 controllers configured in Multibus\_Failover Mode
- Up to 500 meters per cable segment
- Up to 72 disks with 2 optional expansion cabinets
- Up to 8 LUNs with 4 active controller host ports. Note D0 must be the Command Console LUN.
- Set the Unit Preferred Path. Use the Preferred Path unit parameter to assign units to the specific controller that is connected to the primary path. Primary path should be distributed equally across controllers/ports for load balancing:

Path 1, FC HBA1, D1 – D3 (Controller A Port 1)

Path 2, FC HBA2, D4 – D7 (Controller B Port 1)

*For Example:* HSG80> **SET D1 PREFERRED\_PATH = OTHER\_CONTROLLER**

- In this configuration 4 logical connections are available to each LUN
- Rename the connection names, i.e., SERV1P1A, SERV1P1B, SERV2P1A, SERV2P1B

### HA Configuration 1 Parts List

- One HP9000 Server with HP-UX 10.20 or 11.00

Compaq Part #	Description	Quantity										
	RA8000/ESA12000 FC Solution Software V8.3A for HP-UX	1										
380577-001	ACS V8.3 FC-AL Controller SW	2										
	<table border="0"> <tr> <td><u>Machine Type</u></td> <td><u>HP Adapter PN</u></td> </tr> <tr> <td>K-Class:</td> <td>A3404A</td> </tr> <tr> <td>D-Class:</td> <td>A3491A or A3491B</td> </tr> <tr> <td>T-600 Class:</td> <td>A3636A</td> </tr> <tr> <td>V-Class:</td> <td>A3740A</td> </tr> </table>	<u>Machine Type</u>	<u>HP Adapter PN</u>	K-Class:	A3404A	D-Class:	A3491A or A3491B	T-600 Class:	A3636A	V-Class:	A3740A	2
<u>Machine Type</u>	<u>HP Adapter PN</u>											
K-Class:	A3404A											
D-Class:	A3491A or A3491B											
T-600 Class:	A3636A											
V-Class:	A3740A											
380550-B21 (Blue) 380550-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)	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	Choose 1										
242795-B21 295573-B22	FC Optical Hub, 7-Port (no GBICs) FC Optical Hub, 12-Port (no GBICs)	Choose 2										
380561-B21	FC Optical GBIC	6										
234457-B21 234457-B22 234457-B23 234457-B24 234457-B25	FC 2 Meter Optical Cable FC 5 Meter Optical Cable FC 15 Meter Optical Cable FC 30 Meter Optical Cable FC 50 Meter Optical Cable	Choose 6										
380691-B21 380595-B21 380588-B21 380694-B21 380589-B21	4GB UW 10000 RPM Disk 9GB UW 7200 RPM Disk 9GB UW 10000 RPM Disk 18GB UW 7200 RPM Disk 18GB UW 10000 RPM Disk	Choose up to 24 (48 for ESA 48 slot)										


### HA Configuration 1 Optional Components

Compaq Part #	Description	Quantity
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	Up to 2 OR Choose 1
See above	Additional Disks Drives	Up to 72
380563-B21	Expansion Cable Kit	1 per expansion cabinet

## HA Configuration 2 – Two Servers, HP-UX, MC/ServiceGuard, Dual Controller Multibus\_Failover Mode

HA Configuration 2 (Figure 2) combines the benefits of MCSG clustering, HP-UX's PV Links, and controller redundancy to provide the highest level of availability. The configuration utilizes four physical paths (utilizing two independent fibre channel paths) to the storage. In the event of a failure of any component within the access path – FC HBA, cables, hub, or controller host port interface, all units on the failing path will be accessible through an alternate path. In the event of a server failure, cluster failover will be invoked to provide access to units through the surviving server.

**Figure 2 HA Configuration 2**



### HA Configuration 2 Notes

- Two Fibre Channel Loops
- Uses HP-UX PV Links and MC/Service Guard
- Dual HSG80 controllers configured in Multibus\_Failover Mode
- Up to 500 meters per cable segment
- Up to 72 disks with 2 optional expansion cabinets
- Up to 8 LUNs with 4 active controller host ports. Suggested controller unit number assignment: Port 1 units D0 – D7. Note D0 needs to be the Command Console LUN.
- Set the Unit Preferred Path. Use the Preferred Path unit parameter to assign units to the specific controller that is connected to the primary path. Primary path should be distributed equally across controllers/ports for load balancing:

Path 1, FC HBA1, D1 – D3 (Controller A Port 1)

Path 2, FC HBA2, D4 – D7 (Controller B Port 1)

*For Example:* HSG80> **SET D1 PREFERRED\_PATH = OTHER\_CONTROLLER**

- In this configuration 4 logical connections are available to each LUN
- Rename the connection names, i.e., SERVAP1A, SERVAP1B, SERVBP1A, SERVBP1B, SERVAP2A, SERVAP2B, SERVBP2A, SERVBP2B

## HA Configuration 2 Parts List

- Two HP9000 Servers with HP-UX 10.20 or 11.00 and MC/Service Guard 10.10 or 11.00

Compaq Part #	Description	Quantity
	RA8000/ESA12000 FC Solution Software V8.3A for HP-UX	1
380577-001	ACS V8.3 FC-AL Controller SW	2
	Machine Type HP Adapter PN K-Class: A3404A D-Class: A3491A or A3491B T-600 Class: A3636A V-Class: A3740A	4
380550-B21 (Blue) 380550-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)	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	Choose 1
242795-B21 295573-B22	FC Optical Hub, 7-Port (no GBICs) FC Optical Hub, 12-Port (no GBICs)	Choose 2
380561-B21	FC Optical GBIC	8
234457-B21 234457-B22 234457-B23 234457-B24 234457-B25	FC 2 Meter Optical Cable FC 5 Meter Optical Cable FC 15 Meter Optical Cable FC 30 Meter Optical Cable FC 50 Meter Optical Cable	Choose 8
380691-B21 380595-B21 380588-B21 380694-B21 380589-B21	4GB UW 7200 RPM Disk 9GB UW 7200 RPM Disk 9GB UW 10000 RPM Disk 18GB UW 7200 RPM Disk 18GB UW 10000 RPM Disk	Choose up to 24 (48 for ESA 48 slot)

## HA Configuration 2 Optional Components

Compaq Part #	Description	Quantity
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	Up to 2 OR Choose 1
See above	Additional Disks Drives	Up to 48
380563-B21	Expansion Cable Kit	1 per expansion cabinet



## The “mkvlinks” Utility

The “mkvlinks” utility will create volume groups with physical volume on the RA8000 / ESA12000 Storage Subsystem and link all the paths to that volume. It will also link all the paths to the physical volume on existing volume groups that use the RA8000 / ESA12000 subsystem. Logical volumes must then be created according to user requirements.

### Installing the “mkvlinks” Utility

- Copy mkvlinks\_v1.tar to /tmp
- Enter: # tar xf /tmp/mkvlinks\_v1.tar

A copy of this document (AA-FCALM-TE.PDF) and the utility will be put in directory /opt/steam/bin.

### Configuring PV Links using the “mkvlinks” Utility

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

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

Set the preferred path using the following HSG80 Command Line Interface command:

```
HSG80> SET unit-number PREFERRED_PATH = THIS_CONTROLLER (or  
OTHER_CONTROLLER)
```

### Adding Links to the Physical Volume

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

#### NOTE

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

### Using the “mkvlinks” Utility on Existing Volume Groups

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

```
# /opt/steam/bin/mkvlinks <volume-group> [/dev/dsk/<physical-volume>...]
```

If the physical volume already has volume group information on it, the utility will ask whether to overwrite the volume. Enter “yes” to continue adding the volume.

### Using the mkvlinks Utility to Create a New Volume Group

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

```
# /opt/steam/bin/mkvlinks <volume-group> /dev/dsk/<physical-volume> ...]
```

If the physical volume already has volume group information on it, the utility will ask whether to overwrite the volume. Enter “yes” to continue adding the volume.

## Logical Volumes

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

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

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

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

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

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

For shared volume groups between hosts, copy the map file to the other host and perform steps 2 through 6, above.

Note that when using LVM maintenance commands that reference devices directly, (i.e., `pvmove`), you must first verify that the path is active using the following command:

```
# /opt/steam/bin/ctrlinfo . <device_special_file>
```

Where: `device_special_file` does not include the path.

If the message “`0 unk unk unk 0 unk`” is returned, then the path is not active, and you must use a different path. To produce a list of alternate paths, type:

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

## Configuration Reference Material

Table 3 lists the documents to reference for further information on the configuration of RA8000 and ESA12000 Fibre Channel Storage Systems.

**Table 3 Related Configuration Documentation**

<b>Topic</b>	<b>Document Title</b>	<b>Order Number</b>
<i>HSG80 Array Controller</i>	<i>Release Notes, HSG80 Array Controller (ACS V8.3) for HP-UX</i>	<i>AA-RFBDC-TE</i>
<i>Quick Setup Guide</i>	<i>RAID Array 8000/ESA12000 Fibre Channel Storage Subsystem for HP-UX</i>	<i>AA-RFBCA-TE</i>
<i>RA8000/ESA12000 Solution Software Installation</i>	<i>RA8000/ESA12000 HSG80 Solution Software V8.3 for HP-UX Installation Reference Guide</i>	<i>AA-RFBEB-TE</i>
<i>HSG80 Array Controller Configuration</i>	<i>HSG80 Array Controller ACS Version 8.3 Configuration and CLI Reference Guide</i>	<i>EK-HSG80-RG</i>
<i>HSG80 Array Controller Maintenance</i>	<i>HSG80 Array Controller ACS Version 8.3 Maintenance and Service Guide</i>	<i>EK-HSG80-SV</i>
<i>StorageWorks Command Console</i>	<i>Command Console V2.1 (HSG80) for RAID Array 8000/ESA12000, User's Guide</i>	<i>AA-RFA2A-TE</i>
<i>Warranty Information</i>	<i>DIGITAL StorageWorks Warranty Terms and Conditions</i>	<i>EK-HSXS-WC</i>
<i>HSG80 Array Controller Parts</i>	<i>HSG80 Array Controller Illustrated Parts Map</i>	<i>EK-HSG80-MP or 355217-001</i>
<i>RA8000/ESA12000 Configuration</i>	<i>RA8000/ESA12000 Fibre Channel HSG80 Storage Subsystem Configuration Poster</i>	<i>EK-HSG80-CP or 355207-001</i>

HP-UX is a registered trademark of Hewlett-Packard Corporation.  
 Intel is a registered trademark of Intel Corporation.  
 Windows NT and MSCS are trademarks of Microsoft Corporation.  
 STORAGEWORKS and the Compaq Logo are trademarks of Compaq Computer Corporation.  
 All other trademarks and registered trademarks are the property of their respective owners.