

Storage Area Network Configurations For RA8000/ESA12000 on OpenVMS

COMPAQ
STORAGeworks

Application Note
EK-SMA34-AN. A01

Visit Our Websites for the Latest Information

At Compaq we are continually making additions to our storage solutions product line. Please check our websites 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 pages at:

www.compaq.com/storageworks

www.compaq.com/openvms

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 steps in delivering on Compaq's ENSA vision. ENSA addresses the issues that our 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.

With the addition of fibre channel switched fabric support in the RA8000 and ESA12000 FC storage products, Compaq has taken the next evolutionary step in delivery of the ENSA vision. Our customers can now realize the full benefits of a Storage Area Network (SAN) providing exceptional connectivity while increasing the effective bandwidth of the network. Fibre Channel Fabric provides the robust backbone needed to address dispersed server operation with shared storage access across the enterprise.

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 support utilizing both FC-Fabric (switch) and FC-AL (arbitrated loop) technology. (OpenVMS requires that FC storage systems be configured for Fabric.)

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.

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 or minimize the affect on 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.

Performance Considerations

The performance of an application on a system that includes RA8000 or ESA12000 FC storage can be limited by many different components of the system. Some of the possible component limitations include the host CPU(s), memory size, the host bus adapter(s) (HBA), the fibre channel bus, the RAID controller, or the specific configuration of disks used behind the controller. The limiting factor can move to any of these components depending on the workload. Identifying the limits will assist in determining the best configuration for a given application.

Table 1 lists the upper limit performance specifications for the components of the RA8000 and ESA12000 FC storage systems based on testing using standard storage performance tests and methodologies. These numbers should be used to compare component level performance as a means to determine the best configuration from a performance perspective. User application tests may not necessarily reach these levels of performance as applications may perform additional levels of processing for each I/O. The controller specifications listed show both cache (no disk access) and media (with disk access) limitations.

The limits are based on I/O performance (I/O's per second) - typical of small transfer applications such as databases and mail, and bandwidth performance (MB's per second) - typical of large transfer applications such as video and graphics.

Table 1 Performance Limits of FC-Fabric Components

Transfer Size	IO/sec (small transfer sizes, random access)		MB/sec (large transfer sizes, sequential access)	
Operation Type	Read	Write	Read	Write
FC HBA - KGPSA-BC	15,500	14,250	86	51
FC Switch	400,000	400,000	1600	1600
Single HSG80 Controller (1 Active Port)	12,000 Cache 4,100 Media	9,000 Cache 4,000 Media	77 Cache 54 Media	50 Cache 44 Media
Single HSG80 Controller (2 Active Ports)	12,000 Cache 4,100 Media	9,000 Cache 4,000 Media	98 Cache 54 Media	100 Cache 47 Media
Dual HSG80 Controllers (2 Active/2 Standby Ports)	24,000 Cache 8,200 Media	18,000 Cache 8,000 Media	154 Cache 102 Media	101 Cache 88 Media
Dual HSG80 Controllers (4 Active Ports)	24,000 Cache 8,200 Media	18,000 Cache 8,000 Media	195 Cache 102 Media	178 Cache 88 Media

NOTE

FC Switch performance limits are theoretical. Actual measured performance is lower due to limits unrelated to the switch.

General Configuration Guidelines

One of the major benefits of fibre channel fabric technology is the ability to optimize and tailor configurations to best meet your specific needs today, while at the same time provide you the flexibility to expand your storage network as needed to address future needs. Consider future as well as present needs when determining the optimal initial configurations.

Configuration Maximums

This application note describes an RA8000/ESA12000 FC OpenVMS fabric configuration based on the configuration maximums for servers and storage systems. This document is not intended to describe all possible configurations, as the potential for fibre channel fabric configurations is too numerous to define or illustrate. The configuration maximums should be used as a guide to determine the viability of a proposed configuration. It is strongly recommended that you not exceed the maximums listed.

- The maximum total distance from server to storage is 11 kilometers in normal operation with a worst-case single fault reroute distance of 21 kilometers. This requires that long wave optical fibre cables and long wave GBICS be used for the Inter-Switch Links (ISLs).
- The maximum number of FC HBAs per server is 4 (or total allowed in the server), maximum of 2 HBAs per zone
- Fibre channel switches - Any combination of OpenVMS Servers and RA8000/ESA12000 storage systems is allowed per switch, provided the following rules are followed:
 1. A maximum of 8 OpenVMS servers per RA8000 or ESA12000. If there are more than 8 servers per fabric, then it may be necessary to use zoning to limit the number of servers connected to a storage subsystem to 8.
 2. A maximum of 8 RA8000 or ESA12000 (with 1 controller pair) can be configured on a

- FC 16-port switch (Note: If you configure all 4 HSG80 ports on each storage system the maximum is 3 RA8000 or ESA12000 storage systems per 16-port switch).
- 3. A maximum of 8 RA8000 or ESA12000 (with 1 controller pair) per FC HBA.
- 4. Cascading of switches is allowed. Refer to *Heterogeneous Storage Area Networks Application Note*, part number EK-SMA30-AN.
- Maximum storage capacity stated (2.6 TB) assumes 36 GB disks; increases in disk drive capacity will increase overall storage capacity for the same enclosure footprint.
- Intermixing of standalone, and multi-host cluster configurations is allowed, provided the rules listed above are followed.

General Configuration Notes

Refer to the following website location for OpenVMS FC support documentation for the minimum required revision of the operating system, firmware, and hardware:

<http://www.openvms.digital.com/openvms/fibre/index.html>

- Dual redundant controllers must be configured for Multiple-Bus Failover Mode
- Each storage system can be configured using SWCC or the Command Line Interface (CLI) through either the HSG80 controller maintenance serial port or a TCP/IP interface.
- All configurations utilize ACS V8.5F or ACS V8.5S (Array Controller Software) configured for FC-FABRIC topology. ACS V8.5F is for fabric operation, and V8.5S adds support for snapshot volumes.
- All configurations utilize short wavelength lasers and multi-mode fibre channel optical cables in each cable segment. Label both ends of each FC cable with the same unique identification marking (number, letter, etc.) prior to installation.
- All configurations require the *Connection Name Operating System* parameter set to “VMS”.
- By default the HSG80 Command Console LUN (CCL) is enabled.
- The CCL must be given a unique device identifier.
- Set the *Connection Name Unit Offset* value to 0 for OpenVMS.
- Use the Unit *Connection Name* parameter to allow exclusive access to units from each server (using the CLI, “SET unit-number DISABLE_ACCESS_PATH = ALL
ENABLE_ACCESS_PATH = *connection-name*”).
- OpenVMS requires the controllers to use SCSI-3 mode
- A device identifier must be set for each logical unit. This identifier must be unique throughout the OpenVMS cluster.
- Record the FC HBA *Worldwide ID* (WWID) IEEE address shown on the back of each FC HBA prior to installing into a server.
- 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 64. After the controllers detect all paths, rename each connection name to be more meaningful to the configuration i.e., SERVER1P1, SERVER1P2, etc. (connection names can be a maximum of 9 characters)
- Each active controller host port presents one SCSI Target ID with up to 128 LUNs to each server (FC HBA) on the same zone. LUNs (logical units) can consist of single disks (JBOD), a storageset of multiple disks, configured for a specific RAID level of 0, 1, 0 + 1, or 3/5

Please refer to the OpenVMS website <http://www.compaq.com/openvms/> for detailed information about the required operating system version and update kits, that are required to support SANs and Host Based Volume Shadowing. Fibre Channel information can be found by following the links to “Product Information” and “VMSClusters”.

Configuration Descriptions

Minimum Configuration

The minimum OpenVMS SAN configuration is shown in figure 1. One RA8000 or ESA12000 storage system may be connected to a single stand-alone host or a two node cluster using one FC switch.

Maximum Configuration

The RA8000 and ESA12000 FC-Fabric SAN maximum configurations for OpenVMS are shown in Figure 2, SAN Configuration. The configuration is based on the maximums listed in the general configuration guideline section. The configuration shows the limits with regard to the number of servers and storage systems per FC fabric.

SAN configuration 2 describes maximums with the RA8000 and ESA12000 HSG80 controllers configured in Multiple Bus Failover Mode. All configurations support a mix of both clustered and non-clustered servers. All configurations illustrated utilize 16-port FC switches. The 8-port FC switch is supported in the same configuration types with less connectivity due to the reduced number of ports available.

Configuration Parts List

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

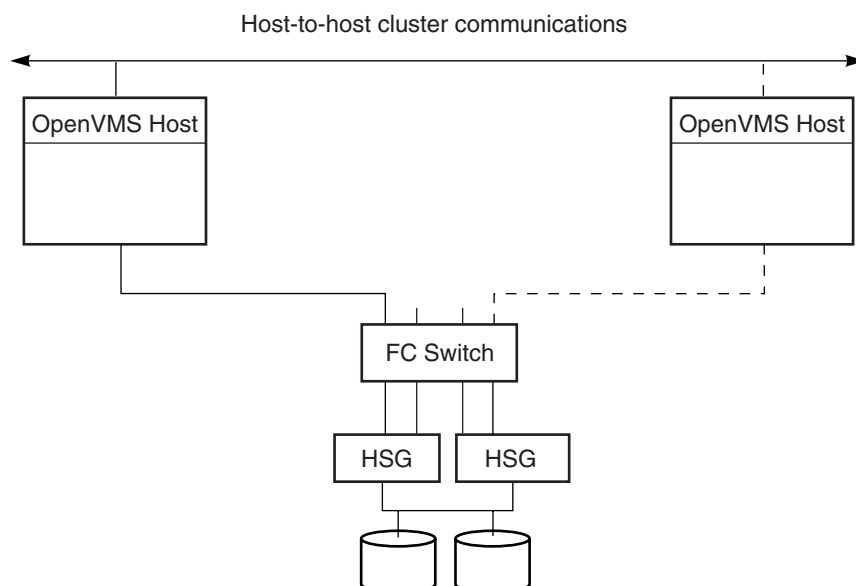
Additional Information/References

A table of available product documentation is included at the end of this application note for reference. Refer to the documents listed for more details on the product installation and configuration.

SAN Configuration 1 – Minimum SAN Configuration

The configuration below shows the minimum OpenVMS SAN. A single stand-alone host or a two node OpenVMS cluster can be connected through a single 8-port switch providing access from the host or hosts to the storage system. The HSG80 array controllers in the RA8000 or ESA12000 Storage System may be configured in Multiple-Bus Failover mode providing controller redundancy in the unlikely event of a controller failure.

Figure 1 SAN Configuration



SHR-1605

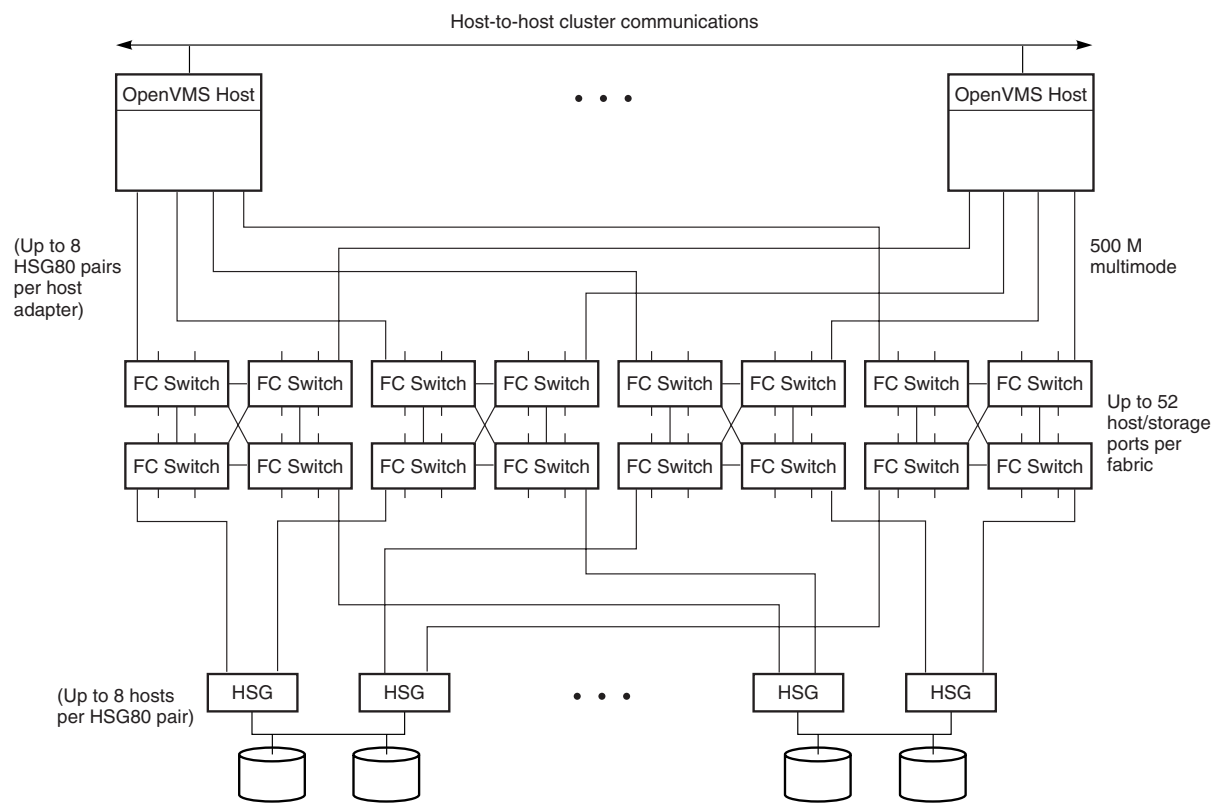
SAN Configuration 1 Notes

- One Fibre Channel Switch
- One host or a two-node cluster
- One storage system
- HSG80 controllers configured in Multiple-Bus Failover mode
- Rename the connection names, for example SERVER1P1, SERVER2P1, SERVER1P2 and so on.
- Unique device identifiers for the storage units and the CCL
- Controllers in SCSI-3 mode
- Each cluster and non-clustered host must have exclusive access to its storage devices.
- Set the connection unit offsets to 0.

SAN Configuration 2 - Multiple FC 16-Port Switches

The SAN Configuration below shows the maximum server connectivity currently supported. Up to eight servers can be connected through the switches providing access from all servers to the storage systems. This configuration utilizes the high connectivity and high performance capabilities of the FC 16-port switch. The HSG80 array controllers within the RA8000 or ESA12000 storage enclosure are configured in Multiple-Bus Failover Mode providing full controller redundancy in the unlikely event of either controller failing.

Figure 2 SAN Configuration



SHR-1606

SAN Configuration 2 Notes

- Multiple Fibre Channel Switches
- Up to 8 servers, clustered or non-clustered per HSG80 pair
- Up to 8 shared storage enclosures per host adapter
- HSG80 controllers configured in Multiple Bus Failover Mode
- Up to 500 meters per cable segment (ISLs up to 10 Km using long wave cables and GBICS)
- Up to 72 disks with 2 optional expansion enclosures per storage system
- Rename the connection names, i.e., SERVER1P1, SERVER2P1, SERVER1P2, etc.
- Unique device identifiers for the storage units and the CCL
- Controllers in SCSI-3 mode
- Each cluster and non-clustered server must have exclusive access to its storage devices.
- Set exclusive unit access for each unit to a specific server connection name, i.e., for unit D0 ENABLE = SERVER1P1, for unit D100 ENABLE = SERVER7P2. If you are using SWCC to create units, by default units are enabled on all known connection names. You must explicitly disable connection names enabled for units to prevent access.
- Set the connection unit offsets to 0.

Parts List

Compaq Part #	Description
380555-001	RA8000/ESA12000 FC Solution Software V8.5 for OpenVMS
TBS	ACS V8.5F Controller SW
380574-001	KGPSA-BC PCI FC HBA
380560-B21 (Blue) 380560-B22 (Opal) 380670-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-B21 (Blue) 380630-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 ESA12000 w/2 pairs/dual HSG80 48 Slot 60HZ ESA12000 w/2 pairs/dual HSG80 48 Slot 60HZ ESA12000 w/2 pairs/dual HSG80 48 Slot 50HZ ESA12000 w/2 pairs/dual HSG80 48 Slot 50HZ
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
380578-B21 (Blue) 380578-B22 (Opal) 380591-B21 (Blue) 380591-B22 (Opal) 158223-B21 158222-B21	FC Switch 16 port (no GBICs) FC Switch 16 port (no GBICs) FC Switch 8 port (no GBICs) FC Switch 8 port (no GBICs) FC SAN Switch 16 port firmware V2.0.3A FC SAN Switch 8 port firmware V2.0.3A
380561-B21 127508-B21	FC Optical GBIC Optical Long Wave GBIC
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
380691-B21 380595-B21 380694-B21 380588-B21 380589-B21	4GB UW 7200 RPM Disk 9GB UW 7200 RPM Disk 18GB UW 7200 RPM Disk 9GB UW 10000 RPM Disk 18GB UW 10000 RPM Disk

Product Details

The RA8000/ESA12000 storage systems utilize the Compaq HSG80 RAID controller running Array Controller Software (ACS) V8.5. 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 such 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 huge 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 for their application. 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.

You can configure and monitor the RA8000 and ESA12000 FC using the StorageWorks Command Console (SWCC). SWCC is included in the HSG80 Software Solutions Platform kit. Using this client/server tool, you can configure RAID sets locally or over a TCP/IP network. The client provides an easy to use graphical user interface. It can be used to monitor your storage system and notify you of events by updating a graphical 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 host platform.

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

Configuration References

The following documents provide further information on the configuration of RA8000 and ESA12000 Fibre Channel Storage Systems:

Table 3 Configuration Reference Material

Topic	Document Title	Order Number
<i>OpenVMS HSG80 Array Controller</i>	<i>Release Notes, RA8000/ESA12000 HSG80 (ACS V8.5) for OpenVMS</i>	<i>AA-RH4AB-TE</i>
<i>OpenVMS RA8000/ESA12000 Installation</i>	<i>RA8000/ESA12000 HSG80 Solution Software V8.5 for OpenVMS Installation Reference Guide</i>	<i>AA-RH4BB-TE</i>
<i>HSG80 Array Controller CLI Commands</i>	<i>HSG80 Array Controller ACS Version 8.5 CLI Reference Guide</i>	<i>EK-HSG85-RG</i>
<i>HSG80 Array Controller Maintenance</i>	<i>HSG80 Array Controller ACS Version 8.5 Maintenance and Service Guide</i>	<i>EK-HSG84-SV</i>
<i>StorageWorks Command Console</i>	<i>Command Console V2.2 (HSG80) for RAID Array 8000/ESA12000, User's Guide</i>	<i>AA-RFA2D-TE</i>
<i>Warranty Information</i>	<i>DIGITAL StorageWorks Warranty Terms and Conditions</i>	<i>EK-HSXS-WC</i>
<i>Heterogeneous SAN Information</i>	<i>Heterogeneous Storage Area Networks Application Note</i>	<i>EK-SMA30-AN</i>
<i>HSG80 Array Controller Configuration</i>	<i>HSG80 Array Controller ACS Version 8.5 Configuration Guide</i>	<i>EK-HSG85-CG</i>