SANworks by Compaq

Tru64 UNIX Kit V2.0 for Enterprise Virtual Array Installation and Configuration Guide

Part Number: AA-RQZYB-TE

Second Edition (July 2002)

Product Version: 2.0

This guide describes how to install and configure Tru64 UNIX servers with an Enterprise Virtual Array.

COMPAQ

Contents

3 Testing Connections to the Enterprise Virtual Array

4 Additional Host Considerations

Glossary

Index

Figures

Tables

About this Guide

The following sections are covered:

- Intended Audience
- Related Documentation
- Text Conventions
- Symbols in Text
- Symbols on Equipment
- Rack Stability
- Getting Help
- Compaq Authorized Reseller

Intended Audience

This guide is intended for storage administrators who are experienced with the following:

- Configuration of SAN fabrics.
- Customer host environments, such as Windows 2000, Windows NT, Sun Solaris, OpenVMS, Tru64 UNIX, HP-UX, and IBM AIX.

Related Documentation

This section lists documents that provide additional information:

- *SANworks by Compaq Tru64 UNIX Kit V2.0 for Enterprise Virtual Array Release Notes*

- *StorageWorks by Compaq Enterprise Virtual Array Users Guide*

- *StorageWorks by Compaq Enterprise Virtual Array - How to Get More Information*

- *StorageWorks by Compaq Enterprise Virtual Array Release Notes*

- *StorageWorks by Compaq Enterprise Virtual Array Storage System Read Me First*

- *StorageWorks by Compaq Enterprise Virtual Array Storage System World Wide Name Label*

- *StorageWorks by Compaq Interactive Help for HSV Element Manager*

- *StorageWorks by Compaq Enterprise Virtual Array Configuration Guide*

- *StorageWorks by Compaq Enterprise Virtual Array Upgrade Instructions*

- *StorageWorks by Compaq Enterprise Virtual Array License Instructions*

- *StorageWorks by Compaq Enterprise Virtual Array Catalog of Associated Documentation*

- *StorageWorks by Compaq Enterprise Virtual Array Storage System Installation Instructions*

- *StorageWorks by Compaq Heterogeneous Open SAN Design Reference Guide*

Text Conventions

The conventions included in Table 1 apply in most cases.

Table 1: Text Conventions

Element	Convention	Examples
• **Named Keys**	**Bold**	**Home, Print Screen, Num Lock, Esc, PgUp**
• **Key Sequences**	A plus sign (**+**) between two keys means that you should press them simultaneously.	**Ctrl+A**, **Ctrl+Home**, **Alt+Ctrl+Del**
		To back up files, click **Backup Now**.

Table 1: Text Conventions (Continued)

Element	Convention	Examples
• **Menu Items and Sequences** • **Buttons** • **Dialog Box Names**	**Bold** A right angle bracket (**>**) between items refers to the sequence for navigating through menu selections.	Choose **Start > Programs > Backup.** Click **Make Backup**. In the **Save As** dialog box, choose the drive, then the folder.
• Directory Names and Paths	Initial Caps (for UNIX-based directory names, the exact case is displayed).	Save the file in the C:\StorageSets\Default directory. (UNIX-based): Save the file in the /home/newuser/practice directory.
• *file names*	*Italics* unless the file name is included in a directory name/path.	To configure storage, edit *storageset.ini*. (Directory name/path): Errors are logged to \syslog\errors\config_errors.txt.

• User Input and System Responses (Output and Error Messages)
 • COMMAND NAMES
 • Drive Names
	`Monospace font.`
	
	 `COMMAND NAMES` appear in uppercase, unless they are case sensitive (UNIX-based command names are case sensitive and will not appear in uppercase).
	
	 Entered `<variables>` are displayed in angle brackets (< >) and all lowercase.
		User Input and System Responses:
		 • To exit from the program, type `exit`.
		 • At the prompt, type this command:
		 `SHOW THIS_CONTROLLER`
		 (no variable).
		 • To see your settings, give the command:
		 `SHOW <storagesets> FULL`
		 (with variable).You will see the `Disk Full` message.
		 Command Names
		 • Use `SET THIS_CONTROLLER` to change parameters.
		 • To manage storage, `RUN SYSMGR.EXE`.
		 • (UNIX-based): To list files, use the `ls` command.
		 • Drive Names:
		 Navigate to your CD-ROM drive (usually `D:` or `E:`).

URLs	Sans serif font.	For update notices, visit: http://www.compaq.com/products/updates

Symbols in Text

These symbols may be found in the text of this guide. They have the following meanings.

WARNING: Text set off in this manner indicates that failure to follow directions in the warning could result in bodily harm or loss of life or damage to equipment.

CAUTION: Text set off in this manner indicates that failure to follow directions could result in damage to equipment or loss of information.

IMPORTANT: Text set off in this manner presents clarifying information or specific instructions.

NOTE: Text set off in this manner presents commentary, sidelights, or interesting points of information.

Symbols on Equipment

Any enclosed surface or area of the equipment marked with these symbols indicates the presence of electrical shock hazards. Enclosed area contains no operator serviceable parts.

WARNING: To reduce the risk of injury from electrical shock hazards, do not open this enclosure.

Any RJ-45 receptacle marked with these symbols indicates a network interface connection.

WARNING: To reduce the risk of electrical shock, fire, or damage to the equipment, do not plug telephone or telecommunications connectors into this receptacle.

Any surface or area of the equipment marked with these symbols indicates the presence of a hot surface or hot component. Contact with this surface could result in injury.

WARNING: To reduce the risk of injury from a hot component, allow the surface to cool before touching.

Power supplies or systems marked with these symbols indicate the presence of multiple sources of power.

WARNING: To reduce the risk of injury from electrical shock, remove all power cords to completely disconnect power from the supplies and systems.

Any product or assembly marked with these symbols indicates that the component exceeds the recommended weight for one individual to handle safely.

WARNING: To reduce the risk of personal injury or damage to the equipment, observe local occupational health and safety requirements and guidelines for manually handling material.

Rack Stability

⚠️ **WARNING: To reduce the risk of personal injury or damage to the equipment, be sure that:**

- **The leveling jacks are extended to the floor.**
- **The full weight of the rack rests on the leveling jacks.**
- **In single rack installations, the stabilizing feet are attached to the rack.**
- **In multiple rack installations, the racks are coupled together.**
- **Only one rack component is extended at any time. A rack may become unstable if more than one rack component is extended for any reason.**

Getting Help

If you still have a question after reading this guide, contact your service representatives or visit our website.

Compaq Technical Support

In North America, call the Compaq technical support at 1-800-OK-COMPAQ. This service is available 24 hours a day, 7 days a week.

NOTE: For continuous quality improvement, calls may be recorded or monitored.

Outside North America, call Compaq technical support at the nearest location. Telephone numbers for worldwide technical support are listed on the Compaq website: http://www.compaq.com.

Be sure to have the following information available before you call Compaq:

- Technical support registration number (if applicable)
- Product serial numbers
- Product model names and numbers
- Applicable error messages
- Operating system type and revision level
- Detailed, specific questions

Compaq Website

The Compaq website has the latest information on this product as well as the latest drivers. Access the Compaq website at: http://www.compaq.com/storage. From this website, select SANworks.

Compaq Authorized Reseller

For the name of your nearest Compaq Authorized Reseller:

- In the United States, call 1-800-345-1518.

- In Canada, call 1-800-263-5868.

- Elsewhere, see the Compaq website for locations and telephone numbers.

1

Understanding the Enterprise Virtual Array

This chapter introduces the Enterprise Virtual Array components and explains how the host servers are integrated. The following topics are covered:

- About Virtualization, page 1–2
- About the Enterprise Virtual Array, page 1–2
    - Storage Configuration, page 1–3
    - Hosts and the Enterprise Virtual Array, page 1–4
    - SAN Considerations, page 1–6
    - Multiple Path Support, page 1–6
    - OS Unit ID, page 1–6
    - Host Presentation Compared with Enterprise Modular Array, page 1–6
- Virtual RAID Considerations, page 1–7

About Virtualization

Virtualization is a new technology used to simplify the creation, presentation, and administration of storage to multivendor host servers in a Storage Area Network (SAN). Virtualization changes the way the storage administrator interacts with storage—streamlining the work required to manage and implement the storage environment. This section describes how virtualization affects storage configuration.

The storage administrator no longer needs to make decisions about planning, creating, and configuring stripesets, mirrorsets, and RAIDsets. The software now automates these decisions. The decisions are simplified to basic choices on virtual disk capacity and redundancy levels. All of this work is done from a central location—the HSV Element Manager. See the HSV Element Manager online help for more information.

Three levels of virtualization are possible within a SAN—server, fabric, and storage system. The Enterprise Virtual Array implements storage system level virtualization.

- **Server level**—useful for small systems—*SANworks*™ Virtual Replicator by Compaq implements small scale virtualization of storage in a Windows NT and 2000 environment

- **Fabric level**—SAN-wide virtualization with increased efficiency

- **Storage system level**—provides large volumes of pooled storage in virtual disks and simplifies management tasks

Virtualization technology, at the storage system level, creates "virtual disks." These virtual disks are created using all the available physical disk drives, not individual or grouped sets of disks. The host recognizes and uses these virtual disks like any other disk device.

About the Enterprise Virtual Array

Storage system level virtualization is a new concept in the storage industry that allows storage administrators to focus on higher-level concerns regarding their specific storage needs.

With the Enterprise Virtual Array, the storage administrator no longer needs to manually present storage pools to the host servers. That is, you do not choose specific disks and sets of disks to create levels of redundancy. No decisions need to be made as to which physical disks are involved in each storage unit. When the virtual disks are created, the entire set of disks in the cabinet are used for load balancing and sparing,

which sets aside extra disk space for failure protection. The Enterprise Virtual Array improves performance as the data is written across many disks and not directed towards a single or specific set of disks.

Setup and management of virtualization is achieved with new software and hardware resources. The storage administrator has greater freedom and control with the following benefits:

• Faster performance with improved system response time

• All SAN and storage management done from a web browser

• Simplified load balanced storage

• Simplified decisions about physical disk setup and partitioning

• Increased bandwidth—use of striping algorithms across many disks accessed with multiple spindles

• Simplified high-availability storage techniques

• Recovery from disk failures includes automatic load balancing

Storage Configuration

All complex choices previously made manually, regarding physical disk usage, are now performed by the software. The end result is overall better performance and less intervention from the storage administrator.

The HSV Element Manager software automatically makes complex choices about physical disk usage. Virtual disks are created behind the scenes with this software. The storage administrator only makes relatively simple choices regarding virtual disk capacity, which host to present to, and the level of redundancy.

The following list provides an overview of the storage configuration process for the Enterprise Virtual Array:

1. Initialize the storage system.

2. Create disk groups.

3. Add hosts to the storage system.

4. Create virtual disks.

5. Present virtual disks to hosts.

For more information on setting up the Enterprise Virtual Array, see the *Enterprise Virtual Array User Guide* and the HSV Element Manager online help system.

Hosts and the Enterprise Virtual Array

This section describes how the host servers fit in the overall Enterprise Virtual Array. Below is a list of the hosts that can attach and interact with the Enterprise Virtual Array:

- Windows NT and Windows 2000

- Compaq *Tru64 UNIX*

- Compaq *OpenVMS*

- Sun Solaris

- HP-UX

- IBM AIX

Be sure you are running a supported version of each operating system in your SAN (see the platform-specific release notes for details).

Figure 1–1 shows an overall SAN architecture. The control center for the Enterprise Virtual Array is the HSV Element Manager, which resides on the StorageWorks Management Appliance (SWMA) and is accessed from a browser anywhere on the network.

The hosts are components of the Enterprise Virtual Array storage system (see Figure 1–1). These host servers attach to the storage pools of the Enterprise Virtual Array and use the virtual disks just like any other disk resource. To the host server, virtual disks appear the same as other storage system disk resources.
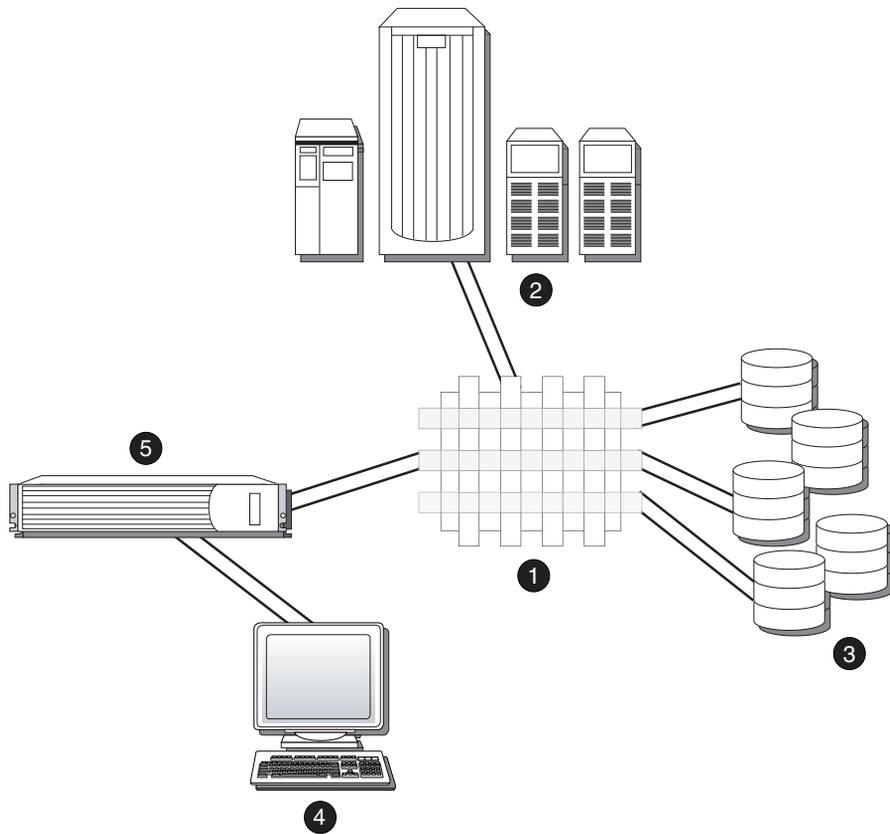
SHR-2503AA

Figure 1−1: SAN architecture

❶   Fabric

❷   Host servers

❸   Enterprise Virtual Array storage systems

❹   Browser for controlling the HSV Element Manager

❺   StorageWorks Management Appliance (SWMA)—where the HSV
        Element Manager resides

Become familiar with the configuration requirements and the physical layout of your
Enterprise Virtual Array (see the documentation that came with your hardware). Also,
read the documentation and release notes for the additional hardware and software
used throughout your SAN.

SAN Considerations

Ensure that your SAN components are all supported for use with the Enterprise Virtual Array. Design your SAN with a Compaq standard topology or by following the Compaq SAN design rules for creating custom topologies. See the *Heterogeneous Open SAN Design Reference Guide* for help with topology rules. The most up-to-date version of this guide can be found on the Compaq website at http://www.compaq.com.

Multiple Path Support

Enterprise Virtual Array requires a multiple path environment. Tru64 UNIX has multiple path as part of the operating system. No additional software is necessary.

OS Unit ID

The HSV Element Manager software displays a field for the OS Unit ID. This field is required for OpenVMS host machines. Compaq recommends using the OS Unit ID with Tru64 UNIX (see "Testing Connectivity to Virtual Disks" on page 3–2).

Host Presentation Compared with Enterprise Modular Array

In previous technology (Enterprise Modular Array family), the term "virtual disks" was used to describe what is presented to the host. Those virtual disks were derived from logical slices of storage, using specific sets of physical disks designed by the administrator.

With Enterprise Virtual Array, the disk pool used to derive virtual disks can be up to the entire set of physical disks available to the controllers. The storage administrator does not need to be concerned with which physical disks are used to create the virtual disks. However, the storage administrator does make decisions regarding disk groups, which are the number of physical disks used.

Virtual disks presented to the host servers with either technology appear like any other SAN disk resource. From the server side, there is no difference. Hosts access the virtual disks as a LUN with capacity knowledge.

Virtual RAID Considerations

Virtual RAID helps determine the level to which user data is protected—VraidX, where X can be 0, 1, or 5. Redundancy is directly proportional to cost in terms of storage usage—the greater the level of data protection, the more storage space is required. There are three types of redundancy available with the Enterprise Virtual Array:

- **Vraid0**—No failure tolerance of data is provided.

- **Vraid1**—All data is duplicated within the storage system. This is the highest level of storage use with the lowest amount of read/write overhead.

- **Vraid5**—All data is protected by parity. This is the lowest level of storage use while maintaining redundancy at a cost of a higher amount of read/write overhead.

2

Installing the Tru64 UNIX Kit

This chapter provides instructions for installing the Tru64 UNIX Kit for Enterprise Virtual Array. This kit lets you connect host servers to the virtual disks of the Enterprise Virtual Array. The following topics are covered:

- Installing the Fibre Channel Adapter, page 2–2

- Installing the Host Kit, page 2–2

  — Installing the Storage System Scripting Utility, page 2–3

  — Checking the Tru64 UNIX Driver Version, page 2–3

  — Updating the Tru64 UNIX Driver, page 2–4

- Upgrading a V1.0 Server to V2.0, page 2–5

Installing the Fibre Channel Adapter

Supported Fibre Channel Adapters (FCAs) must be installed in the host server in order to communicate with the Enterprise Virtual Array. See the platform-specific release notes for a list of supported FCAs. Since the Enterprise Virtual Array requires a multiple path environment, a minimum of two FCAs are required for each host.

NOTE: Traditionally, the adapter used to connect the host server to the fabric is called a Host Bus Adapter (HBA). The HBA used in the Enterprise Virtual Array is called a Fibre Channel Adapter (FCA). You may also see the adapter referred to as a Fibre Channel Host Bus Adapter (FC HBA) in other related documents.

You should follow the rules and conventions for installing hardware for your server type. The Fibre Channel Adapter is shipped with its own documentation for installation. Refer to that documentation for complete instructions.

You need the following items to begin:

- FCA boards and the installation instructions included with the adapter kit

- The computer's hardware manual for instructions on installing adapters

- Appropriate tools to service your computer

The FCA board plugs into a standard PCI slot in the host computer. Refer to the system manual for instructions on plugging in boards.

Installing the Host Kit

The Tru64 UNIX Kit V2.0 for Enterprise Virtual Array installs the Storage System Scripting Utility (SSSU).

The HSV controllers can be configured with SSSU as an alternate to the HSV Element Manager. SSSU allows a command line interface for issuing commands directly to the controller. Complex configuration requests and operations can be handled by either the HSV Element Manager or SSSU. Simple or initial configuration requests can be handled easily and expediently through the element manager, but repetitious and complex configuration steps can be scripted and executed through the command line interface.

IMPORTANT: The SSSU CAPTURE CONFIGURATION command (see the entry on the CAPTURE CONFIGURATION command in the Storage System Scripting Utility reference guide) is the only way to capture, save, and recreate a storage system's configuration.

See the *Storage System Scripting Utility V2.0 for Enterprise Virtual Array Reference Guide* for more information on using SSSU (included on the CD-ROM that comes with this kit).

The supported versions of the Tru64 UNIX operating system include the latest FCA device drivers. No additional setup is required. Check the platform-specific release notes to ensure your operating system meets the minimum requirements for working with the Enterprise Virtual Array. The following section describes how to install SSSU. Refer to "Checking the Tru64 UNIX Driver Version" on page 2–3 to ensure you have the correct version of the FCA driver.

Installing the Storage System Scripting Utility

Follow this procedure to install the SSSU application.

1. Insert the Tru64 UNIX Kit V2.0 for Enterprise Virtual Array CD-ROM.

2. Enter the following command to extract the file (for V5.x):

   # mount –r –t cdfs –o rrip /dev/disk/cdrom0c /mnt

NOTE: Substitute your CD device for cdrom0c if necessary.

3. Change directories on the CD-ROM by entering:

   # cd /mnt

4. Enter the following command to run the installation program:

   # setld –l .

NOTE: The –l is a lowercase L.

   The installation asks if you want to install the listed subsets.

5. Enter the corresponding number of the software and press **Return**.

6. Select option y and press **Return**.

7. The SSSU executable is installed in the following directory:

   #/usr/opt/ENTP002/sbin

   To execute SSSU:

   • Command line version—simply type: #sssu

Checking the Tru64 UNIX Driver Version

The FCA device driver for Tru64 UNIX (Emx) is installed as part of the operating system. You may wish to check the version to ensure compatibility with the Enterprise Virtual Array (see the platform-specific release notes for supported versions).

As your system starts up, startup information is logged in a file called `messages` where, among many other items, the version of the installed FCA drivers are recorded.

1. View the `messages`

   # cd /var/adm
   # view *messages*

2. Inspect the latest entry line for KGPSA (the FCA name).

   The following is an example:

   Jan 17 vmunix: KGPSA-CA : Driver Rev 1.32: F/W Rev 3.81A4(2.01A0) : wwn 1000-0000-c923-102

3. Ensure that you have version 1.29 or above for each of your FCAs.

   If your driver version is not at 1.29 or above, you should update it.

Updating the Tru64 UNIX Driver

This driver is for an Emulex LP8000 or LP9002, using the SLI2 programming interface. The adapter minimum firmware revision that should be run in conjunction with the driver is:

- LP8000 - V3.81a4

- LP9002 - V3.81a4

NOTE: LP8000s require console V5.6 or above on the Alpha platforms. Check your platform configuration to make sure that all components are up to date.

To prepare the system for the Emulex device driver and potentially prepare it to build the Emulex sources, invoke the corresponding BINSTALL script for your operating system revision.

   BINSTALL.V51a

If you must manually add support for the LP8000 to your system, add an entry for the LP8000 to /etc/sysconfigtab under the "emx:" section.

1. Edit /etc/sysconfigtab with a text editor.

2. Search for "emx:" in the file. You should see the following lines:

   emx:

   #   PCI_Option = <..., Device_Id - 0x1ae5, ...>

   #   PCI_Option = <..., Device_Id - 0xf700, ...>

3. Copy and paste a copy of the line with Device_ID 0xf700, replacing the 0xf700 with 0xf800 and removing the opening '#' character.

The results should look like the following:

emx:

   PCI_Option = <..., Device_Id - 0xf800, ...>

   #   PCI_Option = <..., Device_Id - 0x1ae5, ...>

   #   PCI_Option = <..., Device_Id - 0xf700, ...>

4. Write the contents of the file.

5. If you are not already logged in as root, change to root with su root.

6. Rebuild the kernel (run the doconfig command), copy the new kernel to /vmunix, and reboot the system.

Upgrading a V1.0 Server to V2.0

Follow this procedure to upgrade an existing V1.0 host server to V2.0.

1. Upgrade the Virtual Controller Software (VCS) and the HSV Element Manager software from V1.0 to V2.0 (see the upgrade instructions that come with the VCS kit).

2. Locate and delete the V1.0 SSSU executable.

3. Install the V2.0 kit (see "Installing the Host Kit" on page 2–2).

   You may also need to upgrade the FCA firmware. Check the platform-specific release notes for supported FCA versions.

3

Testing Connections to the Enterprise Virtual Array

This chapter describes how to test that the Tru64 UNIX host servers can connect to the Enterprise Virtual Array. This chapter also describes how to test that the Storage System Scripting Utility (SSSU) can connect to the Enterprise Virtual Array. The following topics are covered:

- Testing the Storage System Scripting Utility, page 3–2

- Testing Connectivity to Virtual Disks, page 3–2

  — Host Connectivity to the Fabric, page 3–2

  — Adding Hosts, page 3–2

  — Creating and Presenting Virtual Disks, page 3–3

  — Verifying Virtual Disks from the Host, page 3–3

- Configuring Virtual Disks from the Host, page 3–5

  — Creating Partitions on a Virtual Disk Using disklabel, page 3–5

  — Creating a File System on a Virtual Disk, page 3–6

  — Mounting the File System, page 3–6

Testing the Storage System Scripting Utility

The Storage System Scripting Utility (SSSU) is installed from the Tru64 UNIX Kit V2.0 for Enterprise Virtual Array (see "Installing the Storage System Scripting Utility" on page 2–3). You should verify that SSSU runs from your host server.

The SSSU executable is installed in the following directory by default:

```
# /usr/opt/ENTP002/sbin
```

To start SSSU, type the following at the command prompt:

```
# sssu
```

Testing Connectivity to Virtual Disks

Virtual disks need to be set up and presented to your host server in order to test the connectivity and begin using them. This section briefly describes how to create, present, and access virtual disks.

Host Connectivity to the Fabric

Once you have installed the FCA, you need to connect the host to the switches. See "Connecting Hosts to Switches" on page 4–2 for the specific cabling instructions.

Adding Hosts

You add a host using the HSV Element Manager. You need to add each FCA installed in the host system in order for the host to work with the Enterprise Virtual Array.

1. Collect information on the WWN for each FCA on your server. You need this information when choosing the host FCAs in the HSV Element Manager.

   Adding hosts through the HSV Element Manager software consists of adding each FCA adapter installed in the host. The first step uses the **Add Host** function and each subsequent adapter uses the **Add Port** function. Ensure that you add a port for each active FCA.

2. Add the host from the HSV Element Manager.

3. Ensure that the host FCAs have been added by inspecting the **Host** folder in the **Navigation** tree of the HSV Element Manager.

Creating and Presenting Virtual Disks

This section describes how to create and present additional virtual disks to host servers and verify that the hosts can access them.

1. Create a virtual disk family on the Enterprise Virtual Array using the HSV Element Manager.

2. Set values for the following:

   - Virtual Disk Name

   - Vraid level

   - Size

   - Present to Host (preferably to the host you just created)

   - OS Unit ID should be set to a unique value for each virtual disk (see OS Unit ID on page 3–3)

3. If you chose a specific LUN on the **Virtual Disk Properties** page, you are prompted to select a LUN number.

4. Restart the server or rescan the bus.

OS Unit ID

There is a field on the **Virtual Disk** properties page for OS Unit ID. This is an optional field for Tru64 UNIX. Compaq recommends that each virtual disk have a unique ID number. You can take advantage of this field to easily track the virtual disks presented to the Tru64 UNIX hosts. When you list the drives using hwmgr, they appear with the number you give it in the OS Unit ID field instead of the P/T/L information.

You need to make the OS Unit ID values unique across the entire SAN, not just the controller subsystem.

Verifying Virtual Disks from the Host

This section describes how to verify that Tru64 UNIX hosts can access the virtual disks created on the Enterprise Virtual Array with the HSV Element Manager. The host can recognize Enterprise Virtual Array devices either through a system restart or by rescanning the bus.

Scanning the Bus

1. You scan the bus for the virtual disk and then display information about the Enterprise Virtual Array devices. Start the rescan of the FCA with `hwmgr`.

   # hwmgr -scan scsi

   All new virtual disks become visible to the host.

2. Assign device special files to the new virtual disks with `dsfmgr`.

   # dsfmgr -K

Uppercase "K" reassigns device special files to all devices. Lowercase "k" only assigns device special files to the new devices—in this case, the virtual disks.

NOTE: You could also restart the host machine to rescan for newly presented virtual disks, but restarting the host is not a requirement.

Sample Output from a Rescan

   # hwmgr -view dev

HWID: Device Name	Mfg	Model	Location
4: /dev/kevm			
30: /dev/disk/floppy0c		3.5in floppy	fdi0-unit-0
39: /dev/disk/cdrom0c	COMPAQ	CRD-8402B	bus-1-targ-0-lun-0
40: /dev/disk/dsk0c	COMPAQ	BB00923468	bus-2-targ-0-lun-0
41: /dev/disk/dsk1c	COMPAQ	BB00923468	bus-2-targ-1-lun-0
42: /dev/dmapi/dmapi			
47: /dev/cport/scp0	COMPAQ	HSV110 (C) COMPA	bus-3-targ-0-lun-0
48: /dev/disk/dsk2c	COMPAQ	HSV110 (C) COMPA	IDENTIFIER=2

View the `hwmgr` and `dsfmgr` man pages for more information on using these commands.

If you are unable to access the virtual disk, check the following:

- Verify all cabling to the switch, Enterprise Virtual Array, and host.

- Verify all firmware levels (check the Enterprise Virtual Array QuickSpecs and associated release notes for details).

- Ensure you are running a supported version of the host operating system (see the release notes included with this kit).

- Ensure that the correct host is selected for this virtual disk and that a unique OS Unit ID is used in the HSV Element Manager.

- Ensure Tru64 UNIX is selected as the operating system for the virtual disk in the HSV Element Manager.

Configuring Virtual Disks from the Host

Once you have set up the virtual disks on the Enterprise Virtual Array and rescanned or restarted the host, you need to follow the host-specific conventions for configuring these new disk resources. These disk resources then become usable just like any other disk resource to the host system. Refer to the documentation that came with your server for specific instructions on setting up disk resources.

Tru64 UNIX treats an HSV110 virtual disk much like a SCSI disk; therefore, to prepare your virtual disk for access by the Tru64 UNIX file system, you must do the following:

- Create a file system on the virtual disk

- Mount the file system (to gain access to the virtual disk)

Creating Partitions on a Virtual Disk Using disklabel

Create the partitions on a virtual disk by issuing a disklabel command. The disklabel command partitions the virtual disk for access by Tru64 UNIX. Tru64 UNIX defines only partitions a, b, c, and g for the HSV110 controller. In addition, refer to the rz and disktab man pages for more information about the disklabel utility.

To create the partitions on a virtual disk using the default partition sizes, enter the following:

    # disklabel -rw *disk*

Where:

    disk is the character device name found in /dev/rdisk

For example, to create partitions on block device, *dsk8,* enter:

    # disklabel -rw dsk8

To view a virtual disk partition, enter:

    # disklabel -r *device*

Creating a File System on a Virtual Disk

The `newfs` command is given here as an example for creating a UNIX File System (UFS). For Advanced File System (AdvFS) and for making devices available for Logical Storage Manager (LSM), similar types of commands exist. For additional information, consult the related documentation.

Use the `newfs` command to create a UFS file system on a virtual disk the same way that you would create a file system on a disk device by entering the following:

    # newfs *special-device disk-type*

For example:

    # newfs   /dev/rdisk/dsk8c

Where:

• *special-device* is the appropriate device special file (see the example below)

To create a UFS file system using the above examples, enter the following:

    # newfs /dev/rdisk/dsk8c   HSV110

This creates a UFS file system on the `c` partition (whole disk-device).

Mounting the File System

To access the virtual disk, mount it as a device file system to a mount point. For example:

    # mount   /dev/disk/dsk8c   /mnt

To view the mounted file system, enter:

    # df

The virtual disk is now accessible to the file system just as a disk device would be. The file system cannot see the RAID functionality and number of physical devices attached to the HSV110 controller. This device appears as a single virtual disk or "disk" to the user as viewed by the file system.

4

Additional Host Considerations

This chapter includes additional considerations for the host operating systems in an Enterprise Virtual Array. The following topics are covered:

- Connecting Hosts to Switches, page 4–2

- Working with Clusters, page 4–3

- Expanding Volumes, page 4–4

Connecting Hosts to Switches

This section describes how to connect your host servers to the SAN switches in order to access the virtual disks. The documentation that came with the Enterprise Virtual Array hardware describes how to cable the controllers to the SAN switches.

Figure 4–1 provides an overview of connecting the Enterprise Virtual Array components.

FP = Fibre (Host) Port
FCA = Fibre Channel Adapter

CXO7603A
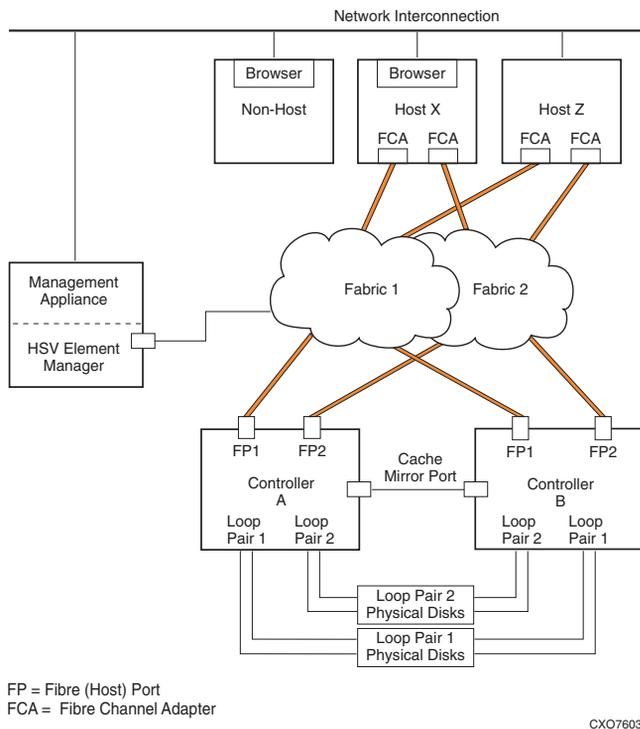
Figure 4–1: Overview of Enterprise Virtual Array component connections

Each host must attach to two redundant switches (fabrics) using standard fiber channel cables. These two switches in turn attach to the two HSV110 controllers on the Enterprise Virtual Array.

NOTE: This is the most basic configuration. More than two SAN switches can be used.

1. Plug one end of the fiber channel cables into the FCAs on the servers. There must be at least two FCAs.

2. Plug the other ends of the cables into the two switches. You must plug each cable from the host's FCAs into different switches.

Figure 4–2 displays the connections between a single server with two FCAs and two switches.

SHR-2457A
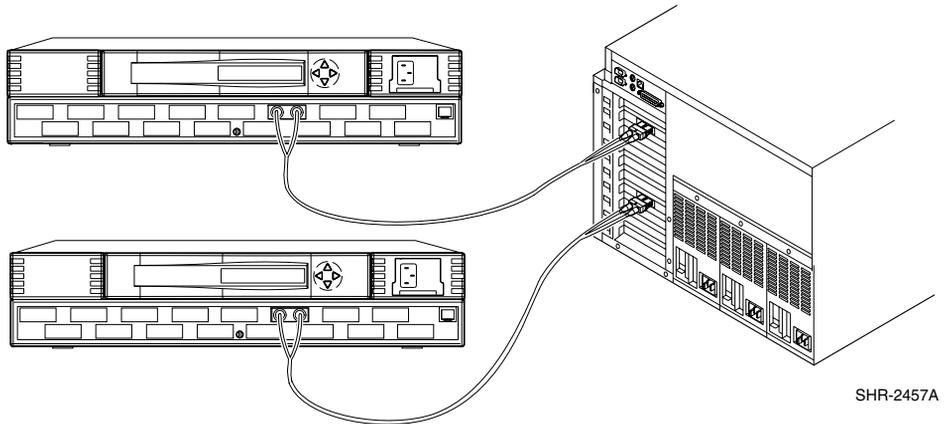
Figure 4–2: Cabling Hosts to Switches

Working with Clusters

Clustering is the connecting of two or more computers together in such a way that they behave like a single computer. Clustering is used for parallel processing, load balancing, and fault tolerance.

The supported clustering software for Tru64 UNIX is TruCluster. See the platform-specific release notes for the specific version of the supported clustering software.

Be sure to install the Tru64 UNIX Kit on each server in the cluster. Refer to the documentation that came with the clustering software for more specific information.

Expanding Volumes

You can increase the volume size of the virtual disk with the HSV Element Manager. You simply change the size parameter on the particular virtual disk's property page.

NOTE: Once a virtual disk is created, you can only increase the size, not decrease it.

Tru64 UNIX does not support dynamically growing a virtual disk volume. If you need to increase the size of a virtual disk, use the workaround that follows.

IMPORTANT: As with any large-scale operation, it is always best to first create a backup of your data.

Volume Expansion Workaround

1. Copy any existing data to a safe place (tape or other device).

2. Stop all applications using the virtual disk.

3. Unattach the host server from the virtual disk.

4. Expand the virtual disk using the HSV Element Manager.

5. Reconnect the host to the virtual disk.

6. Format and partition as necessary.

7. Restore data from the backup.

   The expanded virtual disk is available for use.

Glossary

This glossary defines Enterprise Virtual Array terms used in this publication or related to this product and is not a comprehensive glossary of computer terms.

active virtual disk

A virtual disk (VD) is a simulated disk drive created by the controllers as storage for one or more hosts. An active virtual disk is accessible by one or more hosts for normal storage. An active virtual disk and its snapshot, if one exists, constitute a virtual disk family. An active virtual disk is the only necessary member of a virtual disk family.

See also **virtual disk, virtual disk copy, virtual disk family,** *and* **snapshot.**

adapter

See **controller.**

array

All the physical disk drives in a storage system that are known to and under the control of a controller pair.

array controller

See **controller.**

cable assembly

A fiber optic cable that has connectors installed on one or both ends. General use of these cable assemblies includes the interconnection of multimode fiber optic cable assemblies with either LC or SC type connectors.

- When there is a connector on only one end of the cable, the cable assembly is referred to as a pigtail.

- When there is a connector on both ends of the cable, the cable assembly is referred to as a jumper.

cache

High-speed memory that sets aside data as an intermediate data buffer between a host and the storage media. The purpose of cache is to improve performance.

See also **read cache, write cache,** *and* **mirrored cache.**

communication logical unit number (LUN)

See **console LUN.**

console LUN

A SCSI-3 virtual object that makes a controller pair accessible by the host before any virtual disks are created. *Also* called a communication LUN.

console LUN ID

The ID that can be assigned when a host operating system requires a unique ID. The console LUN ID is assigned by the user, usually when the storage system is initialized.

See also **console LUN.**

controller

A hardware/firmware device that manages communications between host systems and other devices. Controllers typically differ by the type of interface to the host and provide functions beyond those the devices support.

controller pair

Two interconnected controller modules which together control a physical disk array. A controller pair and the disk array together constitute a storage system.

Element Manager

The graphical user interface through which a user controls and monitors a storage system. The HSV Element Manager software can be installed on more than one management appliance in a fabric. Each installation of the HSV Element Manager software is a management agent.

Enterprise Virtual Array

The Compaq name used to describe the storage system that includes HSV controllers, storage devices, enclosures, cables, and power supplies. *Also* called the Enterprise Storage System.

fabric

A Fibre Channel fabric switch or two or more interconnected Fibre Channel switches allowing data transmission.

fabric port

A port which is capable of supporting an attached arbitrated loop. This port on a loop will have the AL_PA hexadecimal address 00 (loop ID 7E), giving the fabric the highest priority access to the loop. A loop port is the gateway to the fabric for the node ports on a loop.

failover

The process that takes place when one controller assumes the workload of a failed companion controller. Failover continues until the failed controller is operational.

FC HBA

Fibre Channel Host Bus Adapter. An interchangeable term for Fibre Channel adapter.

See also **FCA.**

FCA

Fiber Channel Adapter. An adapter used to connect the host server to the fabric. *Also* called a Host Bus Adapter (HBA) or a Fibre Channel Host Bus Adapter (FC HBA).

See also **FC HBA**

fiber

The optical media used to implement Fibre Channel.

fibre

The international spelling that refers to the Fibre Channel standards for optical media.

Fibre Channel

A data transfer architecture designed for mass storage devices and other peripheral devices that require very high bandwidth.

Fibre Channel adapter

See **FCA.**

GB

Gigabyte. A unit of measurement defining either:

- A data transfer rate.

  See also **GBps**

- A storage or memory capacity of 1,073,741,824 (2^{30}) bytes.

GBps

Gigabytes per second. A measurement of the rate at which the transfer of bytes of data occurs. A GBps is a transfer rate of 1,000,000,000 (10^9) bits per second.

See also **GB.**

Giga (G)

The notation to represent 10^9 or 1 billion (1,000,000,000).

HBA

Host Bus Adapter.

See **FCA.**

host

A computer that runs user applications and uses (or can potentially use) one or more virtual disks created and presented by the controller pair.

Host Bus Adapter

See **FCA.**

host computer

See **host.**

host ports

A connection point to one or more hosts through a Fibre Channel fabric. A host is a computer that runs user applications and that uses (or can potentially use) one or more of the virtual disks that are created and presented by the controller pair.

host-side ports

See **host ports.**

HSV Element Manager

See **Element Manager.**

initialization

A process that prepares a storage system for use. Specifically, the system binds controllers together as an operational pair and establishes preliminary data structures on the disk array. Initialization also sets up the first disk group, called the default disk group.

K

Kilo. A scientific notation denoting a multiplier of one thousand (1,000).

KB

Kilobyte. A unit of measurement defining either storage or memory capacity.

1. For storage, a KB is a capacity of 1,000 (10^3) bytes of data.

2. For memory, a KB is a capacity of 1,024 (2^{10}) bytes of data.

LUN

Logical Unit Number. A unique identifier used on a SCSI bus to distinguish between devices that share the same bus. SCSI is a parallel interface that allows up to eight devices to be connected along a single cable.

management agent

The HSV Element Manager software that controls and monitors the Enterprise Storage System. The software can exist on more than one management appliance in a fabric. Each installation of the HSV Element Manager software is a management agent.

MB

Megabtye. A term defining either:

*   A data transfer rate.

    See also **MBps.**

*   A measure of either storage or memory capacity of 1,048,576 (2^{20}) bytes.

MBps

Megabytes per second. A measure of bandwidth or data transfers occurring at a rate of 1,000,000 (10^6) bytes per second.

Mega

A notation denoting a multiplier of 1 million (1,000,000).

metadata

Information that a controller pair writes on the disk array. This information is used to control and monitor the array and is not readable by the host.

mini-port driver

The Windows driver type used for connecting to the Enterprise Virtual Array through the SAN fabric.

mirrored caching

A process in which half of each controller's write cache mirrors the companion controller's write cache. The total memory available for cached write data is reduced by half, but the level of protection is greater.

mirroring

The act of creating an exact copy or image of data.

node port

A device port that can operate on the arbitrated loop topology.

NSC

Network Storage Controller. The HSV Controllers used by the Enterprise Storage System.

OSM

Open SAN Manager. A centralized, appliance-based monitoring and management interface that supports multiple applications, operating systems, hardware platforms, storage systems, tape libraries and SAN-related interconnect devices. It is included and resides on the SANworks Management Appliance, a single aggregation point for data management.

password

A security interlock whose purpose is to allow:

- a management agent control of only certain storage systems

- only certain management agents control of a storage system

physical disk

A disk drive mounted in a disk drive enclosure that communicates with a controller pair through the device-side Fibre Channel loops. A physical disk is hardware with embedded software, as opposed to a virtual disk, which is constructed by the controllers. Only the controllers can communicate directly with the physical disks.

The physical disks, in aggregate, are called the array and constitute the storage pool from which the controllers create virtual disks.

physical disk array

See **array**.

port

A Fibre Channel connector on a Fibre Channel device.

port_name

A 64-bit unique identifier assigned to each Fibre Channel port. The port_name is communicated during the login and port discovery processes.

preferred path

A preference for which controller of the controller pair manages the virtual disk. This preference is set by the user through the HSV Element Manager when creating the virtual disk. A host can change the preferred path of a virtual disk at any time. The primary purpose of preferring a path is load balancing.

read caching

A cache method used to decrease subsystem response times to a read request by allowing the controller to satisfy the request from the cache memory rather than from the disk drives. Reading data from cache memory is faster than reading data from a disk. The read cache is specified as either on or off for each virtual disk. The default state is on.

read ahead caching

A cache management method used to decrease the subsystem response time to a read request by allowing the controller to satisfy the request from the cache memory rather than from the disk drives.

redundancy

1. Element Redundancy—The degree to which logical or physical elements are protected by having another element that can take over in case of failure. For example, each loop of a device-side loop pair normally work independently but can take over for the other in case of failure.

2. Data Redundancy—The level to which user data is protected. Redundancy is directly proportional to cost in terms of storage usage; the greater the level of data protection, the more storage space is required.

SANworks by Compaq HSV Element Manager

See **HSV Element Manager.**

SCSI

1. Small Computer System Interface. An American National Standards Institute (ANSI) interface which defines the physical and electrical parameters of a parallel I/O bus used to connect computers and a maximum of 16 bus elements.

2. The communication protocol used between a controller pair and the hosts. Specifically, the protocol is FC-AL or SCSI on a Fibre Channel. SCSI is the higher command-level protocol and Fibre Channel is the low-level transmission protocol. The controllers have full support for SCSI-2; additionally, they support some elements of SCSI-3.

SCSI-3

The ANSI standard that defines the operation and function of Fibre Channel systems.

small computer system interface

See **SCSI.**

SMART

Self-Monitoring, Analysis, and Reporting Technology. An industry standard for running internal diagnostics on disk drives and making predictive failure analysis. If a physical disk drive reports an imminent failure, the controllers redistribute data to other drives. The SMART message is reported and should trigger the immediate corrective action to replace the drive.

snapshot

A temporary virtual disk (VD) that reflects the contents of another virtual disk at a particular point in time. A snapshot operation is only done on an active virtual disk. Only one snapshot of an active virtual disk can exist at any point. The active disk and its snapshot constitute a virtual family.

See also **active virtual disk, virtual disk copy,** *and* **virtual disk family.**

SSN

Storage System Name. An HSV Element Manager-assigned, unique 20-character name that identifies a specific storage system.

storage pool

The aggregated blocks of available storage in the total physical disk array.

storage system

The controllers, storage devices, enclosures, cables, and power supplies and their software.

Storage System Scripting Utility (SSSU)

A command-line application supplied in the host operating system kits that allows for control of the Enterprise Virtual Array.

Storage System Name

See **SSN.**

StorageWorks by Compaq

The Compaq trademarked name used to describe the set of rack-mounted enclosures containing controllers, transceivers, I/O modules, EMUs, disk drives, cables, blowers, and power supplies used to design and configure a solution-specific storage system.

switch

An electro-mechanical device that initiates an action or completes a circuit.

TB

Terabyte. A term defining either:

- A data transfer rate.

  See also **TBps.**

- A measure of either storage or memory capacity of 1,099,5111,627,776 (2^{40}) bytes.

TBps

Terabytes per second. A data transfer rate of 1,000,000,000,000 (10^{12}) bytes per second.

uninitialized system

A state in which the storage system is not ready for use.

See also **initialization.**

Virtual Disk

A simulated disk drive created by the controllers as storage for one or more hosts. The virtual disk characteristics, chosen by the storage administrator, provide a specific combination of capacity, availability, performance, and accessibility. A controller pair simulates the characteristics of the virtual disk by deploying the disk group from which the virtual disk was created.

The host computer sees the virtual disk as "real," with the characteristics of an identical physical disk.

See also **active virtual disk, virtual disk copy, virtual disk family,** *and* **virtual disk snapshot**

virtual disk copy

A clone or exact replica of another virtual disk at a particular point in time. Only an active virtual disk can be copied. A copy immediately becomes the active disk of its own virtual disk family.

See also **active virtual disk, virtual disk family,** *and* **virtual disk snapshot**

virtual disk family

A virtual disk and its snapshot, if a snapshot exists, constitute a family. The original virtual disk is called the active disk. When you first create a virtual disk family, the only member is the active disk.

See also **active virtual disk, virtual disk copy,** *and* **virtual disk snapshot.**

virtual disk snapshot

See **snapshot.**

Vraid0

A virtualization technique that provides no data protection. Data from the host is broken down into chunks and distributed on the disks comprising the disk group from which the virtual disk was created. Reading and writing to a Vraid0 virtual disk is very fast and makes the fullest use of the available storage, but there is no data protection (redundancy) unless there is parity.

Vraid1

A virtualization technique that provides the highest level of data protection. All data blocks are mirrored or written twice on separate physical disks. For read requests, the block can be read from either disk, which can increase performance. Mirroring takes the most storage space because twice the storage capacity must be allocated for a given amount of data.

Vraid5

A virtualization technique that uses parity striping to provide moderate data protection. Parity is a data protection mechanism for a striped virtual disk. A striped virtual disk is one whose data to and from the host is broken down into chunks and distributed on the physical disks comprising the disk group in which the virtual disk was created. If the striped virtual disk has parity, another chunk (a parity chunk) is calculated from the set of data chunks and written to the physical disks. If one of the data chunks becomes corrupted, the data can be reconstructed from the parity chunk and the remaining data chunks.

World Wide Name

See **WWN.**

write back caching

A controller process that notifies the host that the write operation is complete when the data is written to the cache. This occurs before transferring the data to the disk. Write back caching improves response time since the write operation completes as soon as the data reaches the cache. As soon as possible after caching the data, the controller then writes the data to the disk drives.

write caching

A process when the host sends a write request to the controller, and the controller places the data in the controller cache module. As soon as possible, the controller transfers the data to the physical disk drives.

WWN

World Wide Name. A unique Fibre Channel identifier consisting of a 16-character hexadecimal number. A WWN is required for each Fibre Channel communication port.

Index

W

Warning
  electrical shock hazard symbol, defined viii
  excessive weight symbol, defined ix
  hot surface symbol, defined ix
  multiple power source symbol, defined ix

  network interface connection symbol, defined
   viii
  symbol and definition viii
websites
  Compaq storage xi
  Compaq technical support x
Windows 2000 1–4
Windows NT 1–4