HP EVA 4000/6000/8000 and EVA 4100/6100/8100 User Guide

Abstract

This document is intended for customers who operate and manage the EVA 4000/6000/8000 and EVA 4100/6100/8100 storage systems. These models are sometimes referred to as EVA4x00, EVA6x00, and EVA8x00 or as EVAx000 and x100.

IMPORTANT: With the release of the P6300/P6500 EVA, the EVA family name has been rebranded to HP P6000 EVA. The names for all existing EVA array models will not change. The rebranding also affects related EVA software. The following product names have been rebranded:

- HP P6000 Command View (formerly HP StorageWorks Command View EVA)
- HP P6000 Business Copy (formerly HP StorageWorks Business Copy EVA)
- HP P6000 Continuous Access (formerly HP StorageWorks Continuous Access EVA)
- HP P6000 Performance Data Collector (formerly EVAPerf)

HP Part Number: 5697-1119
Published: January 2012
Edition: 12

Warranty

To obtain a copy of the warranty for this product, see the warranty information website:

http://www.hp.com/go/storagewarranty

Acknowledgements

Contents

1 Enterprise Virtual Array startup

This chapter describes the procedures to install and configure the Enterprise Virtual Array. When these procedures are complete, you can begin using your storage system.

NOTE: Installation of the Enterprise Virtual Array should be done only by an HP authorized service representative. The information in this chapter provides an overview of the steps involved in the installation and configuration of the storage system.

This chapter consists of:

EVA8000/8100 storage system connections

Figure 1 (page 11) shows how the storage system is connected to other components of the storage solution.

- The HSV210-A and HSV210-B controllers connect via four host ports (FP1, FP2, FP3, and FP4) to the Fibre Channel fabrics. The hosts that will access the storage system are connected to the same fabrics.
- The HP P6000 Command View management server also connects to the fabric.
- The controllers connect through two loop pairs to the drive enclosures. Each loop pair consists of two independent loops, each capable of managing all the disks should one loop fail. Four FC loop switches are used to connect the controllers to the disk enclosures.

Figure 1 EVA8000/8100 configuration

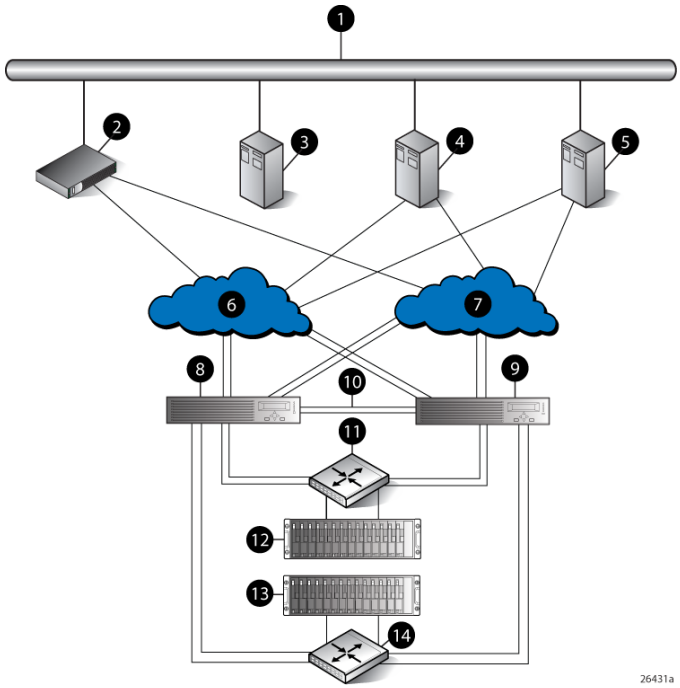

1 Network interconnection
2 Management server
3 Non-host
4 Host X
5 Host Z
6 Fabric 1
7 Fabric 2
8 Controller A
9 Controller B
10 Cache mirror ports
11 FC loop switch
12 Drive enclosure 1
13 Drive enclosure 2
14 FC loop switch

EVA6000/6100 storage system connections

Figure 2 (page 12) shows a typical EVA6000/6100 SAN topology:

- The HSV200-A and HSV200-B controllers connect via two host ports (FP1 and FP2) to the Fibre Channel fabrics. The hosts that will access the storage system are connected to the same fabrics.
- The HP Command View EVA management server also connects to both fabrics.
- The controllers connect through one loop pair to the drive enclosures. The loop pair consists of two independent loops, each capable of managing all the disks should one loop fail. Two FC loop switches are used to connect the controllers to the disk enclosures.

Figure 2 EVA6000/6100 configuration

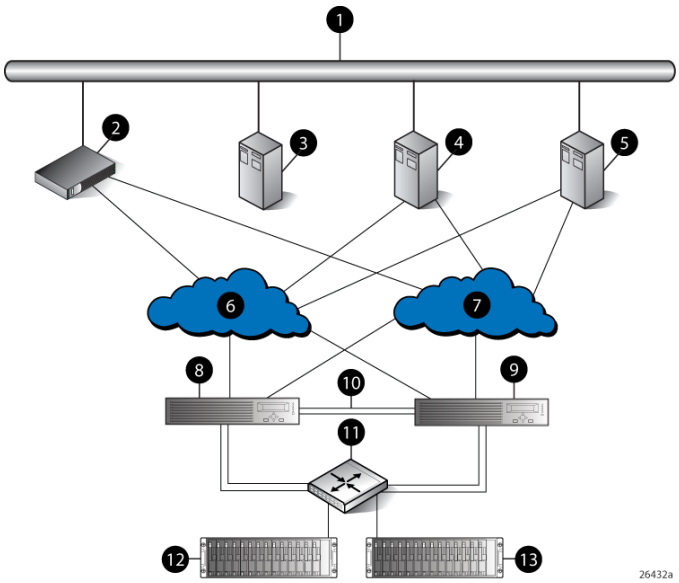

1 Network interconnection
2 Management server
3 Non-host
4 Host X
5 Host Z
6 Fabric 1
7 Fabric 2
8 Controller A
9 Controller B
10 Cache mirror ports
11 FC loop switch
12 Drive enclosure 1
13 Drive enclosure 2

EVA4000/4100 storage system connections

Figure 3 (page 13) shows a typical EVA 4000/4100 SAN topology:

- The HSV200-A and HSV200-B controllers connect via two host ports (FP1 and FP2) to the Fibre Channel fabrics. The hosts that will access the storage system are connected to the same fabrics.
- The HP P6000 Command View management server also connects to both fabrics.
- The controllers connect through one loop pair to the drive enclosures. The loop pair consists of two independent loops, each capable of managing all the disks should one loop fail. The controllers connect directly to the disk enclosures.

Figure 3 EVA4000/4100 configuration

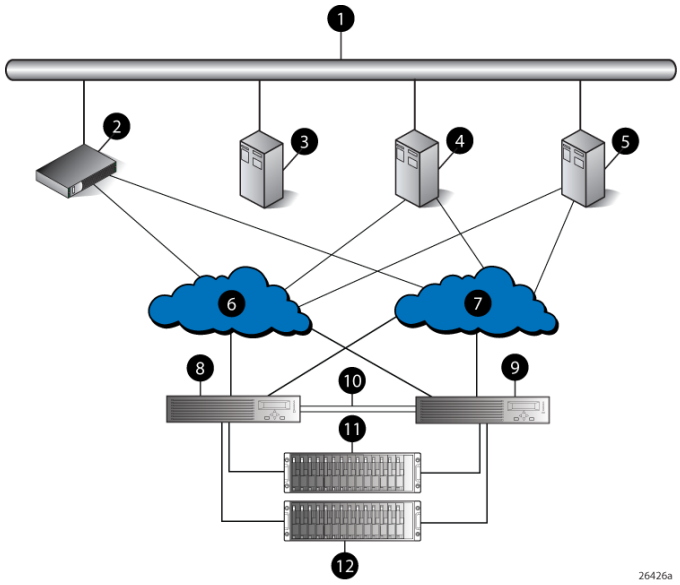

1 Network interconnection	7 Fabric 2
2 Management server	8 Controller A
3 Non-host	9 Controller B
4 Host X	10 Cache mirror ports
5 Host Z	11 Drive enclosure 1
6 Fabric 1	12 Drive enclosure 2

Direct connect

NOTE: Direct connect is currently supported on Microsoft Windows only. For more information on direct connect, go the Single Point of Connectivity Knowledge (SPOCK) at: http://www.hp.com/storage spock.

Direct connect provides a lower cost solution for smaller configurations. When using direct connect, the storage system controllers are connected directly to the host(s), not to SAN Fibre Channel switches. Make sure the following requirements are met when configuring your environment for direct connect:

- A management server running HP P6000 Command View must be connected to one port on each EVA controller. The management host must use dual HBAs for redundancy.
- To provide redundancy, it is recommended that dual HBAs be used for each additional host connected to the storage system. Using this configuration, up to four hosts (including the management host) can be connected to an EVA8x00, and up to two hosts can be connected to an EVA6x00 or EVA4x00.
- The Host Port Configuration must be set to Direct Connect using the OCP.
- HP P6000 Continuous Access cannot be used with direct connect configurations.
- The HSV controller firmware cannot differentiate between an empty host port and a failed host port in a direct connect configuration. As a result, the Connection state dialog box on the Controller Properties window displays *Connection failed for an empty host port*. To fix this problem, insert an optical loop-back connector into the empty host port; the Connection state will display *Connected*. For more information about optical loop-back connectors, contact your HP-authorized service provider.

iSCSI connection configurations

The EVA4x00/6x00/8x00 support iSCSI attach configurations using the HP MPX100. Both fabric connect and direct connect are supported for iSCSI configurations. For complete information on iSCSI configurations, go to the following website:

http://h18006.www1.hp.com/products/storageworks/evaiscsiconnect/index.html

NOTE: An iSCSI connection configuration supports mixed direct connect and fabric connect.

Fabric connect iSCSI

Fabric connect provides an iSCSI solution for EVA Fibre Channel configurations that want to continue to use all EVA ports on FC or if the EVA is also used for HP P6000 Continuous Access.

Make sure the following requirements are met when configuring your MPX100 environment for fabric connect:

- A maximum of two MPX100s per storage system are supported
- Each storage system port can connect to a maximum of two MPX100 FC ports.
- Each MPX100 FC port can connect to a maximum of one storage system port.
- In a single MPX100 configuration, if both MPX100 FC ports are used, each port must be connected to one storage system controller.
- In a dual MPX100 configuration, at least one FC port from each MPX100 must be connected to one storage system controller.
- The Host Port Configuration must be set to Fabric Connect using the OCP.
- HP P6000 Continuous Access is supported on the same storage system connected in MPX100 fabric connect configurations.

Direct connect iSCSI

Direct connect provides a lower cost solution for configurations that want to dedicate controller ports to iSCSI I/O. When using direct connect, the storage system controllers are connected directly to the MPX100(s), not to SAN Fibre Channel switches.

Make sure the following requirements are met when configuring your MPX100 environment for direct connect:

- A maximum two MPX100s per storage system are supported.
- In a single MPX100 configuration, if both MPX100 FC ports are used each port must be connected to one storage system controller.
- In a dual MPX100 configuration, at least one FC port from each MPX100 must be connected to one storage system controller.
- The Host Port Configuration must be set to Direct Connect using the OCP.
- HP P6000 Continuous Access cannot be used with direct connect configurations.
- EVAs cannot be directly connected to each other to create an HP P6000 Continuous Access configuration. However, hosts can be directly connected to the EVA in an HP P6000 Continuous Access configuration. At least one port from each array in an HP P6000 Continuous Access configuration must be connected to a Fabric connection for remote array connectivity.

Procedures for getting started

Step	Responsibility
1. Gather information and identify all related storage documentation.	Customer
2. Contact an authorized service representative for hardware configuration information.	Customer
3. Enter the World Wide Name (WWN) into the OCP.	HP Service Engineer
4. Configure HP P6000 Command View.	HP Service Engineer
5. Prepare the hosts.	Customer
6. Configure the system through HP P6000 Command View.	HP Service Engineer
7. Make virtual disks available to their hosts. See the storage system software documentation for each host's operating system.	HP Service Engineer

Gathering information

The following items should be available when installing and configuring an Enterprise Virtual Array. They provide information necessary to set up the storage system successfully.

- *HP 4x00/6x00/8x00 Enterprise Virtual Array World Wide Name label*, which is shipped with the system
- *HP EVA 4000/6000/8000 and EVA 4100/6100/8100 Read Me First*
- *HP EVA 4000/6000/8000 and EVA 4100/6100/8100 Release Notes (XCS 6.250)*
- The latest HP P6000 Command View software (Check the *HP P6000 Enterprise Virtual Array Compatibility Reference* for controller software and HP P6000 Command View compatibility.)

Locate these items and keep them handy. You will need them for the procedures in this manual.

Host information

Make a list of information for each host computer that will be accessing the storage system. You will need the following information for each host:

- The LAN name of the host
- A list of World Wide Names of the FC adapters, also called host bus adapters, through which the host will connect to the fabric that provides access to the storage system, or to the storage system directly if using direct connect.
- Operating system type
- Available LUN numbers

Setting up a controller pair using the OCP

NOTE: This procedure should be performed by an HP authorized service representative.

Two pieces of data must be entered during initial setup using the controller OCP:

- World Wide Name (WWN) — Required to complete setup. This procedure should be performed by an HP authorized service representative.
- Storage system password — Optional. A password provides security allowing only specific instances of HP P6000 Command View to access the storage system.

The OCP on either controller can be used to input the WWN and password data. For more information about the OCP, see "Operator control panel" (page 43).

Table 1 (page 16) lists the push-button functions when entering the WWN, WWN checksum, and password data.

Table 1 Push button functions

Button	Function
▲	Selects a character by scrolling up through the character list one character at a time.
▶	Moves forward one character. If you accept an incorrect character, you can move through all 16 characters, one character at a time, until you display the incorrect character. You can then change the character.
▼	Selects a character by scrolling down through the character list one character at a time.
◀	Moves backward one character.
ESC	Returns to the default display.
ENTER	Accepts all the characters entered.

Entering the WWN

Fibre Channel protocol requires that each controller pair have a unique WWN. This 16-character alphanumeric name identifies the controller pair on the storage system. Two WWN labels attached to the rack identify the storage system WWN and checksum. See Figure 4 (page 16).

NOTE:

- The WWN is unique to a controller pair and cannot be used for any other controller pair or device anywhere on the network.
- This is the only WWN applicable to any controller installed in a specific physical location, even a replacement controller.
- Once a WWN is assigned to a controller, you cannot change the WWN while the controller is part of the same storage system.

Figure 4 Location of the World Wide Name labels

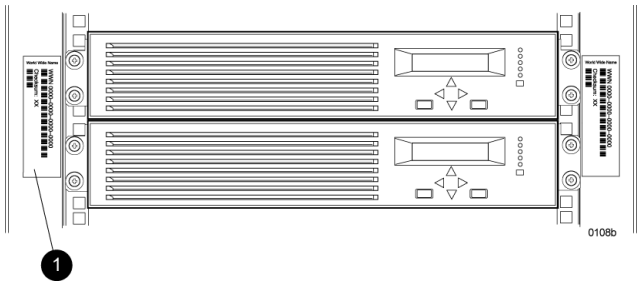

1. World Wide Name labels

Complete the following procedure to assign the WWN to each pair of controllers.

1. Turn the power switches on both controllers off.
2. Apply power to the rack.
3. Turn the power switch on both controllers on.

NOTE: Notifications of the startup test steps that have been executed are displayed while the controller is booting. It may take up to two minutes for the steps to display. The default WWN entry display has a 0 in each of the 16 positions.

4. Press ▲ or ▼ until the first character of the WWN is displayed. Press ▶ to accept this character and select the next.
5. Repeat Step 4 to enter the remaining characters.
6. Press **Enter** to accept the WWN and select the checksum entry mode.

Entering the WWN checksum

The second part of the WWN entry procedure is to enter the two-character checksum, as follows.

1. Verify that the initial WWN checksum displays 0 in both positions.
2. Press ▲ or ▼ until the first checksum character is displayed. Press ▶ to accept this character and select the second character.
3. Press ▲ or ▼ until the second character is displayed. Press **Enter** to accept the checksum and exit.
4. Verify that the default display is automatically selected. This indicates that the checksum is valid.

NOTE: If you enter an incorrect WWN or checksum, the system will reject the data and you must repeat the procedure.

Entering the storage system password

The storage system password feature enables you to restrict management access to the storage system. The password must meet the following requirements:

- 8 to 16 characters in length
- Can include upper or lower case letters
- Can include numbers 0 - 9
- Can include the following characters: ! " # $ % &amp ' ( ) * + , - . / : ; < = > ? @ [ ] ^ _ ` { | }
- Cannot include the following characters: space ~ \

NOTE: You must be running HP Command View EVA 6.0 or later to use passwords of more than eight characters. HP Command View EVA 8.0.1 is required with XCS 6.200. If you set a password longer than eight characters, you will no longer be able to manage the storage system with an earlier version of HP P6000 Command View. In this case, it will be necessary to clear the long password and reenter a password of no more than eight characters.

Complete the following procedure to enter the password:

1. Select a unique password of 8 to 16 characters.
2. With the default menu displayed, press ▼ three times to display *System Password*.
3. Press ▶ to display *Change Password?*
4. Press **Enter** for yes.

   The default password, *AAAAAAAA~~~~~~~~*, is displayed.
5. Press ▼ or ▲ to select the desired character.
6. Press ▶ to accept this character and select the next character.
7. Repeat the process to enter the remaining password characters.
8. Press **Enter** to enter the password and return to the default display.

Installing HP P6000 Command View

HP P6000 Command View is installed on a management server. Installation may be skipped if the latest version of HP P6000 Command View is running. Verify the latest version at the HP website: http://h18006.www1.hp.com/storage/software.html.

See the *HP P6000 Command View Installation Guide* for information on installing the software.

Installing optional EVA software licenses

If you purchased optional EVA software, it will be necessary to install the license. Optional software available for the Enterprise Virtual Array includes HP Business Copy EVA and HP P6000 Continuous Access. Installation instructions are included with the license.

2 Enterprise Virtual Array hardware components

The Enterprise Virtual Array includes the following hardware components:

- Fibre Channel drive enclosure — Contains disk drives, power supplies, blowers, I/O modules, and an Environmental Monitoring Unit (EMU).
- Fibre Channel loop switches — Provides twelve-port central interconnect for Fibre Channel drive enclosure FC Arbitrated Loops. The loop switches are required for EVA6000/6100 and EVA8000/8100 configurations with more than four disk enclosures.
- HSV controller — Manages all aspects of storage system operation, including communications between host systems and other devices. A pair of HSV controllers is included in the Enterprise Virtual Array.
- Rack — A variety of free-standing racks are available.

Physical layout of the storage system

The basic physical components are shown in Figure 5 (page 19). The disk drives are installed in the disk enclosures, which connect to Fibre Channel loop switches, except on the EVA4000/4100 which does not use switches. The controller pair also connects to the loop switches.

Figure 5 Storage system hardware components

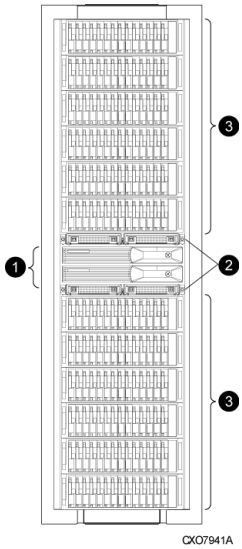

1. controllers
2. loop switches
3. disk enclosures

The EVA8000/8100, EVA6000/6100, and EVA4000/4100 are available as follows:

- EVA8000/8100 — available in multiple configurations ranging from the single-rack 2C2D configuration to the multi-rack 2C18D. The EVA8000 includes two HSV210-A controllers and four Fibre Channel loop switches. The EVA8100 includes two HSV210-B controllers and four Fibre Channel loop switches.
- EVA6000/6100 — available in configurations ranging from the 2C4D configuration to the 2C8D configuration. The EVA6000 includes two HSV200-A controllers and two Fibre Channel

loop switches. The EVA6100 includes two HSV200-B controllers with two Fibre Channel loop switches.

- EVA4000/4100 — available in configurations ranging from the 2C1D configuration to the 2C4D configuration without loop switches. The EVA4000 includes two HSV200-A controllers. The EVA4100 includes two HSV200-B controllers. Multiple EVA4000/4100s can be installed in a single rack.

See the *HP 4x00/6x00/8x00 Enterprise Virtual Array Hardware Configuration Guide* for more information about configurations. See "Related information" (page 101) for links to this document.

Fibre Channel drive enclosures

The drive enclosure contains the disk drives used for data storage. A storage system includes multiple drive enclosures. The major components of the enclosure are:

- 3U enclosure
- Dual redundant, active-to-active 2 Gbps FC loops
- 2.125-Gbps, dual loop, 14-drive enclosure
- Dual 2 Gbps FC I/O modules (A and B loops)
- Copper Fibre Channel cables
- Fibre Channel disk drives and drive blanks
- Dual redundant power supplies
- Dual redundant blowers
- Environmental Monitoring Unit (EMU)

Enclosure layout

The disk drives mount in bays in the front of the enclosure. The bays are numbered sequentially from left to right. A drive is referred to by its bay number. Enclosure status indicators are located in the lower-right, front corner. Figure 6 (page 20) shows the front and rear views of the FC drive enclosure.

Figure 6 FC drive enclosure—front and rear views

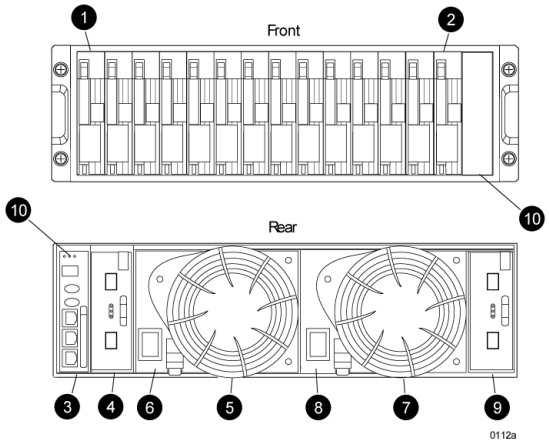

1. Drive bay 1
2. Drive bay 14
3. EMU
4. I/O module B
5. Blower 1
6. Power supply 1

7. Blower 2
8. Power supply 2
9. I/O module A
10. Status indicators (EMU, enclosure power, enclosure fault)

I/O modules

Two I/O modules provide the interface between the drive enclosure and the host controllers. See Figure 7 (page 21). They route data to and from the disk drives using Loop A and Loop B, the dual-loop configuration. For redundancy, only dual-controller, dual-loop operation is supported. Each controller is connected to both I/O modules in the drive enclosure.

Figure 7 I/O module

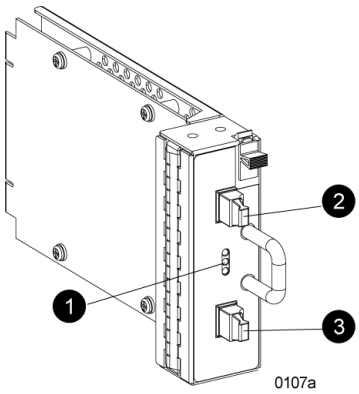

1. Status indicators (Upper port, Power, and Lower port)
2. Upper port
3. Lower port

The I/O modules are functionally identical, but are not interchangeable. Module A can only be installed at the right end of the enclosure, and module B can only be installed at the left end of the enclosure. See Figure 6 (page 20).

Each I/O module has two ports that can both transmit and receive data for bidirectional operation. Activating a port requires connecting a FC cable to the port. The port function depends upon the loop. See Figure 8 (page 21).

Figure 8 Input and output ports

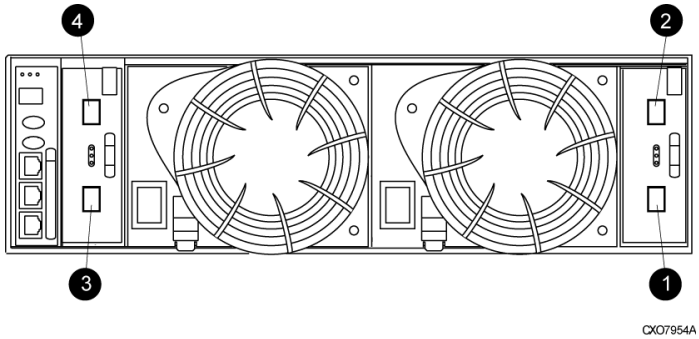

1. Loop A lower port
2. Loop A upper port
3. Loop B lower port
4. Loop B upper port

I/O module status indicators

There are three status indicators on the I/O module. See Figure 7 (page 21). The status indicator states for an operational I/O module are shown in Table 2 (page 22). Table 3 (page 22) shows the status indicator states for a non-operational I/O module.

Table 2 Operational I/O module status indicators

Upper	Power	Lower	Descriptions
Off	On	Off	• I/O Module is operational.
On	Flashing, then On	On	• Top port—Fibre Channel drive enclosure signal detected. • Power—Flashes for about 90 seconds after initial power application, then remains constant. • Bottom port—Fibre Channel drive enclosure signal detected.
On	On	On	• Top port—Fibre Channel drive enclosure signal detected. • Power—Present. • Bottom port—Fibre Channel drive enclosure signal detected.
Flashing	Flashing	Flashing	• When the locate function is active, all three indicators flash simultaneously. The Locate function overrides all other indicator functions. Therefore, an error could be detected while the Locate function is active and not be indicated until the Locate action terminates.

Table 3 Non-operational I/O module status indicators

Upper	Power	Lower	Descriptions
On	On	Off	• Top port—Fibre Channel drive enclosure signal detected. • Power—Present. • Bottom port—No Fibre Channel drive enclosure signal detected. Check transceiver and fiber cable connections. **NOTE:** This status applies to configurations with and without FC loop switches.
Off	On	On	• Top port—No Fibre Channel drive enclosure signal detected. Check transceiver and fiber cable connections. • Power—Present. • Bottom port—Fibre Channel drive enclosure signal detected .
Flashing	On	On	• Top port—EMU detected possible transceiver problem. Check transceiver and fiber cable connections. • Power—Present. • Bottom port—Fibre Channel drive enclosure signal detected .
On	On	Flashing	• Top port—Fibre Channel drive enclosure signal detected. • Power—Present. • Bottom port—EMU detected possible transceiver problem. Check transceiver and fiber cable connections. **NOTE:** The EMU will not flash the lower indicator on its own. It will flash only in response to a locate command. You can flash each of the lights independently during a locate action.
Off	Off	Off	• No I/O module power. • I/O module is nonoperational. • Check power supplies. If power supplies are operational, replace I/O module.

Fiber Optic Fibre Channel cables

The Enterprise Virtual Array uses orange, 50-µm, multi-mode, fiber optic cables for connection to the SAN. The fiber optic cable assembly consists of two 2-m fiber optic strands and small form-factor connectors on each end. See Figure 9 (page 23).

To ensure optimum operation, the fiber optic cable components require protection from contamination and mechanical hazards. Failure to provide this protection can cause degraded operation. Observe the following precautions when using fiber optic cables.

- To avoid breaking the fiber within the cable:
  - Do not kink the cable
  - Do not use a cable bend-radius of less than 30 mm (1.18 in)
- To avoid deforming, or possibly breaking the fiber within the cable, do not place heavy objects on the cable.
- To avoid contaminating the optical connectors:
  - Do not touch the connectors
  - Never leave the connectors exposed to the air
  - Install a dust cover on each transceiver and fiber cable connector when they are disconnected

If an open connector is exposed to dust, or if there is any doubt about the cleanliness of the connector, clean the connector as described in "Handling fiber optic cables" (page 71).

Figure 9 Fiber Optic Fibre Channel cable

Copper Fibre Channel cables

The Enterprise Virtual Array uses copper Fibre Channel cables to connect the drive enclosures to each other, or to the loop switches and to the HSV controllers. The cables are available in 0.6-meter and 2.0-meter lengths. Copper cables provide performance comparable to fiber optic cables. Copper cable connectors differ from fiber optic small form-factor connectors (see Figure 10 (page 23)).

Figure 10 Copper Fibre Channel cable

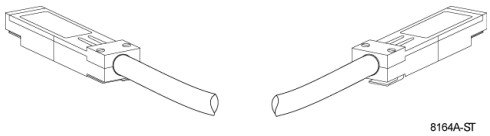

Fibre Channel disk drives

The Fibre Channel disk drives are hot-pluggable and include the following features:

- Dual-ported 2-Gbps Fibre Channel drive enclosure interface that allows up to 120 disk drives to be supported per Fibre Channel drive enclosure pair
- Compact, direct-connect design for maximum storage density and increased reliability and signal integrity
- Both online high-performance disk drives and FATA disk drives in a variety of capacities and spindle speeds
- Better vibration damping for improved performance

Up to 14 disk drives can be installed in a drive enclosure.

Disk drive status indicators

Three status indicators display the drive operational status. Figure 11 (page 24) shows the disk drive status indicators. Table 4 (page 24) provides a description of each status indicator.

Figure 11 Disk drive status indicators

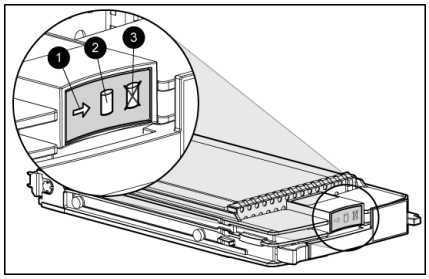

1. Activity
2. Online
3. Fault

Table 4 Disk drive status indicator descriptions

Status indicator	Description
Activity	This green status indicator flashes when the disk drive is being accessed. It is on when the drive is idle.
Online	The green status indicator is on when the disk drive is online and operating normally. This indicator will be off in the following situations: • There is no controller on the bus. • +5.1 VDC is not available. • The drive is not properly installed in the enclosure.
Fault	This amber status indicator is on when there is a disk drive failure. Depending on the host controller, this indicator may flash when the controller detects an error condition. The amber status indicator flashes in synchronization with the other two status indicators in response to the EMU locate command.

Disk drive status displays

The disk drive status indicators can assume three states: on, off, or flashing. The status indicators states for operational drive status are shown in Table 5 (page 24). See Table 6 (page 25) for the non-operational drive status indicator states.

Table 5 Operational disk drive status indications

Activity	Online	Fault	Description
Flashing	On	Off	Initial startup.
On	On	Off	The drive is online but is not being accessed.
Flashing	Flashing	Flashing	The drive is being located.
Flashing	On	Off	The drive is operational and active.

Table 6 Non-operational disk drive status indications

Activity	Online	Fault	Description
On	On	On	Indicates no connection or the controllers are offline. Recommended corrective actions: 1. Check power supplies for proper operation. 2. If defective, replace disk drive.
On	Off	Flashing	Indicates disk drive error/not active. Recommended corrective actions: 1. Verify FC loop continuity. 2. Replace disk drive.

Disk drive blank

To maintain the proper enclosure air flow, a disk drive or a disk drive blank must be installed in each drive bay. The disk drive blank maintains proper airflow within the disk enclosure.

Power supplies and blowers

This section describes the function and operation of the disk enclosure power supplies and blowers. Figure 12 (page 25) illustrates the major power supply and cooling components.

Figure 12 Power supply and blower assembly components

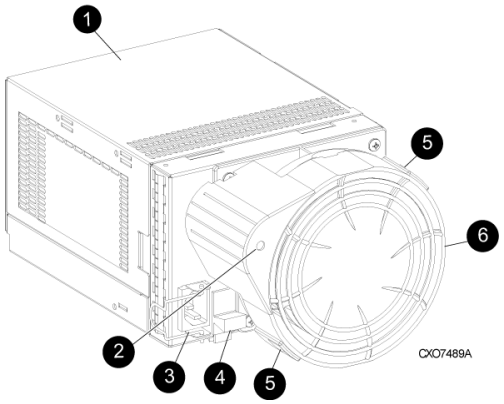

1. Power supply
2. Status indicator
3. AC Input connector with bail
4. Module latch (red wine-colored)
5. Blower tabs (red wine-colored)
6. Blower

Power supplies

The two power supplies mount in the rear of the enclosure. See Figure 23 (page 50). The supplies are autoranging and operate on a country-specific AC input voltage of 202 to 240 VAC ±10%, 50 to 60 Hz, ±5%, (188 to 264 VAC, 47 to 63 Hz).

The DC outputs of this power supply are:

- +5.1 VDC for the EMU, I/O module, backplane, and disk drives
- +12.1 VDC for the disk drives
- +12.5 VDC for the disk drives

The output of each power supply is 499 W, with a peak output of 681 W. A single power supply can support an enclosure with a full complement of disks.

The power supply circuitry provides protection against:

- Overloads
- Short circuits
- Overheating

Power supply status and diagnostic information is reported to the EMU with voltage, current, and temperature signals.

See "Regulatory notices and specifications" (page 104) for the enclosure power specifications.

The power supply temperature sensor provides a temperature range signal to the EMU, which uses this signal to set the blower speed.

The power supply internal temperature can also control the speed of the blower. The higher the power supply temperature, the faster the speed of the blower. If the power supply temperature exceeds a preset value, the power supply automatically shuts down.

Blowers

The power supply-mounted blowers cool the enclosure by circulating air through the enclosure. The blowers, under the control of the EMU or the associated power supply, can operate at multiple speeds. This ensures that, when the enclosure temperature changes, the blowers can automatically adjust the air flow.

If a blower is operating too slowly or has stopped (a "blower failure"), internal circuitry automatically operates the remaining blower at a higher speed. Simultaneously, the error condition is reported in several ways, including the power supply indicator, the audible alarm, the enclosure fault indicators, and the EMU alphanumeric display.

Should both blowers fail, the power supplies automatically shut down.

NOTE: The blowers are field-replaceable units and can be replaced, individually, while the system is running. The blowers are also interchangeable. The failure of a power supply +12.5 VDC circuit disables the associated blower.

The status indicator on the blower displays the status of both the power supply and the blower. See Figure 12 (page 25). See Table 7 (page 26) for definitions of the indicator displays.

Table 7 Power supply/blower status indicators

Blower status indicator	Description
On	Both the power supply and the blower are operational.
Flashing	The power supply or the blower locate function is active.
Off	The power supply or the blower is non-operational. When there is a blower problem, the other blower runs at a higher speed. Recommended corrective actions: • Check blower for proper operation. Replace if defective. • Check power supplies for proper operation. Replace if defective.

Drive enclosure EMU

The EMU provides increased protection against catastrophic failures. The EMU detects conditions such as failed power supplies, failed blowers, elevated temperatures, and external air sense faults and communicates these conditions to the storage system controllers.

The EMU for Fibre Channel-Arbitrated Loop (FC-AL) drive enclosures is fully compliant with SCSI-3 Enclosure Services (SES), and mounts in the left rear bay of a drive enclosure. See Figure 6 (page 20).

Controls and displays

Figure 13 (page 27) illustrates the location and function of the EMU displays, controls, and connectors.

Figure 13 EMU controls and displays

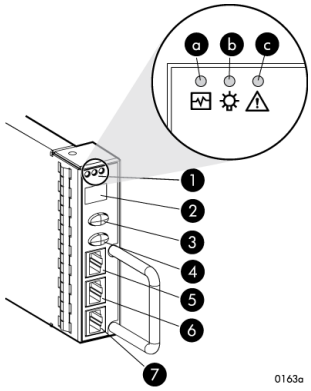

1. Status indicators:
   - **a.** EMU — This flashing green is the heartbeat for an operational EMU.
   - **b.** Enclosure power— When both the +5 VDC and +12 VDC are correct, this green indicator is on.
   - **c.** Enclosure fault — This amber indicator is normally off. The indicator is lit when an enclosure error condition exists.
2. Alphanumeric display — A two-character, seven-segment alphanumeric display of the enclosure functions and status.
3. Function select button — The primary function of this button is to select a display group function. The indicator is on when an error condition exists.
4. Display group select button — This button is used to view display groups and control the audible alarm. The indicator is on when the audible alarm is muted or disabled.
5. RS232 – For use by HP-authorized service representatives
6. LCD ONLY – Unused
7. CAB ONLY – Enclosure address bus connector

⚠ **WARNING!** To reduce the risk of electrical shock, fire, or damage to the equipment, do not plug telephone or telecommunications connectors into the RS232 ONLY receptacle.

EMU functions

The primary functions of the EMU include:

- Using the Enclosure Services Processor (ESP) to control the Enclosure Services Interface (ESI) and communicate with the controllers.
- Assigning the Enclosure Number (En), based upon the cabinet address bus feature.
- Displaying the bay 1 loop ID.
- Monitoring enclosure operation.
- Detecting, reporting, recording, and displaying conditions.
- Displaying EMU, enclosure, and element status.
- Implementing automatic corrective actions for some conditions.

- Providing enclosure status data to the controllers.
- Reporting the WWN and the logical address of all disk drives.

NOTE: Although the EMU can determine the logical address of a drive, the EMU can neither display nor change this information. HP P6000 Command View can display the addresses from the EMU-supplied status information.

EMU monitoring functions

The internal EMU circuitry monitors the enclosure and component functions listed in Table 8 (page 28).

Table 8 EMU monitoring functions

Component	Monitored Functions	
Blowers	• Installation • Removal	• Type • Speed (rpm)
Disk drives	• Installation • Removal • Bypass status	• Loop ID • Temperature • Drive fault
EMU	• Temperature • Operation	• Type • Revision level
Enclosure	• Enclosure power • Enclosure fault	• Backplane type • Backplane revision level
I/O module	• Installation • Removal • Status	• Type • Revision level
Power supplies	• Installation • Removal • Status • Type • Revision level	• +5 VDC voltage and current • +12 VDC voltage and current • Total power • Temperature
Transceiver	• Type	• Link status

EMU displays

The EMU uses a combination of status indicators, alphanumeric display, and an audible alarm to indicate the operational status of the enclosure and its components. See Table 9 (page 28).

Table 9 EMU status displays

Display	Function
Audible alarm (For information on the audible alarm, see "Audible alarm operations " (page 30).)	Any EMU-detected condition causes this alarm to sound.
Status indicators (For a description of the status indicators, see "EMU indicator displays" (page 29).) 	Display enclosure and EMU status.
Alphanumeric display (For a description of the alphanumeric display, see "Using the alphanumeric display" (page 29).)	The two-character, seven-segment display displays alphanumeric characters.

EMU indicator displays

The EMU status indicators are located above the alphanumeric display. See Figure 13 (page 27). These indicators present the same information as those on the front, lower right corner of the enclosure.

You can determine the EMU and enclosure status using the information in Table 10 (page 29).

Table 10 EMU status indications

EMU indicator (green)	Power indicator (green)	Fault indicator (amber)	Status and recommended actions
Flashing	Flashing	Flashing	The EMU Locate function is active. This display has precedence over all others. Fault conditions cannot be displayed when the Locate function is active.
Flashing	On	Off	The EMU is operational. The enclosure power (both +5 VDC and +12 VDC) is present and correct. There are no enclosure faults.
Flashing	On	On	The EMU is operational. There is an enclosure fault. Check the alphanumeric display error code for information about the problem.
Flashing	Off	Off	The EMU is operational. This display may be present when power is initially applied to the enclosure. **Note**: When the +5 VDC is incorrect, all the indicators are off.
On	On	Off	There is an EMU fault. There is no enclosure fault.
Off	On	Off	There is an EMU fault. There is no enclosure fault.
Off	Off	Off	There is an enclosure fault. Either +5 VDC is incorrect, or both +5 VDC and +12 VDC are incorrect. Other error conditions may exist.

Using the alphanumeric display

The two-character alphanumeric display is located at the top of the EMU (see Figure 13 (page 27)). This seven-segment display provides information about multiple enclosure functions. The push-button control the data displayed or entered.

Alphanumeric display description

The top-level, two-character alphanumeric display (En, Li, rG, Au, and Er), is the display group. The function of the other displays is display-group dependent. The default display is the enclosure number, a decimal number in the range 00 through 14. The push-button allow you to select the alphanumeric display or to enter data.

- The bottom push-button sequentially moves between groups and selects a display group.

  See Table 11 (page 30) for a description of these display groups.
- he top push-button moves between the levels within a display group (see 2, Figure 13 (page 27)).

Display groups

When you press and release the bottom push-button, the alphanumeric display selects a different display group. Table 11 (page 30) describes the display groups.

Table 11 EMU display groups

Display	Display group	Description
En	Enclosure Number	The enclosure number is the default display and is a decimal number in the range 00 through 14. See "Enclosure number feature" (page 32) for detailed information.
Li	Bay 1 Loop ID	This display group has a single sublevel display that defines the enclosure bay 1 loop ID. Valid loop IDs are in the range 00 through 7F.
rG	Reporting Group	This display group has two two-digit displays that define the reporting group number in the range 0000 through 4095.
Au	Audible Alarm	This display group provides control over the audible alarm or horn. The sublevel displays are audible alarm enabled (on) or audible alarm disabled (oF). See "Audible alarm operations " (page 30) for detailed information.
Fr	Firmware Revision	This display group defines the EMU code firmware version.
Er	Error Condition	This display group reads Er when there is an error condition.

NOTE: Any time you press and release the bottom pushbutton, the display will change to En, Li, rG, Au, or Er.

A flashing alphanumeric display indicates that you can edit an address or state, or view a condition report.

EMU pushbutton status indicators

The pushbutton status indicators display error conditions and the state of the audible alarm.

- When an error condition exists, the top pushbutton status indicator is On.
  - For a single error condition, the status indicator is On until the error condition is viewed.
  - For multiple errors, the status indicator is On until the last error condition is viewed.
- The bottom pushbutton indicator is On only when the alarm is muted or disabled.

Audible alarm operations

Whenever an error condition exists, the audible alarm automatically sounds until all errors are corrected. You have the option of either muting or disabling the alarm.

- Disabling the audible alarm prevents it from sounding for any error condition.
- Muting the alarm silences it for the existing condition, but any new condition causes the alarm to sound.

Audible alarm patterns

The audible alarm sound pattern differs depending on the type of error condition. See Table 12 (page 30) for the duration and the approximate relationship of these alarms. The most severe, active error condition controls the alarm pattern.

Table 12 Audible alarm sound patterns

Condition type	Cycle 1	Cycle 2
UNRECOVERABLE		
CRITICAL		

Table 12 Audible alarm sound patterns *(continued)*

Condition type	Cycle 1	Cycle 2
NONCRITICAL		
INFORMATION		
Legend	 Alarm On	Alarm Off

Controlling the audible alarm

You can control the alarm with the push-button. This process includes muting, enabling, and disabling. When an error condition exists, the alphanumeric display reads *Er*, the alarm sounds, and you can:

- Correct all errors, thereby silencing the alarm until a new error occurs.
- Mute, or temporarily disable, the alarm by pressing and holding the bottom push-button. The alarm remains off until another error occurs, or until you enable (unmute) the alarm. When a new error occurs, the alarm sounds and the push-button indicator is off.

  Using the mute feature ensures that you are aware of the more severe errors and provides you with the capability of correcting them promptly.
- Disable the alarm to prevent any error condition from sounding the alarm.

NOTE: Disabling the alarm does not prevent the EMU alphanumeric display from displaying Er, nor does it prevent HP P6000 Command View from displaying the error condition report.

When the alarm is enabled (on), the bottom push-button status indicator is off.

Enabling the audible alarm

To enable the alarm:

1. Press and release the bottom push-button until the alphanumeric display is *Au*.
2. Press and hold the top push-button until the alphanumeric display is a Flashing *oF* (Audible Alarm Off).

   NOTE: When the alarm display is flashing, press and hold the top push-button to cause the display to toggle between *On* and *oF*. Press and release the top push-button to cause the display to select the next state.
3. Press and release the top push-button to change the display to a flashing *On* (Audible Alarm On).
4. Press and release the bottom push-button to accept the change and to display *Au*. The bottom push-button indicator is now off.

Muting or unmuting the audible alarm

You may want to mute the alarm in the following situations:

- The error does not require immediate corrective action.
- You cannot correct the error at this time. For example, the error may require a replacement part.

To mute the audible alarm:

NOTE: *Er* is displayed in the alphanumeric display when an error condition is present.

1. Press and hold the bottom push-button until the status indicator is *On*.

   A muted alarm will remain off until a new condition report exists.
2. To unmute the alarm, press and hold the bottom push-button until the status indicator is *Off*. When a new error condition occurs, the alarm will sound.

Disabling the audible alarm

CAUTION: Disabling the audible alarm increases the potential of damage to equipment from a reported but unobserved fault. HP does not recommend disabling the audible alarm.

Disabling the audible alarm affects only one enclosure. This action does not affect condition report displays on the EMU alphanumeric display or errors reported by HP P6000 Command View.

To disable the alarm:

1. Press and release the bottom push-button until the alphanumeric display is *Au*.
2. Press and hold the top push-button until the alphanumeric display is a Flashing *on* (Audible Alarm On).

   NOTE: When the alarm display is flashing, pressing and holding the top push-button causes the display to rapidly change between *on* and *oF* and also causes the display to select the next state.
3. Press and release the top push-button to change the display to a Flashing *oF* (Audible Alarm Off).
4. Press and release the bottom push-button to accept the change and display *Au*. The bottom push-button indicator is now on.

   NOTE: A disabled audible alarm (the bottom push-button indicator is on) cannot sound for any error condition. To ensure that you are immediately alerted to error conditions, it is recommended that the alarm mute function be used rather than the alarm disable function. If you must use the disable function, remember to enable the audible alarm as soon as possible to ensure that you are alerted to errors.

Enclosure number feature

This section provides a description of the purpose, function, and operation of the EMU enclosure number (En) feature.

En description

In a single rack configuration, the En is a decimal number in the range 00 through 14, which is automatically assigned by the enclosure address bus.

NOTE: Your storage system may use an enclosure address bus higher than 14 if your configuration includes an optional expansion cabinet. The enclosure address bus connection determines the En. For a single rack, the display is a decimal number in the range 01 through 14. For a multiple (two) rack configuration, the display is decimal number in the range 01 through 24.

By default, the two-character alphanumeric display shows this number. Pressing the bottom push-button changes the display to *En*, the En display mode.

When the display is *En*, pressing and releasing the top push-button displays the enclosure number.

A display of 00 indicates that the enclosure is not connected to the enclosure address bus. When this condition exists, there is no EMU-to-EMU communication over the enclosure address bus.

A display of 01 through 14 indicates that the enclosure is connected to the enclosure address bus and can exchange information with other enclosures on the enclosure address bus. The decimal number indicates the physical position of the enclosure in relation to the bottom of the rack.

- 01 is the address of the enclosure connected to the lower connector in the first (lower) enclosure ID expansion cable.
- 14 is the address of the enclosure closest to end of the bus, the upper connector in the last (upper) ID expansion cable.

Unless there is an error condition, the display automatically returns to the enclosure number (01 through 14) one minute after a push-button was last pressed.

Enclosure address bus

The enclosure address bus provides a means for managing and reporting environmental conditions within the rack. It is composed of enclosure ID expansion cables that interconnect the drive enclosures and controller enclosures. Two drive enclosures connect to each enclosure ID expansion cable.

The drive enclosure numbers are always assigned by the enclosure address bus. Connecting the EMU CAB connector to an enclosure address bus enclosure ID expansion cable automatically establishes an enclosure number of 01 through 14. Any drive enclosure not connected to the enclosure address bus has the enclosure number 00.

NOTE: The enclosure number is automatically assigned. You cannot manually assign an enclosure number.

The enclosures are numbered as shown in Figure 14 (page 33).

Figure 14 Enclosure numbering with enclosure ID expansion cables

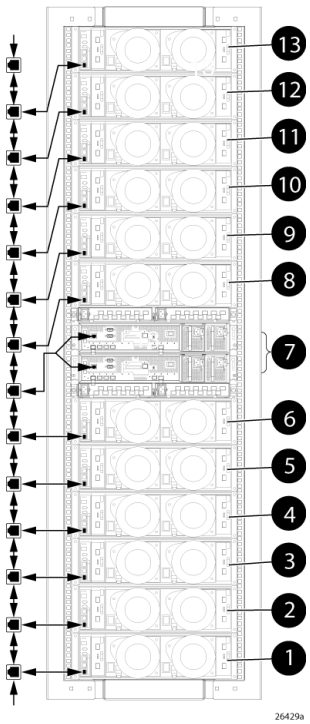

NOTE: If an expansion rack is used, the enclosure numbering shown above may change or contain additional numbering. See the *HP Enterprise Virtual Array Hardware Configuration Guide* for more information.

For more information about the reporting group number, see "Reporting group feature" (page 37).

Enclosure address bus connections

Connecting the enclosures to the enclosure ID expansion cables establishes the enclosure address bus. The enclosures are automatically numbered based on the enclosure ID expansion cable to which they are connected. Figure 15 (page 34) shows the typical configuration of a 42U cabinet with 14 enclosures.

Figure 15 Enclosure address bus components with enclosure ID expansion cables

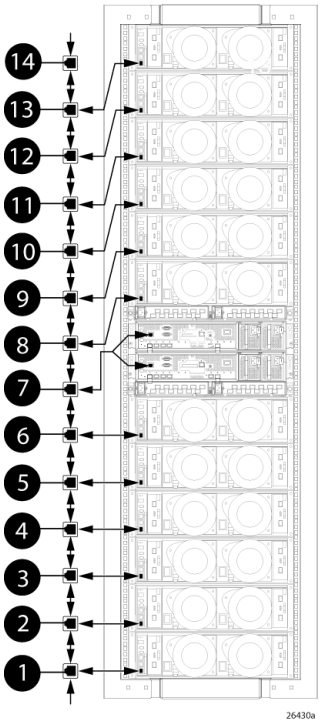

1. Shelf ID expansion cable port 1—Disk enclosure 1
2. Shelf ID expansion cable port 2—Disk enclosure 2
3. Shelf ID expansion cable port 3—Disk enclosure 3
4. Shelf ID expansion cable port 4—Disk enclosure 4
5. Shelf ID expansion cable port 5—Disk enclosure 5
6. Shelf ID expansion cable port 6—Disk enclosure 6
7. Shelf ID expansion cable port 7—Controller enclosures
8. Shelf ID expansion cable port 8—Disk enclosure 8
9. Shelf ID expansion cable port 9—Disk enclosure 9
10. Shelf ID expansion cable port 10—Disk enclosure 10
11. Shelf ID expansion cable port 11—Disk enclosure 11
12. Shelf ID expansion cable port 12—Disk enclosure 12
13. Shelf ID expansion cable port 13—Disk enclosure 13

Error Condition Reporting

The EMU constantly monitors enclosure operation and notifies you of conditions that could affect operation. When an error condition is detected, the following action is taken:

- The EMU alphanumeric display is changed to *Er*. A condition report has precedence over all other displays.
- The audible alarm sounds (if it is not disabled).
- The error is stored in the error queue.
- The error is passed to the controllers for processing and display by HP Command View EVA.

NOTE: An error always generates a condition report. Not all condition reports are generated by errors.

Error condition categories

Each error condition is assigned to a category based on its impact on disk enclosure operation. The following four error categories are used:

- Unrecoverable — the most severe error condition, it occurs when one or more enclosure components have failed and have disabled some enclosure functions. The enclosure may be incapable of correcting, or bypassing the failure, and requires repairs to correct the error.

  NOTE: To maintain data integrity, corrective action should be implemented immediately for an UNRECOVERABLE condition.

- Critical — occurs when one or more enclosure components have failed or are operating outside of their specifications. The failures impact the normal operation of some components in the enclosure. Other components within the enclosure may be able to continue normal operations. Prompt corrective action should be taken to prevent system degradation.
- Noncritical — occurs when one or more components inside the enclosure have failed or are operating outside of their specifications. The failure of these components does not impact continued normal operation of the enclosure. All components in the enclosure continue to operate according to their specifications. The ability of the components to operate correctly may be reduced should other errors occur. Prompt corrective action should be taken to prevent system degradation.
- Information — the least severe condition indicates a condition exists that does not reduce the capability of a component. However, the condition can become an error and require corrective action. An INFORMATION condition provides an early warning, which enables you to prepare to implement corrective action before a component fails. Correction of the reported problem may be delayed.

The error conditions are prioritized by severity—from most severe to least. The most severe condition takes precedence and is reported first when multiple errors are detected. The reporting characteristics for each error condition are listed in Table 13 (page 35).

Table 13 Error condition reporting characteristics

Error condition	Takes precedence over	Audible alarm pattern[1]
UNRECOVERABLE	All other conditions	On continuously
CRITICAL	NONCRITICAL and INFORMATION	Sounds three times per alarm cycle
NONCRITICAL	INFORMATION	Sounds two times per alarm cycle
INFORMATION	No other conditions	Sounds once per alarm cycle

[1] The pattern occurs when the condition is the most severe active condition.

Error queue

The EMU maintains an internal error queue for storing error conditions. Each error condition remains in the error queue until the problem is corrected, or for at least 15 seconds after the error is reported. This ensures that, when there are multiple errors or a recurring error, each can be displayed. Each entry in the error queue can be displayed using a combination of the top and bottom buttons. Each error entry in the queue contains the element type, the element number, and the error code.

Correcting the error removes the associated condition from the error queue. Replacing the EMU will also clear the error conditions. The order in which the EMU displays the error queue information is based on two factors:

- The severity of the error
- The time the error occurred

The most severe error in the queue always has precedence, regardless of how long less severe errors have been in the queue. This ensures that the most severe errors are displayed immediately.

NOTE: When viewing an error, the occurrence of a more severe error takes precedence and the display changes to the most severe error.

The earliest reported condition within an error type has precedence over errors reported later. For example, if errors at all levels have occurred, the EMU displays them in the following order:

1. UNRECOVERABLE errors in the sequence they occurred.
2. CRITICAL errors in the sequence they occurred.
3. NONCRITICAL errors in the sequence they occurred.
4. INFORMATION conditions in the sequence they occurred.

Error condition report format

Each EMU detected condition generates a condition report containing three pieces of information.

- **Element type** The first two-digit hexadecimal display defines the element type reporting the problem. The format for this display is *e.t.* with a period after each character. Valid element types are 0.1. through F.F.
- **Element number** The second display is a two-digit decimal number that defines the specific element reporting the problem. The format for this display is *en.* with a period after the second character.
- **Error code** The third display is a two-digit decimal number that defines the specific error code. The format for this display is *ec* without any periods.

For detailed information about each condition report, including recommended corrective actions, see "EMU-generated condition reports" (page 117).

Navigating the error condition display

When an error condition occurs, the alphanumeric display changes to *Er* and the error menu is active. The buttons are used to display the error condition values.

Perform the following procedure to display error conditions. Figure 16 (page 37) illustrates the sequence for displaying error conditions.

1. With *Er* in the display, press and hold the top push-button until the first element type is displayed. The most severe error in the queue will be displayed.
2. Release the top push-button when the element type is displayed. The element type has both decimal points lit.
3. Press and release the top push-button to display the element number. This display has only the right decimal point lit.
4. Press and release the top push-button again to display the error code. This display has no decimal points lit. Repeated press/release operations will cycle through these three values.
5. Press and hold the top push-button from any of the three display states to move to the element type for the next error condition in the queue.
6. Use the top push-button to display the values for the error condition.
7. When all error conditions have been displayed, press and release the bottom pushbutton to return to the Er display.

Figure 16 Displaying error condition values

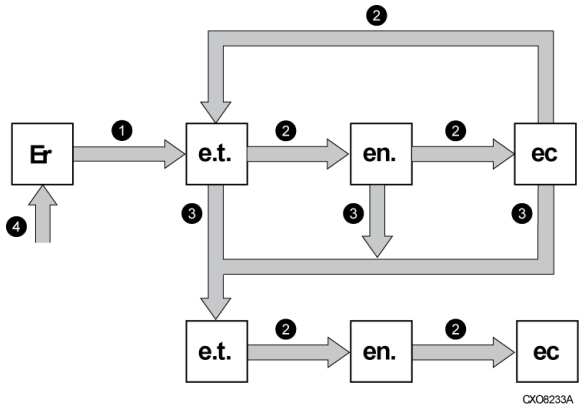

1 Press and hold top push-button to view first error in queue.

2 Press and release top push-button.

3 Press and hold top push-button to view next error.

4 Press and release the bottom push-button at any time to return to the Er display.

e.t. = element type, **en.** = element number, **ec** = error code

Analyzing condition reports

Analyzing each error condition report involves three steps:

1. Identifying the element.
2. Determining the major problem.
3. Defining additional problem information.

Reporting group feature

Another function of the enclosure address bus is to provide communications within a reporting group. A Reporting Group (rG) is an HSV controller pair and the associated drive enclosures. The controller pair automatically assigns a unique (decimal) four-digit Reporting Group Number (RGN) to each EMU on a Fibre Channel drive enclosure.

Each of the drive enclosures on a loop pair are in one reporting group:

- All of the drive enclosures on loop pair 1, both loop 1A and loop 1B, share a unique reporting group number.
- All of the drive enclosures on loop pair 2, both loop 2A and loop 2B, share a unique reporting group number.

Each EMU collects environmental information from the associated enclosure and broadcasts the information to reporting group members using the enclosure address bus. Information from enclosures in other reporting groups is ignored.

Reporting group numbers

The reporting group number (RGN) range is 0000 through 4099, decimal.

- 0000 is reserved for enclosures that are not part of any reporting group.
- 0001 through 0015 are RGNs reserved for use by the EMU.
- 0016 through 4095 are valid RGNs.
- 4096 through 4099 are invalid RGNs.

The reporting group numbers are displayed on the EMU alphanumeric display as a pair of two-digit displays. These two displays are identified as *rH* and *rL*.

- Valid *rH* displays are in the range 00 through 40, and represent the high-order (most significant) two digits of the RGN.
- Valid *rL* displays are in the range 00 through 99, and represent the low-order (least significant) two digits of the RGN.

To view a reporting group number:

1. Press and release the bottom push-button until the alphanumeric display is *rG*.
2. To display the two most significant digits of the Reporting Group Number, press and hold the top pushbutton unit the display is *rH*.
3. Press and release the top push-button to display the first two digits of the RGN.
4. Press and release the top push-button until the alphanumeric display is *rH*.
5. Press and hold the top push-button until the alphanumeric display is *rL*.
6. Press and release the top push-button to display the last two digits of the Reporting Group Number.
7. To exit the display, press and release the bottom push-button until the alphanumeric display is *rG*.

Fibre Channel loop switches

The loop switches act as a central point of interconnection and establish a fault-tolerant physical loop topology between the controllers and the disk enclosures. The loop switches are required in any configuration with more than four disk enclosures. The EVA8000/8100 uses four loop switches and the EVA6000/6100 uses two switches to connect the drive enclosures to the controller pair.

The loop switches provide the following features:

- 2.125-Gbps operating speed
- Twelve ports
- Half-width, 1U size
- System and port status indicators
- Universal power supply that operates between 100 to 250 VAC and 50 to 60 Hz

NOTE: Each bezel covers two FC loop switches in a space of 1U.

The EVA8000 uses four loop switches to connect all of the drive enclosures to the controller pair using FC cables. The EVA 6000 includes two loop switches. Each switch acts as a central point of interconnection and establishes a fault-tolerant physical loop topology. The EVA6100/8100 use the 30-10022-01 loop switch only.

The half-rack form factor switch is controlled by firmware loaded into the on-board flash memory. The switch is designed as a central interconnect following the ANSI FC-AL standard. Disk enclosures are connected to the switch through Small Form-factor Pluggable (SFP) transceivers and cables.

The storage system uses one of the following loop switches:

- 30-10022-01 loop switch–used with 2 Gb and 4 Gb controllers (requires XCS 5.110 or later)
- 30-10010-02 loop switch–used with 2 Gb and 4Gb controllers

30-10022-01 loop switch

The 30-10022-01 loop switch contains both system indicators and port indicators. The system indicators indicate the status of the switch, and the port indicators provide status of a specific port. Figure 17 (page 39) shows the 30-10022-01 loop switch with the system and port indicators.

Figure 17 30-10022-01 loop switch status indicators

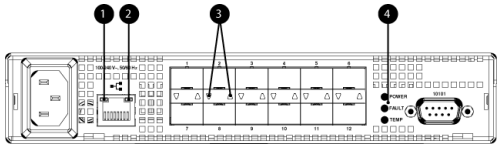

1. Ethernet activity	• Flashing—the Ethernet port is receiving data. • Flashing rapidly—the traffic level is high.
2. Ethernet link	• On—the port is connected to an operational Ethernet.
3. Port status	• Off—SFP is not installed in the port. • On (green)—Normal port operational status when an SFP is installed and a link has been established. • On (yellow)—port has an SFP installed but a link has not been established. • Flashing (green)—activity. Data is being transferred between the port and device.
4. System status	
Power	• On—the switch is plugged in and the internal power supply is functional.
Fault	• On—an event has occurred that meets or exceeds the current Fault threshold setting. The default Fault threshold setting is critical. The switch will continue to operate. Switch functionality may be impaired depending on the event that triggered the Fault condition. Regardless of the cause, the switch requires immediate attention.
Temp	• On—the internal temperature has exceeded acceptable levels. The switch will continue to operate. Switch functionality may be impaired depending on the event that triggered the Temp condition. Regardless of the cause, the switch requires immediate attention.

Power-on self test (POST)

When the switch powers on, it runs Power-On Self-Test (POST) diagnostics to verify the fundamental integrity of the switch ports. All switch LEDs turn on (LEDs illuminate). Then, excluding the Ethernet Link and Power LEDs, the LEDs turn off (LEDs extinguish). Once the switch is operational, the LEDs display current status. See Figure 17 (page 39).

30-10010-02 loop switch

The 30-10010-02 loop switch contains both system indicators and port indicators. The system indicators indicate the status of the switch, and the port indicators provide status of a specific port. Figure 18 (page 40) shows the 30-10010-02 loop switch with the system and port indicators.

The Fibre Channel loop switch acts as a central point of interconnection and establishes a fault-tolerant physical loop topology between the controllers and the disk enclosures.

Figure 18 30-10010-02 loop switch status indicators

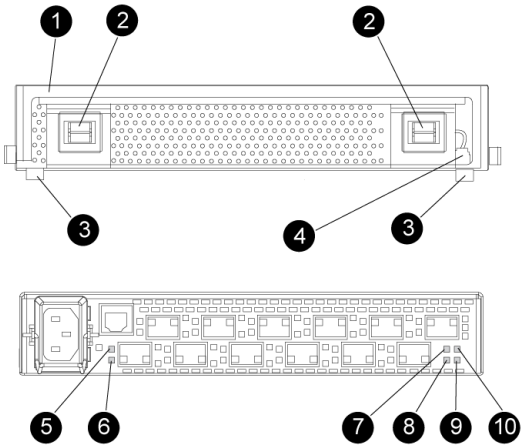

1. Handle
2. Bezel snaps
3. Alignment tabs
4. Walk-up RS232 port
5. SFP status indicator
6. Port Bypassed indicator
7. POST fault indicator
8. Over Temp indicator
9. Power indicator
10. Loop operational indicator

Power-on self test (POST)

When you power on the 30-10010-02 loop switch, it performs a Power-on Self Test (POST) to verify that the switch is functioning properly. During a POST, all of the indicators turn on for approximately two seconds. Then, turn off all of the indicators, except the power indicator.

If the Port Bypass indicators are blinking at a constant rate and the POST Fault indicator is on, the switch detected a fault during the POST. In this case, you need to contact your HP authorized service representative.

Reading the switch status indicators

Figure 18 (page 40) shows the Fibre Channel switch with the system and port indicators. Table 14 (page 40) lists and describes the system indicators.

Table 14 30-10010-02 loop switch status indicators

System indicator	Description
Power	A green indicator. When lit, this indicates that the switch is plugged in and the internal power is functional.
Loop operational	A green indicator. When lit, this indicates that the Fibre Channel loop has completed initialization and is now operational.
POST fault	An amber indicator. When lit, this indicates that the internal hardware self-test failed and the switch will not function.
OverTemp	An amber indicator. When lit, this indicates that the ambient temperature has exceeded 40° C. The switch is still functional; however, you should correct the problem immediately. The OverTemp indicator turns off when the problem is corrected.

Table 15 (page 41) describes the port indicators.

Table 15 30-10010-02 loop switch port status indicators

SFP status indicator (Green)	Port bypass indicator (Amber)	Description
Off	Off	Indicates that the port does not have an SFP installed and is bypassed by the loop.
On	Off	Indicates that the port is operating normally. The port and device are fully operational.
On	On	Indicates the that port is in a bypassed state. The port is non-operational due to loss of signal, poor signal integrity, or the Loop Initialization Procedure (LIP). **NOTE:** This condition is also normal when the SFP is present but not attached to a Fibre Channel drive enclosure node, or when the SFP is present and attached to only a cable assembly. Attaching the SFP to a device and plugging it into the port should initiate the LIP by the attached device.
Off	On	Indicates a Tx fault. The port is non-operational due to an SFP transmitter fault, improperly-seated SFP, or another failed device.

Problem isolation

Table 16 (page 41) lists several basic problems and their solutions.

Table 16 30-10010-02 loop switch basic troubleshooting

Problem	Recommended action
SFPs are installed in ports but no indicators are lit.	1. Verify that the power cord is firmly seated in the switch and is connected to the power outlet. 2. Check the power indicator to verify that the switch is on.
SFP is installed, but the Port Bypassed indicator is lit.	Re-seat the SFP. If the same condition occurs, the SFP is probably faulty and should be replaced.
SFP is installed, but the SFP status indicator and the Port Bypassed indicator are lit.	This condition indicates that the switch is not receiving a valid Fibre Channel signal or that the switch is receiving an LIP. 1. Ensure that the switch is powered on. 2. Contact your HP authorized service representative for further assistance.
SFP is installed and the SFP status indicator is lit, but the devices are not communicating.	This condition indicates that the switch is receiving a valid Fibre Channel signal, but there are no upper level protocols active. 1. Verify that you are running the correct firmware on all storage system hardware. 2. Check the Loop Operational indicator. a. If the Loop Operational indicator is lit, the devices have completed initialization. b. If the Loop Operational indicator is off, the devices were not initialized. Disconnect the devices from the switch. Reconnect the devices one at a time. This allows you to isolate the device that is responsible for the loop failure. 3. Contact your authorized service representative for further assistance.

HSV controllers

Two controllers (HSV210-A/B or HSV200-A/B) are contained in each rack. Each controller is contained in a separate controller and provides the following features:

- High-performance microprocessor
- An Operator Control Panel (OCP)

- Four 4 Gbps Fibre Channel-Switched fabric host ports (two host ports in HSV200-A or HSV200-B controller)
- Four 2 Gbps Fibre Channel drive enclosure device ports (two device ports in HSV200-A or HSV200- B controller)
  - Arranged in redundant pairs
  - Data load/performance balanced across a pair
  - Support for up to 240 disks with HSV210-A or HSV210-B and 112 with HSV200-A or HSV200- B
- 2 GB cache per controller, mirrored, with battery backup (1-GB cache in HSV200-A or HSV200- B controller)
- 2 Gbps FC cache mirroring ports with device port backups
- Dual power supplies

In addition to managing the operation of the storage system, the HSV controllers serve as the interface between the storage system hardware and the SAN. All host I/Os and all HP P6000 Command View management commands are processed by the controllers. Up to 18 drive enclosures are supported by one controller pair.

Figure 19 (page 42) shows the HSV210-A/B controller rear view. Figure 20 (page 43) shows the HSV200-A/B controller rear view. The front view of both controllers is shown in Figure 21 (page 43).

NOTE: The EVA4000/6000/8000 and EVA4100/6100/8100 use controllers with 2 Gb and 4 Gb host port capability. The 4 Gb controller can be distinguished from the earlier 2 Gb controllers by the "-A" and "-B" suffixes used on the controller bezel. The 4 Gb EVA4000 and EVA6000 controllers are identified as the HSV200-A. The 4 Gb EVA4100 and EVA6100 are identified as HSV200-B. The 4 Gb EVA8000 is identified as the HSV210-A and the 4 Gb EVA8100 is identified as HSV210-B.

IMPORTANT: To upgrade from an HSV200-A or HSV210-A controller to an HSV200-B or HSV210-B controller, HP requires that you also upgrade the I/O modules (A and B) to AD623C and AD624C on each shelf. If you are upgrading to an EVA6100 (HSV200-B) or EVA8100 (HSV210-B) and you do not already have the 30-10022-01 loop switches installed, you must also upgrade the loop switches to 30-10022-01.

Figure 19 HSV210-A/B controller—rear view

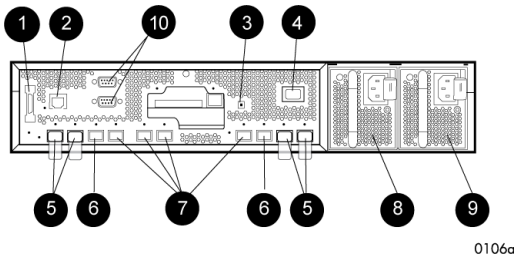

1. Dual controller interconnect
2. CAB (cabinet address bus)
3. Unit ID
4. Power ON
5. FC device ports
6. FC cache mirror ports
7. FC host ports
8. Power supply 0
9. Power supply 1
10. Service connectors (not for customer use)

Figure 20 HSV200-A/B controller—rear view

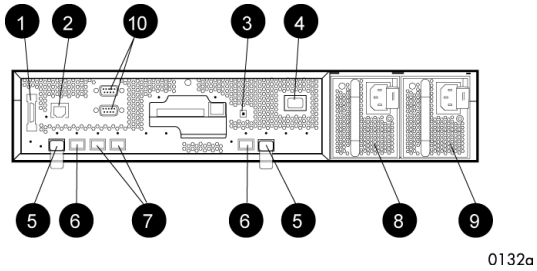

1. Dual controller interconnect
2. CAB (cabinet address bus)
3. Unit ID
4. Power ON
5. FC device ports
6. FC cache mirror ports
7. FC host ports
8. Power supply 0
9. Power supply 1
10. Service connectors (not for customer use)

Figure 21 HSV controller—front view

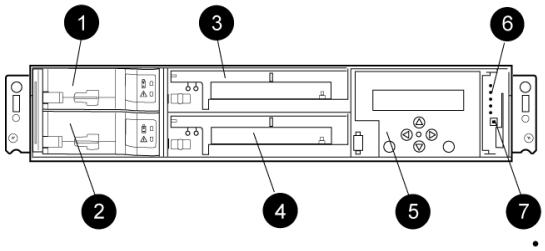

1. Battery 0
2. Battery 1 (EVA8000/8100 only)
3. Blower 0
4. Blower 1
5. Operator Control Panel (OCP)
6. Status indicators
7. Unit ID

High availability features

Two interconnected controllers ensure that the failure of a controller component does not disable the system. A single controller can fully support an entire system until the defective controller, or controller component, is repaired. For EVA8x00 configurations with more than four disk drive enclosures, the complete data redundancy configuration includes device loop switches on the two Loop A and two Loop B data paths. For EVA4x00 and EVA6x00 configurations, data redundancy is accomplished with device loop switches on the two Loop A data paths.

Each HSV210-A/B controller has two lead-acid cache batteries that provide power to the cache memory. Each HSV200-A/B controller has one battery. When the batteries are fully charged, they can provide power to the cache for up to 96 hours.

Operator control panel

The operator control panel (OCP) provides a direct interface to each controller. From the OCP you can display storage system status and configuration information, shut down the storage system, and manage the password.

The OCP includes a 40-character LCD alphanumeric display, six push-buttons, and five status indicators. See Figure 22 (page 44).

HP Command View EVA is the tool you will typically use to display storage system status and configuration information or perform the tasks available from the OCP. However, if HP P6000 Command View is not available, the OCP can be used to perform these tasks.

Figure 22 Controller OCP

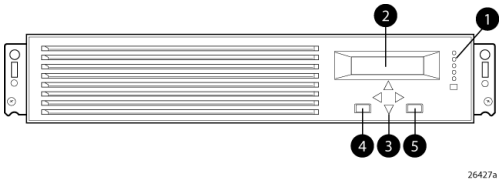

1. Status indicators (see Table 17 (page 44)) and UID button
2. 40-character alphanumeric display
3. Left, right, top, and bottom push-buttons
4. Esc
5. Enter

Status indicators

The status indicators display the operational status of the controller. The function of each indicator is described in Table 17 (page 44). During initial setup, the status indicators might not be fully operational.

The following sections define the alphanumeric display modes, including the possible displays, the valid status indicator displays, and the pushbutton functions.

Table 17 Controller status indicators

Indicator	Description
Fault	When this indicator is on, there is a controller problem. Check either HP P6000 Command View or the LCD Fault Management displays for a definition of the problem and recommended corrective action.
Controller	When this indicator is flashing slowly, a heartbeat, the controller is operating normally. When this indicator is not flashing, there is a problem.
Physical link to hosts established	When this indicator is green, there is at least one physical link between the storage system and hosts that is active and functioning normally. When this indicator is off, there are no links between the storage system and hosts that are active and functioning normally.
Virtual disks presented to hosts	When this indicator is green, all virtual disks that are presented to hosts are healthy and functioning normally. When this indicator is amber, at least one virtual disk is not functioning normally. When this indicator is off, there are no virtual disks presented to hosts and this indicates a problem with the Vdisk on the array.
Cache Battery Assembly	When this indicator is off, the battery assembly is charged. When this indicator is on, the battery assembly is discharged.
Unit ID	Press to light the blue LED on the front and back of the controller. This indicator comes on in response to a Locate command issued by HP P6000 Command View.

Each port on the rear of the controller has an associated status indicator located directly above it. Table 18 (page 45) lists the port and its status description.

Table 18 Controller port status indicators

Port	Description
Fibre Channel host ports	• Green—Normal operation • Amber—No signal detected • Off—No SFP[1] detected or the Direct Connect OCP setting is incorrect
Fibre Channel device ports	• Green—Normal operation • Amber—No signal detected or the controller has failed the port • Off—No SFP [1] detected
Fibre Channel cache mirror ports	• Green—Normal operation • Amber—No signal detected or the controller has failed the port • Off—No SFP[1] detected
Dual controller interconnect port	• Green—Normal operation • Amber—Interconnect cable not connected

[1] On copper Fibre Channel cables, the SFP is integrated into the cable connector.

Navigation buttons

The operation of the navigation buttons is determined by the current display and location in the menu structure. Table 19 (page 45) defines the basic pushbutton functions when navigating the menus and options.

To simplify presentation and to avoid confusion, the pushbutton reference names, regardless of labels, are left, right, top, and bottom.

Table 19 Navigation button functions

Button	Function
▼	Moves down through the available menus and options
▲	Moves up through the available menus and options
▶	Selects the displayed menu or option.
◀	Returns to the previous menu.
Esc	Used for “No” selections and to return to the default display.
Enter	Used for “Yes” selections and to progress through menu items.

Alphanumeric display

The alphanumeric display uses two LCD rows, each capable of displaying up to 20 alphanumeric characters. By default, the alphanumeric display alternates between displaying the Storage System Name and the World Wide Name. An active (flashing) display, an error condition message, or a user entry (pressing a push-button) overrides the default display. When none of these conditions exist, the default display returns after approximately 10 seconds.

Displaying the OCP menu tree

The Storage System Menu Tree lets you select information to be displayed, configuration settings to change, or procedures to implement. To enter the menu tree, press any navigation push-button when the default display is active.

The menu tree is organized into the following major menus:

- System Info—displays information and configuration settings.
- Fault Management—displays fault information. Information about the Fault Management menu is included in "Controller fault management" (page 135).
- Shutdown Options—initiates the procedure for shutting down the system in a logical, sequential manner. Using the shutdown procedures maintains data integrity and avoids the possibility of losing or corrupting data.
- System Password—create a system password to ensure that only authorized personnel can manage the storage system using HP P6000 Command View.

To enter and navigate the storage system menu tree:

1. Press any push-button while the default display is in view. System Information becomes the active display.
2. Press ▼ to sequence down through the menus.

   Press ▲ to sequence up through the menus.

   Press ▶ to select the displayed menu.

   Press ◀ to return to the previous menu.

NOTE: To exit any menu, press **Esc** or wait ten seconds for the OCP display to return to the default display.

Table 20 (page 46) identifies all the menu options available within the OCP display.

△ **CAUTION:** Many of the configuration settings available through the OCP impact the operating characteristics of the storage system. You should not change any setting unless you understand how it will impact system operation. For more information on the OCP settings, contact your HP-authorized service representative.

Table 20 Menu options within the OCP display

System Information	Fault Management	Shutdown Options	System Password
Versions	Last Fault	Restart	Change Password
Host Port Config (Sets Fabric or Direct Connect)	Detail View	Power Off	Clear Password
Device Port Config (Enables/disables device ports)		Uninitialize System	Current Password (Set or not)
IO Module Config (Enables/disables auto-bypass)			
Loop Recovery Config (Enables/disables recoveries)			
Unbypass Devices			
UUID Unique Half			
Debug Flags			
Print Flags			
Mastership Status (Displays controller role — master or slave)			

Displaying system information

NOTE: The purpose of this information is to assist the HP-authorized service representative when servicing your system.

The system information displays show the system configuration, including the XCS version, the OCP firmware and application programming interface (API) versions, and the enclosure address bus programmable integrated circuit (PIC) configuration. You can only view, not change, this information.

Displaying versions system information

When you press ▼, the active display is Versions. From the Versions display you can determine the:

- OCP firmware version
- Controller version
- XCS version

NOTE: The terms PPC, Sprite, Glue, SDC, CBIC, and Atlantis are for development purposes and have no significance for normal operation.

NOTE: When viewing the software or firmware version information, pressing ◀ displays the Versions Menu tree.

To display System Information:

1. The default display alternates between the Storage System Name display and the World Wide Name display.

   Press any push-button to display the Storage System Menu Tree.
2. Press ▼ until the desired Versions Menu option appears, and then press ▼ or ▶ to move to submenu items.

Shutting down the system

CAUTION: To power off the system for more than 96 hours, use HP P6000 Command View.

You can use the Shutdown System function to implement the shutdown methods listed below. These shutdown methods are explained in Table 21 (page 47).

- Shutting down the controller (see "Shutting the controller down" (page 48)).
- Restarting the system (see "Restarting the system" (page 48)).
- Uninitializing the system (see "Uninitializing the system" (page 48)).

To ensure that you do not mistakenly activate a shutdown procedure, the default state is always NO, indicating *do not implement this procedure*. As a safeguard, implementing any shutdown method requires you to complete at least two actions.

Table 21 Shutdown methods

LCD prompt	Description
Restart System?	Implementing this procedure establishes communications between the storage system and HP P6000 Command View. This procedure is used to restore the controller to an operational state where it can communicate with HP P6000 Command View.
Power off system?	Implementing this procedure initiates the sequential removal of controller power. This ensures no data is lost. The reasons for implementing this procedure include replacing a drive enclosure.
Uninitialize?	Implementing this procedure will cause the loss of all data. For a detailed discussion of this procedure, see "Uninitializing the system" (page 48).

Shutting the controller down

Use the following procedure to access the *Shutdown System* display and execute a shutdown procedure.

△ **CAUTION:** If you decide NOT to power off while working in the Power Off menu, *Power Off System NO* must be displayed before you press **Esc**. This reduces the risk of accidentally powering down.

NOTE: HP P6000 Command View offers the preferred method for shutting down the controller. Shut down the controller from the OCP only if HP P6000 Command View cannot communicate with the controller.

Shutting down the controller from the OCP removes power from the controller on which the procedure is performed only. To restore power, toggle the controller's power.

1. Press ▼ three times to scroll to the Shutdown Options menu.
2. Press ▶ to display *Restart*.
3. Press ▼ to scroll to **Power Off**.
4. Press ▶ to select **Power Off**.
5. *Power off system* is displayed. Press **Enter** to power off the system.

Restarting the system

To restore the controller to an operational state, use the following procedure to restart the system.

1. Press ▼ three times to scroll to the Shutdown Options menu.
2. Press ▶ to select **Restart**.
3. Press ▶ to display *Restart system?*.
4. Press **Enter** to go to Startup.

   No user input is required. The system will automatically initiate the startup procedure and proceed to load the Storage System Name and World Wide Name information from the operational controller.

Uninitializing the system

Uninitializing the system is another way to shut down the system. This action causes the loss of all storage system data. Because HP P6000 Command View cannot communicate with the disk drive enclosures, the stored data cannot be accessed.

△ **CAUTION:** Uninitializing the system destroys all user data. The WWN will remain in the controller unless both controllers are powered off. The password will be lost. If the controllers remain powered on until you create another storage system (initialize via GUI), you will not have to re-enter the WWN.

Use the following procedure to uninitialize the system.

1. Press ▼ three times to scroll to the Shutdown Options menu.
2. Press ▶ to display *Restart*.
3. Press ▼ twice to display *Uninitialize System*.
4. Press ▶ to display *Uninitialize?*
5. Select **Yes** and press **Enter**.

   The system displays `Delete all data? Enter DELETE:_____`

6. Press the arrow keys to navigate to the open field and type **DELETE** and then press **ENTER**.

   The system uninitializes.

   NOTE: If you do not enter the word DELETE or if you press **ESC**, the system does not uninitialize. The bottom OCP line displays `Uninit cancelled`.

Password options

The password entry options are:

- Entering a password during storage system initialization (see "Entering the storage system password" (page 17)).
- Displaying the current password.
- Changing a password (see "Changing a password" (page 49)).
- Removing password protection (see "Clearing a password" (page 49)).

Changing a password

For security reasons, you may need to change a storage system password. The password must contain eight to 16 characters consisting of any combination of alpha, numeric, or special. See "Entering the storage system password" (page 17) for more information on valid password characters.

Use the following procedure to change the password.

NOTE: Changing a system password on the controller requires changing the password on any HP P6000 Command View with access to the storage system.

1. Select a unique password of 8 to 16 characters.
2. With the default menu displayed, press ▼ three times to display *System Password*.
3. Press ▶ to display *Change Password?*
4. Press **Enter** for yes.

   The default password, *AAAAAAAA*~~~~~~~~, is displayed.
5. Press ▼ or ▲ to select the desired character.
6. Press ▶ to accept this character and select the next character.
7. Repeat the process to enter the remaining password characters.
8. Press **Enter** to enter the password and return to the default display.

Clearing a password

Use the following procedure to remove storage system password protection.

NOTE: Changing a system password on the controller requires changing the password on any HP P6000 Command View with access to the storage system.

1. Press ▼ four times to scroll to the System Password menu.
2. Press ▶ to display *Change Password?*
3. Press ▼ to scroll to *Clear Password*.
4. Press ▶ to display *Clear Password*.
5. Press **Enter** to clear the password.

   The *Password cleared* message will be displayed.

Power supplies

Two power supplies provide the necessary operating voltages to all controller enclosure components. If one power supply fails, the remaining supply is capable of operating the enclosure.

Figure 23 Power supplies

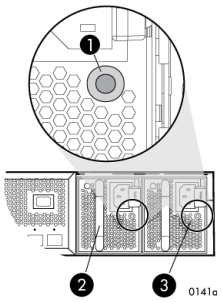

1. Status indicator
2. Power supply 0
3. Power supply 1

Table 22 Power supply status indicators

Status indicator	Description
Green	Normal operation
Amber	Power supply failure
Flashing amber	The power cord is disconnected from the power supply

Blowers

Two blowers provide the cooling necessary to maintain the proper operating temperature within the controller enclosure. If one blower fails, the remaining blower is capable of cooling the enclosure.

Figure 24 Blower

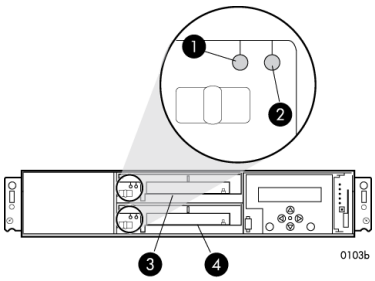

1. Status indicator
2. Fault indicator
3. Blower 0
4. Blower 1

Table 23 Blower status indicators

Status indicator	Fault indicator	Description
On	Off	Normal operation
Off	On	Blower failure

Cache battery

Batteries provide backup power to maintain the contents of the controller cache when AC power is lost and the storage system has not been shutdown properly. When fully charged the batteries can sustain the cache contents for to 96 hours. Two batteries are used on the EVA8x00 and a single battery is used on the EVA6x00 and EVA4x00. Figure 25 (page 51) illustrates the location of the cache batteries and the battery status indicators. See Table 24 (page 51) for additional information on the status indicators.

Figure 25 Cache batteries

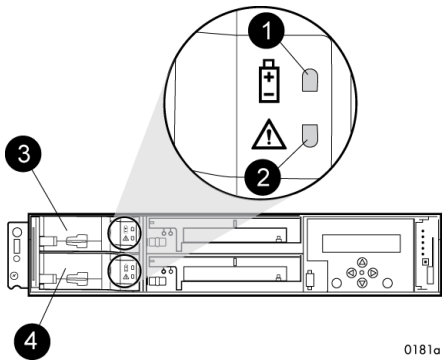

1. Status indicator
2. Fault indicator
3. Battery 0
4. Battery 1

The table below describes the battery status indicators. When a battery is first installed, the fault indicator goes on (solid) for approximately 30 seconds while the system discovers the new battery. Then, the battery status indicators display the battery status as described in the table below.

Table 24 Battery status indicators

Status indicator	Fault indicator	Description
On	Off	Normal operation. A maintenance charge process keeps the battery fully charged.
Flashing	Off	Battery is undergoing a full charging process. This is the indication you typically see after installing a new battery.
Off	On	Battery fault. The battery has failed and should be replaced.
Off	Flashing	The battery has experienced an over temperature fault.
Flashing (fast)	Flashing (fast)	Battery code is being updated. When a new battery is installed, it may be necessary for the controllers to update the code on the battery to the correct version. Both indicators flash rapidly for approximately 30 seconds.
Flashing	Flashing	Battery is undergoing a scheduled battery load test, during which the battery is discharged and then recharged to ensure it is working properly. During the discharge cycle, you will see this display. The load test occurs infrequently and takes several hours.

HSV controller cabling

All data cables and power cables attach to the rear of the controller. Adjacent to each data connector is a two-colored link status indicator. Table 18 (page 45) identifies the status conditions presented by these indicators.

NOTE: These indicators do not indicate whether there is communication on the link, only whether the link can transmit and receive data.

The data connections are the interfaces to the disk drive enclosures or loop switches (depending on your configuration), the other controller, and the fabric. Fiber optic cables link the controllers to the fabric, and, if an expansion cabinet is part of the configuration, link the expansion cabinet drive enclosures to the loop es in the main cabinet. Copper cables are used between the controllers (mirror port) and between the controllers and the drive enclosures or loop switches.

Racks

All storage system components are mounted in a rack. The rack provides the capability for mounting standard 483 mm (19 in) wide controller and drive enclosures. Each configuration includes two controller enclosures (the controller pair), drive enclosures, FC loop switches (if required), and an expansion bulkhead. Each controller pair and all the associated drive enclosures form a single storage system.

The following racks are available for the EVA8000/8100:

- 36U Rack
- 42U Rack

NOTE:

- Although the 22U, 25U, 33U, and 41U rack configurations are no longer available, existing storage systems in these racks are still supported.
- Racks and rack-mountable components are typically described using "U" measurements. "U" measurements are used to designate panel or enclosure heights.

The racks provide the following:

- Unique frame and rail design — Allows fast assembly, easy mounting, and outstanding structural integrity.
- Thermal integrity — Front-to-back natural convection cooling is greatly enhanced by the innovative multi-angled design of the front door.
- Security provisions — The front and rear door are lockable, which prevents unauthorized entry.
- Flexibility — Provides easy access to hardware components for operation monitoring.
- Custom expandability — Several options allow for quick and easy expansion of the racks to create a custom solution.

Rack configurations

Each system configuration depends on the number of disk enclosures included in the storage system. For more information about racks and configurations, including expansion and interconnection, see the *HP Enterprise Virtual Array Hardware Configuration Guide*.

Power distribution

AC power is distributed to the rack through a dual Power Distribution Unit (PDU) assembly mounted at the bottom rear of the rack. The characteristics of the fully-redundant rack power configuration are as follows:

- Each PDU is connected to a separate circuit breaker-protected, 30-A AC site power source (220–240 VAC ±10%, 50 or 60-Hz, ±5%). Figure 26 (page 53) illustrates the compatible 60-Hz and 50-Hz wall receptacles.

Figure 26 60-Hz and 50-Hz wall receptacles

NEMA L6-30R receptacle, 3-wire, 30-A, 60-Hz

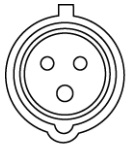

IEC 309 receptacle, 3-wire, 30-A, 50-Hz

- The standard power configuration for any Enterprise Virtual Array rack is the fully redundant configuration. Implementing this configuration requires:
  - Two separate circuit breaker-protected, 30-A site power sources with a compatible wall receptacle (see Figure 26 (page 53)).
  - One dual PDU assembly. Each PDU connects to a different wall receptacle.
  - Six Power Distribution Modules (PDM) per rack. Three PDMs mount vertically on each side of the rack. Each set of PDMs connects to a different PDU.
  - The drive enclosure power supplies on the left (PS 1) connect to the PDMs on the left with a gray, 66 cm (26 in) power cord.
  - The drive enclosure power supplies on the right (PS 2) connect to the PDMs on the right with a black, 66 cm (26 in) power cord.
  - The upper controller connects to a PDM on the left with a gray, 152 cm (60 in) power cord.
  - The lower controller connects to a PDM on the right with a black, 66 cm (26 in) power cord.

NOTE: Drive enclosures, when purchased separately, include one 50 cm black cable and one 50 cm gray cable.

The configuration provides complete power redundancy and eliminates all single points of failure for both the AC and DC power distribution.

△ **CAUTION:** Operating the array with a single PDU will result in the following conditions:

- No redundancy
- Louder controllers and disk enclosures due to increased fan speed
- HP P6000 Command View will continuously display a warning condition, making issue monitoring a labor-intensive task

Although the array is capable of doing so, HP strongly recommends that an array operating with a single PDU should not:

- Be put into production
- Remain in this state for more than 24 hours

PDUs

Each Enterprise Virtual Array rack has either a 50- or 60-Hz, dual PDU mounted at the bottom rear of the rack. The 228481-002/228481-003 PDU placement is back-to-back, plugs facing down, with switches on top.

- The standard 50-Hz PDU cable has an IEC 309, 3-wire, 30-A, 50-Hz connector.
- The standard 60-Hz PDU cable has a NEMA L6-30P, 3-wire, 30-A, 60-Hz connector.

If these connectors are not compatible with the site power distribution, you must replace the PDU power cord cable connector.

Each of the two PDU power cables has an AC power source specific connector. The circuit breaker-controlled PDU outputs are routed to a group of four AC receptacles (see Figure 27 (page 54)). The voltages are then routed to PDMs, sometimes referred to as AC power strips, mounted on the two vertical rails in the rear of the rack.

Figure 27 Dual PDU assembly

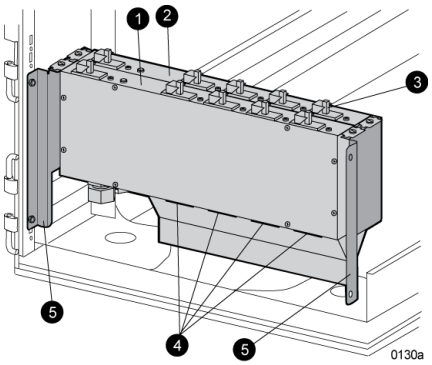

1. PDU 1
2. PDU 2
3. Circuit breakers
4. AC receptacles
5. Mounting hardware

PDU 1

PDU 1 connects to AC power distribution source 1. A PDU 1 failure:

- Disables the power distribution circuit.
- Removes power from PDMs 1, 2, and 3.
- Disables PS 1 in the drive enclosures.
- Disables the upper controller power supply.

PDU 2

PDU 2 connects to AC power distribution source 2. A PDU 2 failure:

- Disables the power distribution circuit.
- Removes power from PDMs 4, 5, and 6.
- Disables PS 2 in the drive enclosures.
- Disables the lower controller power supply.

PDMs

There are six PDMs mounted in the rear of each rack:

- Three mounted on the left vertical rail connect to PDU 1.
- Three mounted on the right vertical rail connect to PDU 2.

Each PDM has eight AC receptacles and one thermal circuit breaker. The PDMs distribute the AC power from the PDUs to the enclosures. Two power sources exist for each controller pair and drive enclosure. If a PDU fails, the system will remain operational.

CAUTION: The AC power distribution within a rack ensures a balanced load to each PDU and reduces the possibility of an overload condition. Changing the cabling to or from a PDM could cause an overload condition. HP supports only the AC power distributions defined in this user guide.

Figure 28 Rack PDM

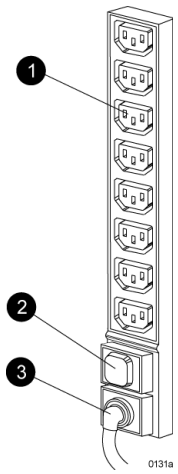

1. Power receptacles
2. Thermal circuit breakers
3. AC power connector

Rack AC power distribution

The power distribution in an Enterprise Virtual Array rack is the same for all variants. The site AC input voltage is routed to the dual PDU assembly mounted in the rack lower rear. Each PDU distributes AC to a maximum of four PDMs mounted on the left and right vertical rails (see Figure 29 (page 56)).

- PDMs 1 through 3 connect to receptacles A through D on PDU 1. Power cords connect these PDMs to the number 1 drive enclosure power supplies and to the upper controller enclosure.
- PDMs 4 through 6 connect to receptacles A through D on PDU 2. Power cords connect these PDMs to the number 2 drive enclosure power supplies and to the lower controller enclosure.

NOTE: The locations of the PDUs and the PDMs are the same in all racks.

Figure 29 Rack AC power distribution

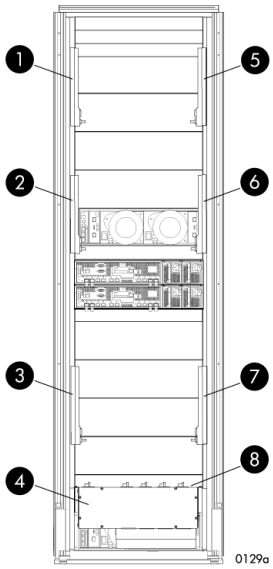

1. PDM 1
2. PDM 2
3. PDM 3
4. PDU 1
5. PDM 4
6. PDM 5
7. PDM 6
8. PDU 2

Rack System/E power distribution components

AC power is distributed to the Rack System/E rack through Power Distribution Units (PDU) mounted on the two vertical rails in the rear of the rack. Up to four PDUs can be mounted in the rack—two mounted on the right side of the cabinet and two mounted on the left side.

Each of the PDU power cables has an AC power source specific connector. The circuit breaker-controlled PDU outputs are routed to a group of ten AC receptacles. The storage system components plug directly into the PDUs.

Rack AC power distribution

The power distribution configuration in a Rack System/E rack depends on the number of storage systems installed in the rack. If one storage system is installed, only two PDUs are required. If multiple storage systems are installed, four PDUs are required.

The site AC input voltage is routed to each PDU mounted in the rack. Each PDU distributes AC through ten receptacles directly to the storage system components.

- PDUs 1 and 3 (optional) are mounted on the left side of the cabinet. Power cords connect these PDUs to the number 1 drive enclosure power supplies and to the controller enclosures.
- PDUs 2 and 4 (optional) are mounted on the right side of the cabinet. Power cords connect these PDUs to the number 2 drive enclosure power supplies and to the controller enclosures.

Moving and stabilizing a rack

⚠ **WARNING!** The physical size and weight of the rack requires a minimum of two people to move. If one person tries to move the rack, injury may occur.

To ensure stability of the rack, always push on the lower half of the rack. Be especially careful when moving the rack over any bump (e.g., door sills, ramp edges, carpet edges, or elevator openings). When the rack is moved over a bump, there is a potential for it to tip over.

Moving the rack requires a clear, uncarpeted pathway that is at least 80 cm (31.5 in) wide for the 60.3 cm (23.7 in) wide, 42U rack. A vertical clearance of 203.2 cm (80 in) should ensure sufficient clearance for the 200 cm (78.7 in) high, 42U rack.

CAUTION: Ensure that no vertical or horizontal restrictions exist that would prevent rack movement without damaging the rack.

Make sure that all four leveler feet are in the fully raised position. This process will ensure that the casters support the rack weight and the feet do not impede movement.

Each rack requires an area 600 mm (23.62 in) wide and 1000 mm (39.37 in) deep (see Figure 30 (page 57)).

Figure 30 Single rack configuration floor space requirements

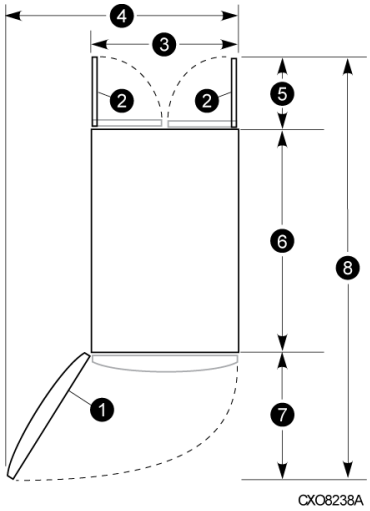

1. Front door
2. Rear door
3. Rack width 600 mm
4. Service area width 813 mm
5. Rear service area depth 300 mm
6. Rack depth 1000 mm
7. Front service area depth 406 mm
8. Total rack depth 1706 mm

If the feet are not fully raised, complete the following procedure:

1. Raise one foot by turning the leveler foot hex nut counterclockwise until the weight of the rack is fully on the caster (see Figure 31 (page 58)).
2. Repeat Step 1 for the other feet.

Figure 31 Raising a leveler foot

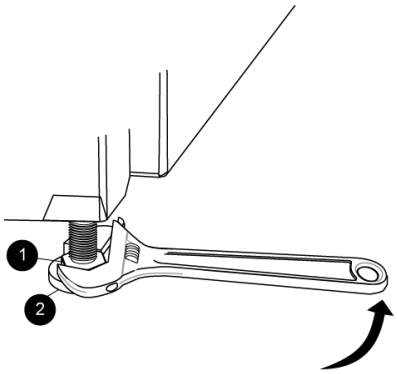

CXO7589A

1. Hex nut
2. Leveler foot

3. Carefully move the rack to the installation area and position it to provide the necessary service areas (see Figure 30 (page 57)).

To stabilize the rack when it is in the final installation location:

1. Use a wrench to lower the foot by turning the leveler foot hex nut clockwise until the caster does not touch the floor. Repeat for the other feet.
2. After lowering the feet, check the rack to ensure it is stable and level.
3. Adjust the feet as necessary to ensure the rack is stable and level.

3 Enterprise Virtual Array operation

This chapter presents the tasks that you might need to perform during normal operation of the storage system.

Best practices

For useful information on managing and configuring your storage system, see the *HP Enterprise Virtual Array configuration best practices* white paper available from

http://h18006.www1.hp.com/storage/arraywhitepapers.html

Operating tips and information

Reserving adequate free space

To ensure efficient storage system operation, a certain amount of unallocated capacity, or free space, should be reserved in each disk group. The recommended amount of free space is influenced by your system configuration. For guidance on how much free space to reserve, see the *HP Enterprise Virtual Array configuration best practices* white paper. See "Best practices" (page 59).

Using FATA disk drives

FATA drives are designed for lower duty cycle applications such as near online data replication for backup. These drives should not be used as a replacement for EVA's high performance, standard duty cycle, Fibre Channel drives. Doing so could shorten the life of the drive.

Changing the host port topology

Before changing the topology settings of an array host port, physically disconnect the host port from its existing connection, change the topology setting, and then reconnect the host port.

For example, to change from a fabric (switched) topology to a direct connect topology, do the following:

1. Disconnect the host port(s) from the Fibre Channel switch.
2. Using the operator control panel (OCP), change the controller host port mode from **`fabric`** to **`direct connect`**.
3. Connect the host HBA(s) directly to the array host port(s).

Host port connection limit on B-series 3200 and 3800 switches

The B-series 3200 and 3800 switches are limited to a maximum of three EVA4x00/6x00/8x00 host ports on a single B-series 3200 and 3800 switch running version 3.2.x. HP recommends not exceeding more than one storage host port connection on a single switch. If you are required to connect more than one storage host port to a single affected switch, separate the connection into different quadrants. Connections are typically dropped following an array controller resynchronization or when an event impacts the fabric, such as rebooting or adding a switch.

Use the following options to avoid or manage the port limitation:

- For all hosts, zone by HBA as defined in the *HP SAN Design Reference Guide*.
- Limit affected switches to only one HBA connection per host.
- Limit placement of the switch as an edge device and not part of the core.

  If the switch drops a connection, reestablish as follows:

  1. Disconnect the Fibre Channel cable from the failed port.
  2. Wait 10 seconds and reconnect the cable. This will cause the port to relog into the fabric and reestablish connection to the array.

Enabling Boot from SAN for Windows direct connect

To ensure that Boot from SAN is successful for Windows hosts that are directly connected to an array, enable the Spin up delay setting in the HBA BIOS. This applies to QLogic and Emulex HBAs.

This workaround applies to all supported Windows operating systems and all supported QLogic and Emulex HBAs. For support details, go to the Single Point of Connectivity Knowledge (SPOCK) website:

http://www.hp.com/support/spock

Windows 2003 MSCS cluster installation

The MSCS cluster installation wizard on Windows 2003 may fail to find the shared quorum device and disk resources may not be auto-created by the cluster setup wizard. This is a known Windows Cluster Setup issue that has existed since Windows 2003 was released.

There are two possible workarounds for this problem:

- The issue and recommended workaround are described in the following Microsoft support article entitled *Shared disks are missing or are marked as "Failed" when you create a server cluster in Windows Server 2003* (ID 886807), which can be downloaded from the following website:

  http://support.microsoft.com/default.aspx?scid=KB;EN-US;886807

- You can bypass this issue by setting the load balancing policy for each LUN to NLB using the MPIO DSM CLI.

Microsoft is currently working on a resolution to address this issue.

Connecting to C-series switches

If C-series switches are not set correctly, the EVA host ports may not log back in to the fabric after changes occur in the fabric. This issue involves the following C-series switch model families: 90xx, 91xx, 92xx, and 95xx. You may also need to restart the controller if you move the array Fibre Channel cable to a different port on a C-series switch, or to a port on a different C-series switch, which causes the corresponding controller host port to become unavailable.

If the switch does not log into the array, disconnect and then reconnect the FC cable on the array or the switch.

If the above recommendations do not correct the problem, it will be necessary to restart the controller to restore host port operation. You only need to restart the controller experiencing the host port problem. This can be done from HP P6000 Command View as follows:

To restart the controller:

1. Open HP P6000 Command View and click the icon of the appropriate storage system. You can select either an initialized or unintialized storage system.
2. On the Initialized Storage System Properties page, click **Shut down**.
3. Under Controller Shutdown, select the appropriate controller (A or B) for restart.
4. Click **Restart**.

   The controller is restarted. After the restart, the host port should be operating normally.

NOTE: If HP P6000 Command View cannot be used to restart the storage system, use the controller operator control panel (OCP). The Restart option is located under the Shutdown Options menu on the OCP.

HP Insight Remote Support software

HP strongly recommends that you install HP Insight Remote Support software to complete the installation or upgrade of your product and to enable enhanced delivery of your HP Warranty,

HP Care Pack Service or HP contractual support agreement. HP Insight Remote Support supplements your monitoring, 24x7 to ensure maximum system availability by providing intelligent event diagnosis, and automatic, secure submission of hardware event notifications to HP, which will initiate a fast and accurate resolution, based on your product's service level. Notifications may be sent to your authorized HP Channel Partner for on-site service, if configured and available in your country. The software is available in two variants:

- HP Insight Remote Support Standard: This software supports server and storage devices and is optimized for environments with 1-50 servers. Ideal for customers who can benefit from proactive notification, but do not need proactive service delivery and integration with a management platform.
- HP Insight Remote Support Advanced: This software provides comprehensive remote monitoring and proactive service support for nearly all HP servers, storage, network, and SAN environments, plus selected non-HP servers that have a support obligation with HP. It is integrated with HP Systems Insight Manager. A dedicated server is recommended to host both HP Systems Insight Manager and HP Insight Remote Support Advanced.

Details for both versions are available at:

http://www.hp.com/go/insightremotesupport

To download the software, go to Software Depot:

http://www.software.hp.com

Select **Insight Remote Support** from the menu on the right.

Failback preference setting for HSV controllers

Table 25 (page 62) describes the failback preference behavior for the controllers.

Table 25 Failback preference behavior

Setting	Point in time	Behavior
No preference	At initial presentation	The units are alternately brought online to Controller A or to Controller B.
	On dual boot or controller resynch	If cache data for a LUN exists on a particular controller, the unit will be brought online there. Otherwise, the units are alternately brought online to Controller A or to Controller B.
	On controller failover	All LUNs are brought online to the surviving controller.
	On controller failback	All LUNs remain on the surviving controller. There is no failback except if a host moves the LUN using SCSI commands.
Path A - Failover Only	At initial presentation	The units are brought online to Controller A.
	On dual boot or controller resynch	If cache data for a LUN exists on a particular controller, the unit will be brought online there. Otherwise, the units are brought online to Controller A.
	On controller failover	All LUNs are brought online to the surviving controller.
	On controller failback	All LUNs remain on the surviving controller. There is no failback except if a host moves the LUN using SCSI commands.
Path B - Failover Only	At initial presentation	The units are brought online to Controller B.
	On dual boot or controller resynch	If cache data for a LUN exists on a particular controller, the unit will be brought online there. Otherwise, the units are brought online to Controller B.
	On controller failover	All LUNs are brought online to the surviving controller.
	On controller failback	All LUNs remain on the surviving controller. There is no failback except if a host moves the LUN using SCSI commands.
Path A - Failover/Failback	At initial presentation	The units are brought online to Controller A.
	On dual boot or controller resynch	If cache data for a LUN exists on a particular controller, the unit will be brought online there. Otherwise, the units are brought online to Controller A.
	On controller failover	All LUNs are brought online to the surviving controller.
	On controller failback	All LUNs remain on the surviving controller. After controller restoration, the units that are online to Controller B and set to Path A are brought online to Controller A. This is a one time occurrence. If the host then moves the LUN using SCSI commands, the LUN will remain where moved.
Path B - Failover/Failback	At initial presentation	The units are brought online to Controller B.

Table 25 Failback preference behavior *(continued)*

Setting	Point in time	Behavior
	On dual boot or controller resynch	If cache data for a LUN exists on a particular controller, the unit will be brought online there. Otherwise, the units are brought online to Controller B.
	On controller failover	All LUNs are brought online to the surviving controller.
	On controller failback	All LUNs remain on the surviving controller. After controller restoration, the units that are online to Controller A and set to Path B are brought online to Controller B. This is a one time occurrence. If the host then moves the LUN using SCSI commands, the LUN will remain where moved.

Table 26 (page 63) describes the failback default behavior and supported settings when AULA-compliant multipath software is running with each operating system. Recommended settings may vary depending on your configuration or environment.

Table 26 Failback Settings by operating system

Operating system	Default behavior	Supported settings
HP-UX	Host follows the unit[1]	No Preference Path A/B – Failover Only Path A/B – Failover/Failback
IBM AIX	Host follows the unit[1]	No Preference Path A/B – Failover Only Path A/B – Failover/Failback
Linux	Host follows the unit [1]	No Preference Path A/B – Failover Only Path A/B – Failover/Failback
OpenVMS	Host follows the unit	No Preference Path A/B – Failover Only Path A/B – Failover/Failback (recommended)
Oracle Solaris	Host follows the unit[1]	No Preference Path A/B – Failover Only Path A/B – Failover/Failback
VMware	Host follows the unit[1]	No Preference Path A/B – Failover Only Path A/B – Failover/Failback
Windows	Failback performed on the host	No Preference Path A/B – Failover Only Path A/B – Failover/Failback

[1] If preference has been configured to ensure a more balanced controller configuration, the Path A/B – Failover/Failback setting is required to maintain the configuration after a single controller reboot.

Changing virtual disk failover/failback setting

Changing the failover/failback setting of a virtual disk may impact which controller presents the disk. Table 27 (page 64) identifies the presentation behavior that results when the failover/failback setting for a virtual disk is changed.

NOTE: If the new setting causes the presentation of the virtual disk to move to a new controller, any snapshots or snapclones associated with the virtual disk will also be moved.

Table 27 Impact on virtual disk presentation when changing failover/failback setting

New setting	Impact on virtual disk presentation
No Preference	None. The disk maintains its original presentation.
Path A Failover	If the disk is currently presented on controller B, it is moved to controller A. If the disk is on controller A, it remains there.
Path B Failover	If the disk is currently presented on controller A, it is moved to controller B. If the disk is on controller B, it remains there.
Path A Failover/Failback	If the disk is currently presented on controller B, it is moved to controller A. If the disk is on controller A, it remains there.
Path B Failover/Failback	If the disk is currently presented on controller A, it is moved to controller B. If the disk is on controller B, it remains there.

Storage system shutdown and startup

The storage system is shut down using HP P6000 Command View. The shutdown process performs the following functions in the indicated order:

1. Flushes cache
2. Removes power from the controllers
3. Disables cache battery power
4. Removes power from the drive enclosures
5. Disconnects the system from HP P6000 Command View

NOTE:

- The storage system may take a long time to complete the necessary cache flush during controller shutdown when snapshots are being used. The delay may be particularly long if multiple child snapshots are used, or if there has been a large amount of write activity to the snapshot source virtual disk.
- Individual EVA storage array components should not be powered off during normal operation. Before powering off any storage system component, contact your HP-authorized service representative for assistance.

Shutting down the storage system

To shut the storage system down, perform the following steps:

1. Start HP P6000 Command View.
2. Select the appropriate storage system in the Navigation pane.

   The Initialized Storage System Properties window for the selected storage system opens.
3. Click **Shut down**.

   The Shutdown Options window opens.

4. Under System Shutdown click **Power Down**. If you want to delay the initiation of the shutdown, enter the number of minutes in the Shutdown delay field.

   The controllers complete an orderly shutdown and then power off. The disk enclosures then power off. Wait for the shutdown to complete.
5. Turn off the power switch (callout 4 in Figure 17 (page 39)) on the rear of each HSV controller.
6. Turn off the circuit breakers on both of the EVA rack Power Distribution Units (PDU).
7. If your management server is an SMA and you are not using it to manage other storage arrays, shut down the SMA. From the SMA user interface, click **Settings > Maintenance > Shutdown**.

Starting the storage system

To start a storage system, perform the following steps:

1. Verify that each fabric Fibre Channel switch to which the HSV controllers are connected is powered up and fully booted. The power indicator on each switch should be on.

   If you must power up the SAN switches, wait for them to complete their power-on boot process before proceeding. This may take several minutes.
2. If the management server you shut down is an SMA, power it on and wait for it to completely boot. Verify the SMA is running by logging into it using the web interface.

   NOTE: Before applying power to the rack, ensure that the power switch on each HSV controller is off.
3. Power on the circuit breakers on both EVA rack PDUs. Verify that all drive enclosures are operating properly. The status indicator and the power indicator should be on (green).
4. Wait three minutes and then verify that all disk drives are ready. The drive ready indicator and the drive online indicator should be on (green).
5. Power on the upper controller. It takes the role of primary controller.
6. If you want the preferred path setting to be applied, wait three seconds and power on the lower controller. (Otherwise, wait 10 seconds before powering on the lower controller.) It takes the role of secondary controller.
7. Verify that the Operator Control Panel (OCP) display on each controller displays the storage system name and the EVA WWN.
8. Start HP P6000 Command View and verify connection to the storage system. If the storage system is not visible, click **HSV Storage Network** in the Navigation pane, and then click **Discover** in the Content pane to discover the array.

   NOTE: If the storage system is still not visible, reboot the management server to re-establish the communication link.
9. Check the storage system status using HP P6000 Command View to ensure everything is operating properly. If any status indicator is not normal, check the log files or contact your HP-authorized service provider for assistance.

Saving storage system configuration data

As part of an overall data protection strategy, storage system configuration data should be saved during initial installation, and whenever major configuration changes are made to the storage system. This includes adding or removing disk drives, creating or deleting disk groups, and adding or deleting virtual disks. The saved configuration data can save substantial time should it ever become necessary to re-initialize the storage system. The configuration data is saved to a series of files stored in a location other than on the storage system.

This procedure can be performed from the Storage Management Appliance (SMA) or management server where HP P6000 Command View is installed, or any host that can run the Storage System Scripting Utility (SSSU) to communicate with the HP P6000 Command View.

NOTE: For more information on using SSSU, see the *HP Storage System Scripting Utility reference*. See "Related information" (page 101).

1. Double-click on the SSSU desktop icon to run the application. When prompted, enter Manager (management server name or IP address), User name, and Password.
2. Enter `LS SYSTEM` to display the EVA storage systems managed by the management server.
3. Enter `SELECT SYSTEM system name`, where *system name* is the name of the storage system.

   The storage system name is case sensitive. If there are spaces between the letters in the name, quotes must enclose the name: for example, `SELECT SYSTEM "Large EVA"`.
4. Enter `CAPTURE CONFIGURATION`, specifying the full path and filename of the output files for the configuration data.

   The configuration data is stored in a series of from one to five files, which are SSSU scripts. The file names begin with the name you select, with the restore step appended. For example, if you specify a file name of `LargeEVA.txt`, the resulting configuration files would be `LargeEVA_Step1A.txt, LargeEVA_Step1B`, etc.

   The contents of the configuration files can be viewed with a text editor.

NOTE: If the storage system contains disk drives of different capacities, the SSSU procedures used do not guarantee that disk drives of the same capacity will be exclusively added to the same disk group. If you need to restore an array configuration that contains disks of different sizes and types, you must manually recreate these disk groups. The controller software and the utility's `CAPTURE CONFIGURATION` command are not designed to automatically restore this type of configuration. For more information, see the *HP Storage System Scripting Utility Reference*.

Example 1 Saving configuration data using SSSU on a Windows Host

To save the storage system configuration:

1. Double-click on the SSSU desktop icon to run the application. When prompted, enter Manager (management server name or IP address), User name, and Password.
2. Enter `LS SYSTEM` to display the EVA storage systems managed by the management server.
3. Enter `SELECT SYSTEM system name`, where *system name* is the name of the storage system.
4. Enter `CAPTURE CONFIGURATION pathname\filename`, where *pathname* identifies the location where the configuration files will be saved, and `filename` is the name used as the prefix for the configurations files: for example, `CAPTURE CONFIGURATION c:\EVAConfig\LargeEVA`
5. Enter `EXIT` to close the command window.

Example 2 Restoring configuration data using SSSU on a Windows Host

To restore the storage system configuration:

1. Double-click on the SSSU desktop icon to run the application.
2. Enter `FILE pathname\filename`, where *pathname* identifies the location where the configuration files are be saved and *filename* is the name of the first configuration file: for example, `FILE c:\EVAConfig\LargeEVA_Step1A.txt`
3. Repeat the preceding step for each configuration file.

Adding disk drives to the storage system

As your storage requirements grow, you may be adding disk drives to your storage system. Adding new disk drives is the easiest way to increase the storage capacity of the storage system. Disk drives can be added online without impacting storage system operation.

△ **CAUTION:** When adding disks to an expansion cabinet on an EVA8000/8100, do not install a disk in bays 12, 13, and 14 in enclosures 17, 20, and 24. These bays in enclosures 17, 20, and 24 do not receive a hard assigned AL_PA. Installing a disk in any of these slots may impact the operation of the storage system. HP also recommends that you keep three additional bays open to maintain the maximum device count of 120. For ease of use and consistency in configurations, HP recommends keeping bays 12, 13, and 14 open in enclosures 16 and 19.

Consider the following best practices to improve availability when adding disk to an array:

- Install high performance and FATA disk drives in separate groups. These different drive types must be in separate disk groups. You may also want to consider separating different drive capacities and spindle speeds into different groups.
- High performance and FATA disk drives can be installed in the same disk enclosure.
- The disk drives should be distributed evenly across the disk enclosures. The number of disks of a given type in each enclosure should not differ by more than one. For example, no enclosure should have two disks until all the other enclosures have at least one.
- Disk drives should be installed in vertical columns within the disk enclosures. Add drives vertically in multiples of eight, completely filling columns if possible. Disk groups are more robust if filled with the same number of disk drives in each enclosure. See Figure 33 (page 68) for an example.
- For growing existing applications, if the operating system supports virtual disk growth, increase virtual disk size. Otherwise, use a software volume manager to add new virtual disks to applications.

- Set the add disk option to manual. See "Changing the Device Addition Policy" (page 69) for more information.
- When adding multiple disk drives, add a disk and wait for its activity indicator (1) to stop flashing (up to 90 seconds) before installing the next disk (see Figure 32 (page 68)). This procedure must be followed to avoid unexpected EVA system behavior.

Figure 32 Disk drive activity indicator

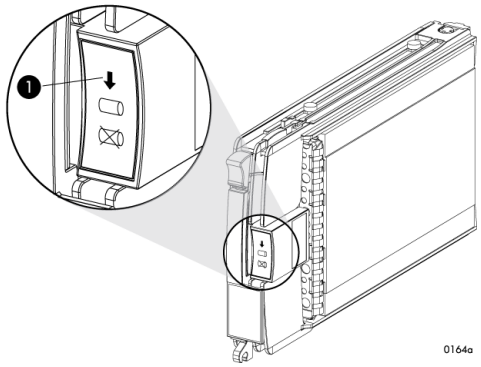

Creating disk groups

The new disks you add will typically be used to create new disk groups. Although you cannot select which disks will be part of a disk group, you can control this by building the disk groups sequentially.

Add the disk drives required for the first disk group, and then create a disk group using these disk drives. Now add the disk drives for the second disk group, and then create that disk group. This process gives you control over which disk drives are included in each disk group. Figure 33 (page 68) shows the sequential building of vertical disk groups.

NOTE: Standard and FATA disk drives must be in separate disk groups. Disk drives of different capacities and spindle speeds can be included in the same disk group, but you may want to consider separating them into separate disk groups.

Figure 33 Sequential building of vertical disk groups

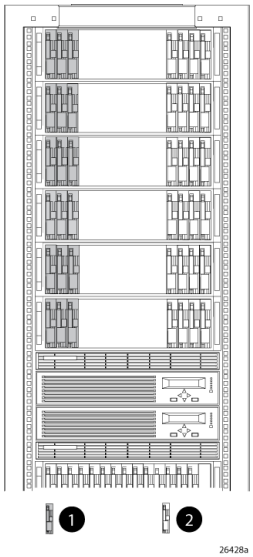

1. Disks installed in first group
2. Disks installed in second group

Adding a disk drive

This section describes the procedure for adding a Fibre Channel disk drive.

Removing the drive blank

1. Grasp the drive blank by the two mounting tabs (see Figure 34 (page 69)).
2. Lift up on the lower mounting tab and pull the blank out of the enclosure.

Figure 34 Removing the drive blank

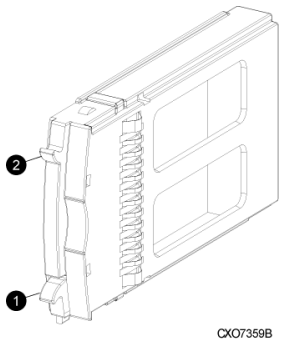

1. Upper mounting tab
2. Lower mounting tab

Changing the Device Addition Policy

To prevent the storage system from automatically grouping a new disk drive that may have the incorrect firmware on it, the Device Addition Policy must be checked and set to manual if necessary:

1. Open HP P6000 Command View and in the navigation pane, select the storage system.

   The Initialized Storage System Properties window opens.
2. Click **System Options**.
3. Select **Set system operational policies**.
4. If the Device Addition Policy is set to **Automatic**, change it to **Manual**.
5. Click **Save changes**.

NOTE: After the Device Addition Policy has been changed to manual mode, HP recommends as a best practice not returning the policy to the automatic device addition setting. This will eliminate the need to make this change for future code load operations. However, if you prefer returning the Device Addition Policy to automatic, repeat steps 1 through 5 after verifying the disk drive has the correct firmware version.

Installing the disk drive

1. Push in the ejector button on the disk drive and pull the release lever down to the full open position.
2. Insert the drive into the enclosure as far as it will go (1, Figure 35 (page 70)).
3. Close the release lever until it engages the ejector button, and the disk drive seats in the backplane (2, Figure 35 (page 70)).
4. Press in firmly on the disk drive to ensure it is seated properly.

Figure 35 Installing the disk drive

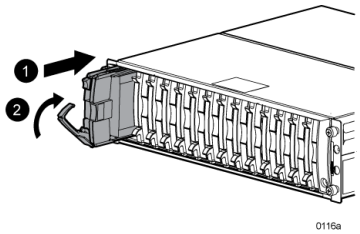

Checking status indicators

Check the following to verify that the disk drive is operating normally:

NOTE: It may take up to 10 minutes for the component to display good status.

- Check the disk drive status indicators. See Figure 36 (page 71).
  - Activity indicator (1) should be on or flashing
  - Online indicator (2) should be on or flashing
  - Fault indicator (3) should be off
- Check the following using HP P6000 Command View.
  - Navigate to the disk drive and check the operational state. It should be ☑.
  - Ensure the disk drive is using the correct firmware. Record the **Model number** and the **Firmware version** of the disk. Check the firmware version against the supported disk firmware in the *HP Enterprise Virtual Array Disk Drive Firmware Support*. See "Related information" (page 101) for a link to this document.

    If the disk drive is using an unsupported version of firmware, download the correct firmware from the following website and install it using the instructions included with the firmware file. Do not add the disk drive to a disk group if it is using an unsupported firmware version.

    http://www.hp.com/support/evadiskfirmware

    NOTE: When downloading the firmware, make sure you use the disk model number to locate the correct firmware file. If you have difficulty locating the correct firmware, contact your HP-authorized service representative for assistance. If you are running XCS 6.000, verify that leveling is not in progress before upgrading the disk drive firmware. In HP P6000 Command View, go to the General tab of the Disk Group Properties window and verify that the Leveling field displays Inactive. If it displays Active, wait for leveling to complete before performing the upgrade. This does not apply if you are running XCS 6.100 or later.

Figure 36 Disk drive status indicators

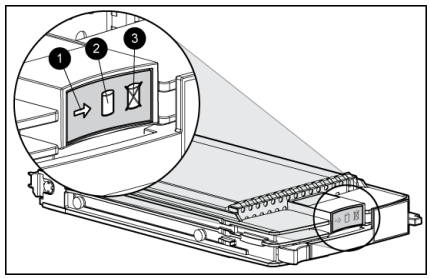

1. Activity
2. Online
3. Fault

Adding the disk to a disk group

After replacing the disk, use HP P6000 Command View to add it to a disk group.

1. In the Navigation pane, select **Storage system > Hardware > Rack > Disk enclosure > Bay**
2. In the Content pane, select the **Disk Drive** tab.
3. Click **Group** to initiate the process for adding the disk to a disk group.

NOTE: If the Device Addition Policy is set to automatic, the disk will automatically be added to a disk group. In this case the **Group** option will not be available.

Handling fiber optic cables

This section provides protection and cleaning methods for fiber optic connectors.

Contamination of the fiber optic connectors on either a transceiver or a cable connector can impede the transmission of data. Therefore, protecting the connector tips against contamination or damage is imperative. The tips can be contaminated by touching them, by dust, or by debris. They can be damaged when dropped. To protect the connectors against contamination or damage, use the dust covers or dust caps provided by the manufacturer. These covers are removed during installation, and are installed whenever the transceivers or cables are disconnected. Cleaning the connectors should remove contamination.

The transceiver dust caps protect the transceivers from contamination. **Do not discard the dust covers.**

△ **CAUTION:** To avoid damage to the connectors, always install the dust covers or dust caps whenever a transceiver or a fiber cable is disconnected. Remove the dust covers or dust caps from transceivers or fiber cable connectors only when they are connected. **Do not discard the dust covers.**

To minimize the risk of contamination or damage, do the following:

- **Dust covers** — Remove and set aside the dust covers and dust caps when installing an I/O module, a transceiver or a cable. Install the dust covers when disconnecting a transceiver or cable.
- **When to clean** — If a connector may be contaminated, or if a connector has not been protected by a dust cover for an extended period of time, clean it.
- **How to clean**:
  1. Wipe the connector with a lint-free tissue soaked with 100% isopropyl alcohol.
  2. Wipe the connector with a dry, lint-free tissue.
  3. Dry the connector with moisture-free compressed air.

4 Configuring application servers

Overview

This chapter provides general connectivity information for all the supported operating systems. Where applicable, an OS-specific section is included to provide more information.

Clustering

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

See the Single Point of Connectivity Knowledge (SPOCK) website (http://www.hp.com/storage/spock) for the clustering software supported on each operating system.

NOTE: For OpenVMS, you must make the Console LUN ID and OS unit IDs unique throughout the entire SAN, not just the controller subsystem.

Multipathing

Multipathing software provides a multiple-path environment for your operating system. See the following website for more information:

http://h18006.www1.hp.com/products/sanworks/multipathoptions/index.html

See the Single Point of Connectivity Knowledge (SPOCK) website (http://www.hp.com/storage/spock) for the multipathing software supported on each operating system.

Installing Fibre Channel adapters

For all operating systems, supported Fibre Channel adapters (FCAs) must be installed in the host server in order to communicate with the EVA.

NOTE: Traditionally, the adapter that connects the host server to the fabric is called a host bus adapter (HBA). The server HBA used with the EVA4x00/6x00/8x00 is called a Fibre Channel adapter (FCA). You might also see the adapter called a Fibre Channel host bus adapter (Fibre Channel HBA) in other related documents.

Follow the hardware installation rules and conventions for your server type. The FCA is shipped with its own documentation for installation. See that documentation for complete instructions. You need the following items to begin:

- FCA boards and the manufacturer's installation instructions
- Server hardware manual for instructions on installing adapters
- Tools to service your server

The FCA board plugs into a compatible I/O slot (PCI, PCI-X, PCI-E) in the host system. For instructions on plugging in boards, see the hardware manual.

You can download the latest FCA firmware from the following website: http://www.hp.com/support/downloads. Enter HBA in the **Search Products** box and then select your product. See the Single Point of Connectivity Knowledge (SPOCK) website (http://www.hp.com/storage/spock) for supported FCAs by operating system.

Testing connections to the EVA

After installing the FCAs, you can create and test connections between the host server and the EVA. For all operating systems, you must:

- Add hosts
- Create and present virtual disks
- Verify virtual disks from the hosts

The following sections provide information that applies to all operating systems. For OS-specific details, see the applicable operating system section.

Adding hosts

To add hosts using HP P6000 Command View:

1. Retrieve the world-wide names (WWNs) for each FCA on your host. You need this information to select the host FCAs in HP P6000 Command View.
2. Use HP P6000 Command View to add the host and each FCA installed in the host system.

   NOTE: To add hosts using HP P6000 Command View, you must add each FCA installed in the host. Select **Add Host** to add the first adapter. To add subsequent adapters, select **Add Port**. Ensure that you add a port for each active FCA.

3. Select the applicable operating system for the host mode.

   Table 28 Operating system and host mode selection

Operating System	Host mode selection in HP P6000 Command View
HP-UX	HP-UX
IBM AIX	IBM AIX
Linux	Linux
Mac OS X	Linux
OpenVMS	OVMS
Oracle Solaris	Oracle Solaris
VMware	VMware
Citrix XenServer	Linux
Microsoft Windows	Microsoft Windows
	Microsoft Windows 2008

4. Check the Host folder in the Navigation pane of HP P6000 Command View to verify that the host FCAs are added.

   NOTE: More information about HP P6000 Command View is available at http://www.hp.com/support/manuals. Click **Storage Software** under Storage, and then select **HP Command View EVA Software** under Storage Device Management Software.

Creating and presenting virtual disks

To create and present virtual disks to the host server:

1. From HP P6000 Command View, create a virtual disk on the EVA4x00/6x00/8x00.
2. Specify values for the following parameters:
   - Virtual disk name
   - Vraid level
   - Size
3. Present the virtual disk to the host you added.
4. If applicable (OpenVMS) select a LUN number if you chose a specific LUN on the Virtual Disk Properties window.

Verifying virtual disk access from the host

To verify that the host can access the newly presented virtual disks, restart the host or scan the bus.

If you are unable to access the virtual disk:

- Verify that all cabling is connected to the switch, EVA, and host.
- Verify that all firmware levels are appropriate for your configuration. For more information, refer to the Enterprise Virtual Array QuickSpecs and associated release notes. See "Related information" (page 101) for the location of these documents.
- Ensure that you are running a supported version of the host operating system. For more information, see the *HP P6000 Enterprise Virtual Array Compatibility Reference*.
- Ensure that the correct host is selected as the operating system for the virtual disk in HP P6000 Command View.
- Ensure that the host WWN number is set correctly (to the host you selected).
- Verify that the FCA switch settings are correct.
- Verify that the virtual disk is presented to the host.
- Verify that the zoning is correct for your configuration.

Configuring virtual disks from the host

After you create the virtual disks on the EVA4x00/6x00/8x00 and rescan or restart the host, follow the host-specific conventions for configuring these new disk resources. For instructions, see the documentation included with your server.

HP-UX

To create virtual disks for HP-UX, scan the bus and then create volume groups on a virtual disk.

Scanning the bus

To scan the FCA bus and display information about the EVA4x00/6x00/8x00 devices:

1. Enter the command `# ioscan -fnCdisk` to start the rescan.

   All new virtual disks become visible to the host.
2. Assign device special files to the new virtual disks using the `insf` command:

   ```
   # insf -e
   ```

   NOTE: Lowercase `e` assigns device special files only to the new devices—in this case, the virtual disks. Uppercase `E` reassigns device special files to all devices.

The following is a sample output from an `ioscan` command:

```
# ioscan -fnCdisk
```

```
# ioscan -fnCdisk
Class     I  H/W Patch          Driver    S/W      H/W Type    Description
```

```
              State
================================================================================
ba        3  0/6                    lba       CLAIMED  BUS_NEXUS  Local PCI Bus
                                                                   Adapter (782)
fc        2  0/6/0/0                td        CLAIMED  INTERFACE  HP Tachyon XL@ 2 FC
                                                                   Mass Stor Adap /dev/td2
fcp       0  0/6/0/0.39             fcp       CLAIMED  INTERFACE  FCP Domain
ext_bus   4  0/6/00.39.13.0.0       fcparray  CLAIMED  INTERFACE  FCP Array Interface
target    5  0/6/0/0.39.13.0.0.0     tgt       CLAIMED  DEVICE
ctl       4  0/6/0/0.39.13.0.0.0.0  sctl      CLAIMED  DEVICE     HP  HSV300  /dev/rscsi/c4t0d0
disk      22 0/6/0/0.39.13.0.0.0.1  sdisk     CLAIMED  DEVICE     HP  HSV300  /dev/dsk/c4t0d1
                                                                                /dev/rdsk/c4t0d
ext_bus   5  0/6/0/0.39.13.255.0     fcpdev    CLAIMED  INTERFACE  FCP Device Interface
target    8  0/6/0/0.39.13.255.0.0   tgt       CLAIMED  DEVICE
ctl       20 0/6/0/0.39.13.255.0.0.0  sctl      CLAIMED  DEVICE     HP HSV300 /dev/rscsi/c5t0d0
ext_bus   10 0/6/0/0.39.28.0.0        fcparray  CLAIMED  INTERFACE  FCP Array Interface
target    9  0/6/0/0.39.28.0.0.0     tgt       CLAIMED  DEVICE
ctl       40 0/6/0/0.39.28.0.0.0.0   sctl      CLAIMED  DEVICE     HP HSV300 /dev/rscsi/c10t0d0
disk      46 0/6/0/0.39.28.0.0.0.2   sdisk     CLAIMED  DEVICE     HP HSV300 /dev/dsk/c10t0d2
                                                                                /dev/rdsk/c10t0d2
disk      47 0/6/0/0.39.28.0.0.0.3   sdisk     CLAIMED  DEVICE     HP HSV300 /dev/dsk/c10t0d3
                                                                                /dev/rdsk/c10t0d3
disk      48 0/6/0/0.39.28.0.0.0.4   sdisk     CLAIMED  DEVICE     HP HSV300 /dev/dsk/c10t0d4
                                                                                /dev/rdsk/c10t0d4
disk      49 0/6/0/0.39.28.0.0.0.5   sdisk     CLAIMED  DEVICE     HP HSV300 /dev/dsk/c10t0d5
                                                                                /dev/rdsk/c10t0d5
disk      50 0/6/0/0.39.28.0.0.0.6   sdisk     CLAIMED  DEVICE     HP HSV300 /dev/dsk/c10t0d
                                                                                /dev/rdsk/c10t0d6
disk      51 0/6/0/0.39.28.0.0.0.7   sdisk     CLAIMED  DEVICE     HP HSV300 /dev/dsk/c10t0d7
                                                                                /dev/rdsk/c10t0d7
```

Creating volume groups on a virtual disk using vgcreate

You can create a volume group on a virtual disk by issuing a `vgcreate` command. This builds the virtual group block data, allowing HP-UX to access the virtual disk. See the `pvcreate`, `vgcreate`, and `lvcreate` man pages for more information about creating disks and file systems. Use the following procedure to create a volume group on a virtual disk:

NOTE: Italicized text is for example only.

1. To create the physical volume on a virtual disk, enter the following command:

   `# pvcreate -f /dev/rdsk/`*`c32t0d1`*

2. To create the volume group directory for a virtual disk, enter the command:

   `# mkdir /dev/`*`vg01`*

3. To create the volume group node for a virtual disk, enter the command:

   `# mknod /dev/vg01/group c 64 0x010000`

   The designation *`64`* is the major number that equates to the 64-bit mode. The *`0x01`* is the minor number in hex, which must be unique for each volume group.

4. To create the volume group for a virtual disk, enter the command:

   `# vgcreate –f /dev/`*`vg01`* `/dev/dsk/`*`c32t0d1`*

5. To create the logical volume for a virtual disk, enter the command:

   `# lvcreate -L1000 /dev/`*`vg01/lvol1`*

   In this example, a 1-Gb logical volume (*`lvol1`*) is created.

6. Create a file system for the new logical volume by creating a file system directory name and inserting a mount tap entry into `/etc/fstab`.
7. Run the command `mkfs` on the new logical volume. The new file system is ready to mount.

IBM AIX

Accessing IBM AIX utilities

You can access IBM AIX utilities such as the Object Data Manager (ODM), on the following website:

http://www.hp.com/support/downloads

In the Search products box, enter **MPIO**, and then click **AIX MPIO PCMA for HP Arrays**. Select **IBM AIX**, and then select your software storage product.

Adding hosts

To determine the active FCAs on the IBM AIX host, enter:

```
# lsdev -Cc adapter |grep fcs
```

Output similar to the following appears:

```
fcs0      Available 1H-08    FC Adapter
fcs1      Available 1V-08    FC Adapter
# lscfg -vl
fcs0  fcs0            U0.1-P1-I5/Q1  FC Adapter
    Part Number.................80P4543
    EC Level....................A
    Serial Number...............1F4280A419
    Manufacturer................001F
    Feature Code/Marketing ID...280B
    FRU Number.................     80P4544
    Device Specific.(ZM)........3
    Network Address.............10000000C940F529
    ROS Level and ID............02881914
    Device Specific.(Z0)........1001206D
    Device Specific.(Z1)........00000000
    Device Specific.(Z2)........00000000
    Device Specific.(Z3)........03000909
    Device Specific.(Z4)........FF801315
    Device Specific.(Z5)........02881914
    Device Specific.(Z6)........06831914
    Device Specific.(Z7)........07831914
    Device Specific.(Z8)........20000000C940F529
    Device Specific.(Z9)........TS1.90A4
    Device Specific.(ZA)........T1D1.90A4
    Device Specific.(ZB)........T2D1.90A4
    Device Specific.(YL)........U0.1-P1-I5/Q1b.
```

Creating and presenting virtual disks

When creating and presenting virtual disks to an IBM AIX host, be sure to:

1. Set the OS unit ID to **0**.
2. Set Preferred path/mode to **No Preference**.
3. Select a LUN number if you chose a specific LUN on the Virtual Disk Properties window.

Verifying virtual disks from the host

To scan the IBM AIX bus and list all EVA devices, enter: `cfgmgr -v`

The `-v` switch (verbose output) requests a full output.

Output similar to the following is displayed:

```
hdisk1  Available 1V-08-01     HP HSV300 Enterprise Virtual Array
hdisk2  Available 1V-08-01     HP HSV300 Enterprise Virtual Array
hdisk3  Available 1V-08-01     HP HSV300 Enterprise Virtual Array
```

Linux

Driver failover mode

If you use the INSTALL command without command options, the driver's failover mode depends on whether a QLogic driver is already loaded in memory (listed in the output of the lsmod command). Possible driver failover mode scenarios include:

- If an hp_qla2x00src driver RPM is already installed, the new driver RPM uses the failover of the previous driver package.
- If there is no QLogic driver module (qla2xxx module) loaded, the driver defaults to failover mode. This is also true if an inbox driver is loaded that does not list output in the /proc/scsi/qla2xxx directory.
- If there is a driver loaded in memory that lists the driver version in /proc/scsi/qla2xxx but no driver RPM has been installed, then the driver RPM loads the driver in the failover mode that the driver in memory currently uses.

Installing a Qlogic driver

NOTE: The HP Emulex driver kit performs in a similar manner; use ./INSTALL -h to list all supported arguments.

1. Download the appropriate driver kit for your distribution. The driver kit file is in the format hp_qla2x00-yyyy-mm-dd.tar.gz.
2. Copy the driver kit to the target system.
3. Uncompress and untar the driver kit using the following command:

   ```
   # tar zxvf hp_qla2x00-yyyy-mm-dd.tar.gz
   ```
4. Change directory to the hp_qla2x00-yyyy-mm-dd directory.
5. Execute the INSTALL command.

   The INSTALL command syntax varies depending on your configuration.

   If a previous driver kit is installed, you can invoke the INSTALL command without any arguments. To use the currently loaded configuration:

   ```
   # ./INSTALL
   ```

   To force the installation to failover mode, use the -f flag:

   ```
   # ./INSTALL -f
   ```

   To force the installation to single-path mode, use the -s flag:

   ```
   # ./INSTALL -s
   ```

   To list all supported arguments, use the -h flag:

   ```
   # ./INSTALL -h
   ```

   The INSTALL script installs the appropriate driver RPM for your configuration, as well as the appropriate fibreutils RPM.
6. Once the INSTALL script is finished, you will either have to reload the QLogic driver modules (qla2xxx, qla2300, qla2400, qla2xxx_conf) or reboot your server.

   To reload the driver use one or more of the following commands, as applicable:

   ```
   # /opt/hp/src/hp_qla2x00src/unload.sh
   # modprobe qla2xxx_conf
   # modprobe qla2xxx
   # modprobe qla2300
   ```

```
# modprobe qla2400
```

To reboot the server, enter the `reboot` command.

CAUTION: If the boot device is attached to the SAN, you must reboot the host.

7. To verify which RPM versions are installed, use the `rpm` command with the `-q` option. For example:

```
# rpm -q hp_qla2x00src
# rpm -q fibreutils
```

Upgrading Linux components

If you have any installed components from a previous solution kit or driver kit, such as the `qla2x00` RPM, invoke the `INSTALL` script with no arguments, as shown in the following example:

```
# ./INSTALL
```

To manually upgrade the components, select one of the following kernel distributions:

- For 2.4 kernel based distributions, use version 7.xx.
- For 2.6 kernel based distributions, use version 8.xx.

Depending on the kernel version you are running, upgrade the driver RPM as follows:

- For the `hp_qla2x00src` RPM:

  ```
  # rpm -Uvh hp_qla2x00src- version-revision.linux.rpm
  ```

- For `fibreutils` RPM, you have two options:
  - To upgrade the driver:

    ```
    # rpm -Uvh fibreutils-version-revision.linux.architecture.rpm
    ```

  - To remove the existing driver, and install a new driver:

    ```
    # rpm -e fibreutils
    # rpm -ivh fibreutils-version-revision.linux.architecture.rpm
    ```

Upgrading qla2x00 RPMs

If you have a `qla2x00` RPM from HP installed on your system, use the `INSTALL` script to upgrade from `qla2x00` RPMs. The `INSTALL` script removes the old `qla2x00` RPM and installs the new `hp_qla2x00src` while keeping the driver settings from the previous installation. The script takes no arguments. Use the following command to run the `INSTALL` script:

```
# ./INSTALL
```

NOTE: IF you are going to use the failover functionality of the QLA driver, uninstall Secure Path and reboot before you attempt to upgrade the driver. Failing to do so can cause a kernel panic.

Detecting third-party storage

The preinstallation portion of the RPM contains code to check for non-HP storage. The reason for doing this is to prevent the RPM from overwriting any settings that another vendor may be using. You can skip the detection process by setting the environmental variable `HPQLAX00FORCE` to **`y`** by issuing the following commands:

```
# HPQLA2X00FORCE=y
# export HPQLA2X00FORCE
```

You can also use the `-F` option of the `INSTALL` script by entering the following command:

```
# ./INSTALL -F
```

Compiling the driver for multiple kernels

If your system has multiple kernels installed on it, you can compile the driver for all the installed kernels by setting the `INSTALLALLKERNELS` environmental variable to `y` and exporting it by issuing the following commands:

```
# INSTALLALLKERNELS=y
# export INSTALLALLKERNELS
```

You can also use the `-a` option of the `INSTALL` script as follows:

```
# ./INSTALL -a
```

Uninstalling the Linux components

To uninstall the components, use the `INSTALL` script with the `-u` option as shown in the following example:

```
# ./INSTALL -u
```

To manually uninstall all components, or to uninstall just one of the components, use one or all of the following commands:

```
# rpm -e fibreutils
# rpm -e hp_qla2x00
# rpm -e hp_qla2x00src
```

Using the source RPM

In some cases, you may have to build a binary `hp_qla2x00` RPM from the source RPM and use that manual binary build in place of the scripted `hp_qla2x00src` RPM. You need to do this if your production servers do not have the kernel sources and `gcc` installed.

If you need to build a binary RPM to install, you will need a development machine with the same kernel as your targeted production servers. You can install the binary RPM-produced RPM methods on your production servers.

NOTE: The binary RPM that you build works only for the kernel and configuration that you build on (and possibly some errata kernels). Ensure that you use the 7.xx version of the `hp_qla2x00` source RPM for 2.4 kernel-based distributions and the 8.xx version of the `hp_qla2x00` source RPM for 2.6 kernel-based distributions.

Use the following procedure to create the binary RPM from the source RPM:

1. Select one of the following options:
   - Enter the **`#./INSTALL -s`** command. The binary RPM creation is complete. You do not have to perform 2 through 4.
   - Install the source RPM by issuing the `# rpm -ivh hp_qla2x00-`*`version-revision`*`.src.rpm` command. Continue with 2.
2. Select one of the following directories:
   - For Red Hat distributions, use the `/usr/src/redhat/SPECS` directory.
   - For SUSE distributions, use the `/usr/src/packages/SPECS` directory.
3. Build the RPM by using the `# rpmbuild -bb hp_qla2x00.spec` command.

   NOTE: In some of the older Linux distributions, the RPM command contains the RPM build functionality.

   At the end of the command output, the following message appears:

   ```
   "Wrote: ...rpm".
   ```

This line identifies the location of the binary RPM.

4. Copy the binary RPM to the production servers and install it using the following command:

   ```
   # rpm -ivh hp_qla2x00-version-revision.architecture.rpm
   ```

Verifying virtual disks from the host

To verify the virtual disks, first verify that the LUN is recognized and then verify that the host can access the virtual disks.

- To ensure that the LUN is recognized after a virtual disk is presented to the host, do one of the following:
  - Reboot the host.
  - Enter the `/opt/hp/hp_fibreutils/hp_rescan -a` command.
- To verify that the host can access the virtual disks, enter the **`# more /proc/scsi/scsi`** command.

  The output lists all SCSI devices detected by the server. An EVA4x00/6x00/8x00 LUN entry looks similar to the following:

  ```
  Host: scsi3 Channel: 00 ID: 00 Lun: 01
  Vendor: HP      Model: HSV300        Rev:
  Type: Direct-Access                      ANSI SCSI revision: 02
  ```

OpenVMS

Updating the AlphaServer console code, Integrity Server console code, and Fibre Channel FCA firmware

The firmware update procedure varies for the different server types. To update firmware, follow the procedure described in the Installation instructions that accompany the firmware images.

Verifying the Fibre Channel adapter software installation

A supported FCA should already be installed in the host server. The procedure to verify that the console recognizes the installed FCA varies for the different server types. Follow the procedure described in the Installation instructions that accompany the firmware images.

Console LUN ID and OS unit ID

HP P6000 Command View software contains a box for the Console LUN ID on the Initialized Storage System Properties window.

It is important that you set the Console LUN ID to a number other than zero (0). If the Console LUN ID is not set or is set to zero (0), the OpenVMS host will not recognize the controller pair. The Console LUN ID for a controller pair must be unique within the SAN. Table 29 (page 81) shows an example of the Console LUN ID.

You can set the OS unit ID on the Virtual Disk Properties window. The default setting is 0, which disables the ID field. To enable the ID field, you must specify a value between 1 and 32767, ensuring that the number you enter is unique within the SAN. An OS Unit ID greater than 9999 is not capable of being served by MSCP.

△ **CAUTION:** It is possible to enter a duplicate Console LUN ID or OS unit ID number. You must ensure that you enter a Console LUN ID and OS Unit ID that is not already in use. A duplicate Console LUN ID or OS Unit ID can allow the OpenVMS host to corrupt data due to confusion about LUN identity. It can also prevent the host from recognizing the controllers.

Table 29 Comparing console LUN to OS unit ID

ID type	System Display
Console LUN ID set to 100	1GGA100:
OS unit ID set to 50	1DGA50:

Adding OpenVMS hosts

To obtain WWNs on AlphaServers, do one of the following:

- Enter the **`show device fg/full`** OVMS command.
- Use the **`WWIDMGR -SHOW PORT`** command at the SRM console.

To obtain WWNs on Integrity servers, do one of the following:

1. Enter the **`show device fg/full`** OVMS command.
2. Use the following procedure from the server console:
   a. From the EFI boot Manager, select **EFI Shell**.
   b. In the EFI Shell, enter "`Shell> drivers`".
      A list of EFI drivers loaded in the system is displayed.
3. In the listing, find the line for the FCA for which you want to get the WWN information.
   For a Qlogic HBA, look for `HP 4 Gb Fibre Channel Driver` or `HP 2 Gb Fibre Channel Driver` as the driver name. For example:

```
            T   D
D           Y C I
R           P F A
V  VERSION  E G G #D #C DRIVER NAME                          IMAGE NAME
== ======== = = = == == ================================== ==================
22 00000105 B X X  1  1 HP 4 Gb Fibre Channel Driver        PciROM:0F:01:01:002
```

4. Note the driver handle in the first column (22 in the example).
5. Using the driver handle, enter the **`drvdfg`** *`driver_handle`* command to find the Device Handle (`Ctrl`). For example:

```
Shell> drvcfg 22
Configurable Components
Drv[22]  Ctrl[25]  Lang[eng]
```

6. Using the driver and device handle, enter the **`drvdfg –s`** *`driver_handle device_handle`* command to invoke the EFI Driver configuration utility. For example:

```
Shell> drvcfg -s 22 25
```

7. From the `Fibre Channel Driver Configuration Utility` list, select item 8 (**`Info`**) to find the WWN for that particular port.
   Output similar to the following appears:

```
Adapter Path:  Acpi(PNP0002,0300)/Pci(01|01)
Adapter WWPN:   50060B00003B478A
Adapter WWNN:   50060B00003B478B
Adapter S/N:    3B478A
```

Scanning the bus

Enter the following command to scan the bus for the OpenVMS virtual disk:

```
$ MC SYSMAN IO AUTO/LOG
```

A listing of LUNs detected by the scan process is displayed. Verify that the new LUNs appear on the list.

NOTE: The EVA4x00/6x00/8x00 console LUN can be seen without any virtual disks presented. The LUN appears as `$1$GGA`*x* (where *x* represents the console LUN ID on the controller).

After the system scans the fabric for devices, you can verify the devices with the `SHOW DEVICE` command:

`$ SHOW DEVICE` *`NAME-OF-VIRTUAL-DISK`*`/FULL`

For example, to display device information on a virtual disk named `$1$DGA50`, enter **`$ SHOW DEVICE $1$DGA50:/FULL`**.

The following output is displayed:

```
Disk $1$DGA50: (BRCK18), device type HSV210, is online, file-oriented device,
    shareable, device has multiple I/O paths, served to cluster via MSCP Server,
    error logging is enabled.

    Error count                    2    Operations completed              4107
    Owner process                 ""    Owner UIC                      [SYSTEM]
    Owner process ID        00000000    Dev Prot            S:RWPL,O:RWPL,G:R,W
    Reference count                0    Default buffer size                512
    Current preferred CPU Id       0    Fastpath                             1
    WWID   01000010:6005-08B4-0010-70C7-0001-2000-2E3E-0000
    Host name                "BRCK18"    Host type, avail AlphaServer DS10 466 MHz, yes
    Alternate host name       "VMS24"    Alt. type, avail HP rx3600  (1.59GHz/9.0MB), yes
    Allocation class               1

  I/O paths to device              9
  Path PGA0.5000-1FE1-0027-0A38 (BRCK18), primary path.
    Error count                    0    Operations completed               145
  Path PGA0.5000-1FE1-0027-0A3A (BRCK18).
    Error count                    0    Operations completed               338
  Path PGA0.5000-1FE1-0027-0A3E (BRCK18).
    Error count                    0    Operations completed               276
  Path PGA0.5000-1FE1-0027-0A3C (BRCK18).
    Error count                    0    Operations completed               282
  Path PGB0.5000-1FE1-0027-0A39 (BRCK18).
    Error count                    0    Operations completed               683
  Path PGB0.5000-1FE1-0027-0A3B (BRCK18).
    Error count                    0    Operations completed               704
  Path PGB0.5000-1FE1-0027-0A3D (BRCK18).
    Error count                    0    Operations completed               853
  Path PGB0.5000-1FE1-0027-0A3F (BRCK18), current path.
    Error count                    2    Operations completed               826
  Path MSCP  (VMS24).
    Error count                    0    Operations completed                 0
```

You can also use the `SHOW DEVICE DG` command to display a list of all Fibre Channel disks presented to the OpenVMS host.

NOTE: Restarting the host system shows any newly presented virtual disks because a hardware scan is performed as part of the startup.

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

- Check the switch zoning database.
- Use HP P6000 Command View to verify the host presentations.
- Check the SRM console firmware on AlphaServers.
- Ensure that the correct host is selected for this virtual disk and that a unique OS Unit ID is used in HP P6000 Command View.

Configuring virtual disks from the OpenVMS host

To set up disk resources under OpenVMS, initialize and mount the virtual disk resource as follows:

1. Enter the following command to initialize the virtual disk:

   `$ INITIALIZE` *`name-of-virtual-disk volume-label`*

2. Enter the following command to mount the disk:

   ```
   MOUNT/SYSTEM name-of-virtual-disk volume-label
   ```

   NOTE: The `/SYSTEM` switch is used for a single stand-alone system, or in clusters if you want to mount the disk only to select nodes. You can use the `/CLUSTER` switch for OpenVMS clusters. However, if you encounter problems in a large cluster environment, HP recommends that you enter a `MOUNT/SYSTEM` command on each cluster node.

3. View the virtual disk's information with the `SHOW DEVICE` command. For example, enter the following command sequence to configure a virtual disk named `data1` in a stand-alone environment:

   ```
   $ INIT $1$DGA1: data1
   $ MOUNT/SYSTEM $1$DGA1: data1
   $ SHOW DEV $1$DGA1: /FULL
   ```

Setting preferred paths

You can use one of the following options for setting, changing, or displaying preferred paths:

- To set or change the preferred path, use the following command:

  ```
  $ SET DEVICE $1$DGA83: /PATH=PGA0.5000-1FE1-0007-9772/SWITCH
  ```

  This allows you to control which path each virtual disk uses.

- To display the path identifiers, use the `SHOW DEV/FULL` command.
- For additional information on using OpenVMS commands, see the OpenVMS help file:

  ```
  $ HELP TOPIC
  ```

  For example, the following command displays help information for the `MOUNT` command:

  ```
  $ HELP MOUNT
  ```

Oracle Solaris

NOTE: The information in this section applies to both SPARC and x86 versions of the Oracle Solaris operating system.

Loading the operating system and software

Follow the manufacturer's instructions for loading the operating system (OS) and software onto the host. Load all OS patches and configuration utilities supported by HP and the FCA manufacturer.

Configuring FCAs with the Oracle SAN driver stack

Oracle-branded FCAs are supported only with the Oracle SAN driver stack. The Oracle SAN driver stack is also compatible with current Emulex FCAs and QLogic FCAs. Support information is available on the Sun website:

http://www.oracle.com/technetwork/server-storage/solaris/overview/index-136292.html

To determine which non-Oracle branded FCAs HP supports with the Oracle SAN driver stack, see the latest MPxIO application notes or contact your HP representative.

Update instructions depend on the version of your OS:

- For Solaris 9, install the latest Oracle StorEdge SAN software with associated patches. To locate the software, log into My Oracle Support:

  https://support.oracle.com/CSP/ui/flash.html

  1. Select the **Patches & Updates** tab and then search for **StorEdge SAN Foundation Software 4.4** (formerly called StorageTek SAN 4.4).
  2. Reboot the host after the required software/patches have been installed. No further activity is required after adding any new LUNs once the array ports have been configured with the `cfgadm -c` command for Solaris 9.

     Examples for two FCAs:

     ```
     cfgadm -c configure c3
     cfgadm -c configure c4
     ```

  3. Increase retry counts and reduce I/O time by adding the following entries to the `/etc/system` file:

     ```
     set ssd:ssd_retry_count=0xa
     set ssd:ssd_io_time=0x1e
     ```

  4. Reboot the system to load the newly added parameters.

- For Solaris 10, go the Oracle Software Downloads website (http://www.oracle.com/technetwork/indexes/downloads/index.html) to install the latest patches. Under Servers and Storage Systems, select **Solaris 10**. Reboot the host once the required software/patches have been installed. No further activity is required after adding any new LUNs, as the controller and LUN recognition are automatic for Solaris 10.
  1. For Solaris 10 x86/64, ensure patch 138889-03 or later is installed. For SPARC, ensure patch 138888-03 or later is installed.
  2. Increase the retry counts by adding the following line to the `/kernel/drv/sd.conf` file:

     ```
     sd-config-list="HP HSV","retries-timeout:10";
     ```

  3. Reduce the I/O timeout value to 30 seconds by adding the following line to the /etc/system file:

     ```
     set sd:sd_io_time=0x1e
     ```

  4. Reboot the system to load the newly added parameters.

Configuring Emulex FCAs with the lpfc driver

To configure Emulex FCAs with the lpfc driver:

1. Ensure that you have the latest supported version of the lpfc driver (see http://www.hp.com/storage/spock).

   You must sign up for an HP Passport to enable access. For more information on how to use SPOCK, see the Getting Started Guide (http://h20272.www2.hp.com/Pages/spock_overview/introduction.html).

2. Edit the following parameters in the `/kernel/drv/lpfc.conf` driver configuration file to set up the FCAs for a SAN infrastructure:

   ```
   topology=2;
   scan-down=0;
   nodev-tmo=60;
   linkdown-tmo=60;
   ```

3. If using a single FCA and no multipathing, edit the following parameter to reduce the risk of data loss in case of a controller reboot:

   ```
   nodev-tmo=120;
   ```

4. If using Veritas Volume Manager (VxVM) DMP for multipathing (single or multiple FCAs), edit the following parameter to ensure proper VxVM behavior:

   ```
   no-device-delay=0;
   ```

5. In a fabric topology, use persistent bindings to bind a SCSI target ID to the world wide port name (WWPN) of an array port. This ensures that the SCSI target IDs remain the same when the system reboots. Set persistent bindings by editing the configuration file or by using the `lputil` utility.

   NOTE: HP recommends that you assign target IDs in sequence, and that the EVA has the same target ID on each host in the SAN.

   The following example for an EVA4x00/6x00/8x00 illustrates the binding of targets 20 and 21 (lpfc instance 2) to WWPNs 50001fe100270938 and 50001fe100270939, and the binding of targets 30 and 31 (lpfc instance 0) to WWPNs 50001fe10027093a and 50001fe10027093b:

   ```
   fcp-bind-WWPN="50001fe100270938:lpfc2t20",
                 "50001fe100270939:lpfc2t21",
                 "50001fe10027093a:lpfc0t30",
                 "50001fe10027093b:lpfc0t31";
   ```

   NOTE: Replace the WWPNs in the example with the WWPNs of your array ports.

6. For each LUN that will be accessed, add an entry to the `/kernel/drv/sd.conf` file. For example, if you want to access LUNs 1 and 2 through all four paths, add the following entries to the end of the file:

   ```
   name="sd" parent="lpfc" target=20 lun=1;
   name="sd" parent="lpfc" target=21 lun=1;
   name="sd" parent="lpfc" target=30 lun=1;
   name="sd" parent="lpfc" target=31 lun=1;
   name="sd" parent="lpfc" target=20 lun=2;
   name="sd" parent="lpfc" target=21 lun=2;
   name="sd" parent="lpfc" target=30 lun=2;
   name="sd" parent="lpfc" target=31 lun=2;
   ```

7. Reboot the server to implement the changes to the configuration files.
8. If LUNs have been preconfigured in the `/kernel/drv/sd.conf` file, use the `devfsadm` command to perform LUN rediscovery after configuring the file.

   NOTE: The lpfc driver is *not* supported for Oracle StorEdge Traffic Manager/Oracle Storage Multipathing. To configure an Emulex FCA using the Oracle SAN driver stack, see "Configuring FCAs with the Oracle SAN driver stack" (page 83).

Configuring QLogic FCAs with the qla2300 driver

Check the Single Point of Connecitivty Knowledge (SPOCK) website or contact your HP representative to determine which QLogic FCAs and which driver version HP supports with the qla2300 driver. To configure QLogic FCAs with the qla2300 driver:

1. Ensure that you have the latest supported version of the qla2300 driver.

   You must sign up for an HP Passport to enable access. For more information on how to use SPOCK, see the Getting Started Guide (http://www.qlogic.com).

2. Edit the following parameters in the `/kernel/drv/qla2300.conf` driver configuration file to set up the FCAs for a SAN infrastructure (HBA0 is used in the example, but the parameter edits apply to all HBAs):

   NOTE: If you are using a Oracle-branded QLogic FCA, the configuration file is `\kernal\drv\qlc.conf`.

   ```
   hba0-connection-options=1;
   hba0-link-down-timeout=60;
   hba0-persistent-binding-configuration=1;
   ```

   NOTE: If you are using Solaris 10, editing the persistent binding parameter is not required.

3. If using a single FCA and no multipathing, edit the following parameters to reduce the risk of data loss in case of a controller reboot:

   ```
   hba0-login-retry-count=60;
   hba0-port-down-retry-count=60;
   hba0-port-down-retry-delay=2;
   ```

   The `hba0-port-down-retry-delay` parameter is *not* supported with the 4.13.01 driver; the time between retries is fixed at approximately 2 seconds.

4. In a fabric topology, use persistent bindings to bind a SCSI target ID to the world wide port name (WWPN) of an array port. This ensures that the SCSI target IDs remain the same when the system reboots. Set persistent bindings by editing the configuration file or by using the `SANsurfer` utility.

   NOTE: Persistent binding is not required for QLogic FCAs if you are using Solaris 10.

   The following example for an EVA4x00/6x00/8x00 illustrates the binding of targets 20 and 21 (hba instance 0) to WWPNs 50001fe100270938 and 50001fe100270939, and the binding of targets 30 and 31 (hba instance 1) to WWPNs 50001fe10027093a and 50001fe10027093b:

   ```
   hba0-SCSI-target-id-20-fibre-channel-port-name="50001fe100270938";
   hba0-SCSI-target-id-21-fibre-channel-port-name="50001fe10027093a";
   hba1-SCSI-target-id-30-fibre-channel-port-name="50001fe100270939";
   hba1-SCSI-target-id-31-fibre-channel-port-name="50001fe10027093b";
   ```

   NOTE: Replace the WWPNs in the example with the WWPNs of your array ports.

5. If the qla2300 driver is version 4.13.01 or earlier, for each LUN that users will access add an entry to the `/kernel/drv/sd.conf` file:

   ```
   name="sd" class="scsi" target=20 lun=1;
   name="sd" class="scsi" target=21 lun=1;
   name="sd" class="scsi" target=30 lun=1;
   name="sd" class="scsi" target=31 lun=1;
   ```

   If LUNs are preconfigured in the `/kernel/drv/sd.conf` file, after changing the configuration file. use the `devfsadm` command to perform LUN rediscovery.

6. If the qla2300 driver is version 4.15 or later, verify that the following or a similar entry is present in the `/kernel/drv/sd.conf` file:

   ```
   name="sd" parent="qla2300" target=2048;
   ```

   To perform LUN rediscovery after configuring the LUNs, use the following command:

   ```
   /opt/QLogic_Corporation/drvutil/qla2300/qlreconfig –d qla2300 -s
   ```

7. Reboot the server to implement the changes to the configuration files.

NOTE: The qla2300 driver is *not* supported for Oracle StorEdge Traffic Manager/Oracle Storage Multipathing. To configure a QLogic FCA using the Oracle SAN driver stack, see "Configuring FCAs with the Oracle SAN driver stack" (page 83).

Fabric setup and zoning

To set up the fabric and zoning:

1. Verify that the Fibre Channel cable is connected and firmly inserted at the array ports, host ports, and SAN switch.
2. Through the Telnet connection to the switch or switch utilities, verify that the WWN of the EVA ports and FCAs are present and online.
3. Create a zone consisting of the WWNs of the EVA ports and FCAs, and then add the zone to the active switch configuration.
4. Enable and then save the new active switch configuration.

NOTE: There are variations in the steps required to configure the switch between different vendors. For more information, see the *HP SAN Design Reference Guide*, available for downloading on the HP website: http://www.hp.com/go/sandesign.

Oracle StorEdge Traffic Manager (MPxIO)/Sun Storage Multipathing

Oracle StorEdge Traffic Manager (MPxIO)/Sun Storage Multipathing can be used for FCAs configured with the Oracle SAN driver and depending on the operating system version, architecture (SPARC/x86), and patch level installed. For configuration details, see the *HP MPxIO application notes*, available on the HP support website: http://www.hp.com/support/manuals.

NOTE: MPxIO is included in the SPARC and x86 Oracle SAN driver. A separate installation of MPxIO is not required.

In the `Search products` box, enter **`MPxIO`**, and then click the search symbol. Select the application notes from the search results.

Configuring with Veritas Volume Manager

The Dynamic Multipathing (DMP) feature of Veritas Volume Manager (VxVM) can be used for all FCAs and all drivers. EVA disk arrays are certified for VxVM support. When you install FCAs, ensure that the driver parameters are set correctly. Failure to do so can result in a loss of path failover in DMP. For information about setting FCA parameters, see "Configuring FCAs with the Oracle SAN driver stack" (page 83) and the FCA manufacturer's instructions.

The DMP feature requires an Array Support Library (ASL) and an Array Policy Module (APM). The ASL/APM enables Asymmetric Logical Unit Access (ALUA). LUNs are accessed through the primary controller. After enablement, use the `vxdisk list <device>` command to determine the primary and secondary paths. For VxVM 4.1 (MP1 or later), you must download the ASL/APM from the Symantec/Veritas support site for installation on the host. This download and installation is *not* required for VxVM 5.0 or later.

To download and install the ASL/APM from the Symantec/Veritas support website:

1. Go to http://support.veritas.com.
2. Enter **Storage Foundation for UNIX/Linux** in the Product Lookup box.
3. Enter **EVA** in the **Enter keywords or phrase** box, and then click the search symbol.
4. To further narrow the search, select **Solaris** in the Platform box and search again.
5. Read `TechNotes` and follow the instructions to download and install the ASL/APM.
6. Run `vxdctl enable` to notify VxVM of the changes.
7. Verify the configuration of VxVM as shown in Example 3 "Verifying the VxVM configuration" (the output may be slightly different depending on your VxVM version and the array configuration).

Example 3 Verifying the VxVM configuration

```
# vxddladm listsupport all | grep HP
libvxhpevale.so    HP  HSV200, HSV210

# vxddladm listsupport libname=libvxhpevale.so
ATTR_NAME            ATTR_VALUE
=======================================================================
LIBNAME              libvxhpevale.so
VID                  HP
PID                  HSV200, HSV210
ARRAY_TYPE           A/A-A-HP
ARRAY_NAME           EVA4K6K, EVA8000

# vxdmpadm listapm all | grep HP
dmphpalua          dmphpalua              1        A/A-A-HP          Active
# vxdmpadm listapm dmphpalua
Filename:               dmphpalua
APM name:               dmphpalua
APM version:            1
Feature:                VxVM
VxVM version:           41
Array Types Supported:  A/A-A-HP
Depending Array Types:  A/A-A
State:                  Active

# vxdmpadm listenclosure all
ENCLR_NAME        ENCLR_TYPE     ENCLR_SNO            STATUS       ARRAY_TYPE
==============================================================================
Disk              Disk           DISKS                CONNECTED    Disk
EVA8100           EVA8100      50001FE1002709E0       CONNECTED    A/A-A-HP
```

By default, the EVA I/O policy is set to `Round-Robin`. For VxVM 4.1 MP1, only one path is used for the I/Os with this policy. Therefore, HP recommends that you change the I/O policy to `Adaptive` in order to use all paths to the LUN on the primary controller. Example 4 "Setting the iopolicy" shows the commands you can use to check and change the I/O policy.

Example 4 Setting the iopolicy

```
# vxdmpadm getattr arrayname EVA8100 iopolicy
ENCLR_NAME     DEFAULT        CURRENT
============================================
EVA8100        Round-Robin    Round-Robin

# vxdmpadm setattr arrayname EVA8100 iopolicy=adaptive

# vxdmpadm getattr arrayname EVA8100 iopolicy
ENCLR_NAME     DEFAULT        CURRENT
============================================
EVA8100        Round-Robin    Adaptive
```

Configuring virtual disks from the host

The procedure used to configure the LUN path to the array depends on the FCA driver. For more information, see "Installing Fibre Channel adapters" (page 72).

To identify the WWLUN ID assigned to the virtual disk and/or the LUN assigned by the storage administrator:

- Oracle SAN driver, with MPxIO enabled:
  - You can use the `luxadm probe` command to display the array/node WWN and associated array for the devices.
  - The WWLUN ID is part of the device file name. For example:

    `/dev/rdsk/c5t`**`600508B4001030E40000500000B20000`**`d0s2`
  - If you use `luxadm display`, the LUN is displayed after the device address. For example:

    `50001fe1002709e9,`**`5`**
- Oracle SAN driver, without MPxIO:
  - The EVA WWPN is part of the file name (which helps you to identify the controller). For example:

    `/dev/rdsk/c3t`**`50001FE1002709E8`**`d5s2`
    `/dev/rdsk/c3t`**`50001FE1002709EC`**`d5s2`
    `/dev/rdsk/c4t`**`50001FE1002709E9`**`d5s2`
    `/dev/rdsk/c4t`**`50001FE1002709ED`**`d5s2`

    If you use `luxadm probe`, the array/node WWN and the associated device files are displayed.
  - You can retrieve the WWLUN ID as part of the `format -e` (scsi, inquiry) output; however, it is cumbersome and hard to read. For example:

    ```
    09 e8 20 04 00 00 00 00 00 00 35 30 30 30 31 46     .........50001F
    45 31 30 30 32 37 30 39 45 30 35 30 30 30 31 46     E1002709E050001F
    45 31 30 30 32 37 30 39 45 38 36 30 30 35 30 38     E1002709E8600508
    42 34 30 30 31 30 33 30 45 34 30 30 30 30 35 30     B4001030E4000050
    30 30 30 30 42 32 30 30 30 30 00 00 00 00 00 00     0000B20000
    ```
  - The assigned LUN is part of the device file name. For example:

    `/dev/rdsk/c3t50001FE1002709E8d`**`5`**`s2`

    You can also retrieve the LUN with `luxadm display`. The LUN is displayed after the device address. For example:

```
50001fe1002709e9,5
```

- Emulex (lpfc)/QLogic (qla2300) drivers:
  - You can retrieve the WWPN by checking the assignment in the driver configuration file (the easiest method, because you then know the assigned target) or by using `HBAnyware/SANSurfer`.
  - You can retrieve the WWLUN ID by using `HBAnyware/SANSurfer`.

    You can also retrieve the WWLUN ID as part of the `format -e` (scsi, inquiry) output; however, it is cumbersome and difficult to read. For example:

```
09 e8 20 04 00 00 00 00 00 00 35 30 30 30 31 46     .........50001F
45 31 30 30 32 37 30 39 45 30 35 30 30 30 31 46     E1002709E050001F
45 31 30 30 32 37 30 39 45 38 36 30 30 35 30 38     E1002709E8600508
42 34 30 30 31 30 33 30 45 34 30 30 30 30 35 30     B4001030E4000050
30 30 30 30 42 32 30 30 30 30 00 00 00 00 00 00     0000B20000
```

  - The assigned LUN is part of the device file name. For example:

    `/dev/dsk/c4t20d5s2`

Verifying virtual disks from the host

Verify that the host can access virtual disks by using the `format` command. See Example 5 “Format command ”.

Example 5 Format command

```
# format
Searching for disks...done
c2t50001FE1002709F8d1: configured with capacity of 1008.00MB
c2t50001FE1002709F8d2: configured with capacity of 1008.00MB
c2t50001FE1002709FCd1: configured with capacity of 1008.00MB
c2t50001FE1002709FCd2: configured with capacity of 1008.00MB
c3t50001FE1002709F9d1: configured with capacity of 1008.00MB
c3t50001FE1002709F9d2: configured with capacity of 1008.00MB
c3t50001FE1002709FDd1: configured with capacity of 1008.00MB
c3t50001FE1002709FDd2: configured with capacity of 1008.00MB

AVAILABLE DISK SELECTIONS:

0. c0t0d0 <SUN18G cyl 7506 alt 2 hd 19 sec 248> /pci@1f,4000/scsi@3/sd@0,0
1. c2t50001FE1002709F8d1 <HSV210-6240 cyl 126 alt 2 hd 128 sec 128>
     /pci@1f,4000/QLGC,qla@4/fp@0,0/ssd@w50001fe1002709f8,1
2. c2t50001FE1002709F8d2 <HSV210-6240 cyl 126 alt 2 hd 128 sec 128>
     /pci@1f,4000/QLGC,qla@4/fp@0,0/ssd@w50001fe1002709f8,2
3. c2t50001FE1002709FCd1 <HSV210-6240 cyl 126 alt 2 hd 128 sec 128>
     /pci@1f,4000/QLGC,qla@4/fp@0,0/ssd@w50001fe1002709fc,1
4. c2t50001FE1002709FCd2 <HSV210-6240 cyl 126 alt 2 hd 128 sec 128>
     /pci@1f,4000/QLGC,qla@4/fp@0,0/ssd@w50001fe1002709fc,2
5. c3t50001FE1002709F9d1 <HSV210-6240 cyl 126 alt 2 hd 128 sec 128>
     /pci@1f,4000/lpfc@5/fp@0,0/ssd@w50001fe1002709f9,1
6. c3t50001FE1002709F9d2 <HSV210-6240 cyl 126 alt 2 hd 128 sec 128>
     /pci@1f,4000/lpfc@5/fp@0,0/ssd@w50001fe1002709f9,2
7. c3t50001FE1002709FDd1 <HSV210-6240 cyl 126 alt 2 hd 128 sec 128>
     /pci@1f,4000/lpfc@5/fp@0,0/ssd@w50001fe1002709fd,1
8. c3t50001FE1002709FDd2 <HSV210-6240 cyl 126 alt 2 hd 128 sec 128>
     /pci@1f,4000/lpfc@5/fp@0,0/ssd@w50001fe1002709fd,2
Specify disk (enter its number):
```

If you cannot access the virtual disks:

- Verify the zoning.
- For Oracle Solaris, verify that the correct WWPNs for the EVA (`lpfc`, `qla2300` driver) have been configured and the target assignment is matched in `/kernel/drv/sd.conf` (`lpfc` and `qla2300 4.13.01`).

Labeling and partitioning the devices

Label and partition the new devices using the Sun format utility:

△ **CAUTION:** When selecting disk devices, be careful to select the correct disk because using the label/partition commands on disks that have data can cause data loss.

1. Enter the **`format`** command at the root prompt to start the utility.
2. Verify that all new devices are displayed. If not, enter `quit` or press **Ctrl+D** to exit the format utility, and then verify that the configuration is correct (see "Configuring virtual disks from the host" (page 89)).
3. Record the character-type device file names (for example, `c1t2d0`) for all new disks.

   You will use this data to create the file systems or to use the file systems with the Solaris or Veritas Volume Manager.
4. When prompted to specify the disk, enter the number of the device to be labeled.
5. When prompted to label the disk, enter **`Y`**.
6. Because the virtual geometry of the presented volume varies with size, select `autoconfigure` as the disk type.

7. For each new device, use the disk command to select another disk, and then repeat Step 1 through Step 6.
8. Repeat this labeling procedure for each new device. (Use the `disk` command to select another disk.)
9. When you finish labeling the disks, enter **`quit`** or press **Ctrl+D** to exit the format utility.

   For more information, see the *System Administration Guide: Devices and File Systems* for your operating system, available on the Oracle website:

   http://www.oracle/com/technetwork/indexes/documentation/index.html

NOTE: Some format commands are not applicable to the EVA storage systems.

VMware

Installing or upgrading VMware

For installation instructions, see the VMware installation guide for your server.

If you have already installed VMware, use the following procedure to patch or upgrade the system:

1. Extract the upgrade-tarball on the system. A sample command extract follows:

   `esx-n.n.n-14182-upgrade.tar.gz`
2. Boot the system in Linux mode by selecting the `Linux` boot option from the boot menu selection window.
3. Extract the tar file and enter the following command:

   `upgrade.pl`
4. Reboot the system using the default boot option (`esx`).

Configuring the EVA with VMware host servers

To configure an EVA4x00/6x00/8x00 on a VMware ESX server:

1. Using HP P6000 Command View, configure a host for one ESX server.
2. Verify that the Fibre Channel Adapters (FCAs) are populated in the world wide port name (WWPN) list. Edit the WWPN, if necessary.
3. Set the connection type to `VMware`.
4. To configure additional ports for the ESX server:
   a. Select a host (defined in Step 1).
   b. Select the **Ports** tab in the Host Properties window.
   c. Add additional ports for the ESX server.
5. Perform one of the following tasks to locate the WWPN:
   - From the service console, enter the **`wwpn.pl`** command.

     Output similar to the following is displayed:

     ```
     [root@gnome7 root]# wwpn.plvmhba0: 210000e08b09402b (QLogic)
     6:1:0vmhba1: 210000e08b0ace2d (QLogic) 6:2:0[root@gnome7 root]#
     ```
   - Check the SCSI device information section of the `/proc/scsi/qla2300/X` directory, where *X* is a bus instance number.

     Output similar to the following is displayed:

     ```
     SCSI Device Information:
     scsi-qla0-adapter-node=200000e08b0b0638;
     scsi-qla0-adapter-port=210000e08b0b0638;
     ```
6. Repeat this procedure for each ESX server.

Configuring an ESX server

This section provides information about configuring the ESX server.

Loading the FCA NVRAM

The FCA stores configuration information in the non-volatile RAM (NVRAM) cache. You must download the configuration for HP Storage products.

Perform one of the following procedures to load the NVRAM:

- If you have a ProLiant blade server:
  1. Download the supported FCA BIOS update, available on http://www.hp.com/support/downloads, to a virtual floppy.

     For instructions on creating and using a virtual floppy, see the *HP Integrated Lights-Out User Guide*.
  2. Unzip the file.
  3. Follow the instructions in the readme file to load the NVRAM configuration onto each FCA.
- If you have a blade server other than a ProLiant blade server:
  1. Download the supported FCA BIOS update, available on http://www.hp.com/support/downloads.
  2. Unzip the file.
  3. Follow the instructions in the readme file to load the NVRAM configuration onto each FCA.

Setting the multipathing policy

You can set the multipathing policy for each LUN or logical drive on the SAN to one of the following:

- Most recently used (MRU)
- Fixed
- Preferred

ESX 2.5.x commands

- The `# vmkmultipath –s vmhba0:0:1 –p mru` command sets `vmhba0:0:1` with an MRU multipathing policy for all LUNs on the SAN.
- The `# vmkmultipath -s vmhba1:0:1 -p fixed` command sets `vmhba1:0:1` with a Fixed multipathing policy.
- The `# vmkmultipath -s vmhba1:0:1 -r vmhba2:0:1 -e vmhba2:0:1` command sets and enables `vmhba2:0:1` with a Preferred multipathing policy.

ESX 3.x commands

- The `# esxcfg-mpath --policy=mru --lun=vmhba0:0:1` command sets `vmhba0:0:1` with an MRU multipathing policy.
- The `# esxcfg-mpath --policy=fixed --lun=vmhba0:0:1` command sets `vmhba1:0:1` with a Fixed multipathing policy.
- The `# esxcfg-mpath --preferred --path=vmhba2:0:1 --lun=vmhba2:0:1` command sets `vmhba2:0:1` with a Preferred multipathing policy.

ESX 4.x commands

- The `# esxcli nmp device setpolicy --device naa.6001438002a56f220001100000710000 --psp VMW_PSP_MRU` command sets `device naa.6001438002a56f220001100000710000` with an MRU multipathing policy.
- The `# esxcli nmp device setpolicy --device naa.6001438002a56f220001100000710000 --psp VMW_PSP_FIXED` command sets `device naa.6001438002a56f220001100000710000` with a Fixed multipathing policy.
- The `# esxcli nmp fixed setpreferred --device naa.6001438002a56f220001100000710000 --path vmhba1:C0:T2:L1` command sets `device naa.6001438002a56f220001100000710000` with a Preferred multipathing policy.

NOTE: Each LUN can be accessed through both EVA storage controllers at the same time; however, each LUN path is optimized through one controller. To optimize performance, if the LUN multipathing policy is Fixed, all servers must use a path to the same controller.

You can also set the multipathing policy from the VMware Management User Interface (MUI) by clicking the **Failover Paths** tab in the Storage Management section and then selecting **Edit... link** for each LUN whose policy you want to modify.

Specifying DiskMaxLUN

The `DiskMaxLUN` setting specifies the highest-numbered LUN that can be scanned by the ESX server.

- For ESX 2.5.x, the default value is 8. If more than eight LUNs are presented, you must change the setting to an appropriate value. To set `DiskMaxLUN`, select **Options> Advanced Settings** in the MUI, and then enter the highest-numbered LUN.
- For ESX 3.x or ESX 4.x, the default value is set to the Max set value of 256. To set `DiskMaxLun` to a different value, in Virtual Infrastructure Client, select **Configuration> Advance Settings> Disk> Disk.MaxLun**, and then enter the new value.

Verifying connectivity

To verify proper configuration and connectivity to the SAN:

- For ESX 2.5.x, enter the **`# vmkmultipath -q`** command.
- For ESX 3.x, enter the **`# esxcfg-mpath -l`** command.
- For ESX 4.x, enter the **`# esxcfg-mpath -b`** command.

For each LUN, verify that the multipathing policy is set correctly and that each path is marked `on`. If any paths are marked `dead` or are not listed, check the cable connections and perform a rescan on the appropriate FCA. For example:

- For ESX 2.5.x, enter the **`# cos-rescan.sh vmhba0`** command.
- For ESX 3.x or ESX 4.x, enter the **`# esxcfg-rescan vmhba0`** command.

If paths or LUNs are still missing, see the VMware or HP documentation for troubleshooting information.

Verifying virtual disks from the host

To verify that the host can access the virtual disks, enter the **`more /proc/scsi/scsi`** command.

The output lists all SCSI devices detected by the server. An EVA8100 LUN entry looks similar to the following:

```
Host: scsi3 Channel: 00 ID: 00 Lun: 01
Vendor: HP     Model: HSV210       Rev:
Type: Direct-Access                      ANSI SCSI revision: 02
```

5 Customer replaceable units

This chapter describes customer replaceable units. Information about initial enclosure installation, ESD protection, and common replacement procedures is also included.

Customer self repair (CSR)

Table 30 (page 97) identifies which hardware components are customer replaceable. Using WEBES, ISEE or other diagnostic tools, a support specialist will work with you to diagnose and assess whether a replacement component is required to address a system problem. The specialist will also help you determine whether you can perform the replacement.

Parts only warranty service

Your HP Limited Warranty may include a parts only warranty service. Under the terms of parts only warranty service, HP will provide replacement parts free of charge.

For parts only warranty service, CSR part replacement is mandatory. If you request HP to replace these parts, you will be charged for travel and labor costs.

Best practices for replacing hardware components

The following information will help you replace the hardware components on your storage system successfully.

CAUTION: Removing a component significantly changes the air flow within the enclosure. All components must be installed for the enclosure to cool properly. If a component fails, leave it in place in the enclosure until a new component is available to install.

Component replacement videos

To assist you in replacing the components, videos have been produced of the procedures. You can view the videos at the following website.

http://www.hp.com/go/sml

Verifying component failure

- Consult HP technical support to verify that the hardware component has failed and that you are authorized to replace it yourself.
- Additional hardware failures can complicate component replacement. Check HP P6000 Command View and/or HP remote support software as follows to detect any additional hardware problems:
  - When you have confirmed that a component replacement is required, you may want to clear the Real Time Monitoring view. This makes it easier to identify additional hardware problems that may occur while waiting for the replacement part.
  - Before installing the replacement part, check the Real Time Monitoring view for any new hardware problems. If additional hardware problems have occurred, contact HP support before replacing the component.
  - See the HP remote support software online help for additional information.

Procuring the spare part

Parts have a nine-character spare component number on their label (Figure 37 (page 97)). The first six characters (123479) identify the element; the last three characters (002) define the revision

level. The replacement component revision level must be the same as, or greater than, the number on the element being replaced. The higher the revision level, the later the revision.

Figure 37 Typical product label

8053A-ST

The spare part number for each disk drive is listed on the capacity label attached to each drive. See Figure 38 (page 97).

Figure 38 Disk drive label

Replaceable parts

This product contains the replaceable parts listed in Table 30 (page 97). Parts that are available for customer self repair (CSR) are indicated as follows:

✓ Mandatory CSR as enabled by XCS 6.000 or later and HP Command View EVA 6.0.2 or later and where geography permits. Order the part directly from HP and repair the product yourself. On-site or return-to-depot repair is not provided under warranty.

• Optional CSR. You can order the part directly from HP and repair the product yourself, or you can request that HP repair the product. If you request repair from HP, you may be charged for the repair depending on the product warranty.

– No CSR. The replaceable part is not available for self repair. For assistance, contact an HP-authorized service provider.

Table 30 Hardware component CSR support

Description	Spare part number (non RoHS/RoHS)	CSR
Cache battery (non-CSR)	30-10013-S1/30-10013-T1	–
Controller blower	390852-001,12-10008-S1 or 390852-005,12-10008-T1 (all RoHS)	✓
Controller power supply	349800-001/406442-001	✓
Disk enclosure blower	123482-001,70-40085-S1/ 123482-005,70-40085-T1	•
Disk enclosure power supply	212398-001 or 30-50872-S1	•
Disk drive – 72 GB 10K	244448-001/244448-002	✓
Disk drive – 146 GB 10K	300590-001/300590-002	✓

Table 30 Hardware component CSR support *(continued)*

Description	Spare part number (non RoHS/RoHS)	CSR
Disk drive – 300 GB 10K	366023-001/366023-002	✓
Disk drive – 450 GB 10K	518736-001	✓
Disk drive – 600 GB 10K	518737-001	✓
Disk drive – 72 GB 15K	300588-001/300588-002	✓
Disk drive – 146 GB 15K	366024-001/366024-002	✓
Disk drive – 300 GB 15K	416728-001	✓
Disk drive – 450 GB 15K	454415-001	✓
Disk drive – 600 GB 15K	531995-001	✓
Disk drive – 250 GB FATA	366022-001/366022-002	✓
Disk drive – 400 GB FATA	382262-001	✓
Disk drive – 500 GB FATA	371142-001	✓
Disk drive – 1 TB. 7.2K, FATA	454416-001	✓
Chassis with backplane	408515-001, 70-41260-T1 (both RoHS)	--
Operator control panel	390859-001,70-41143-S1 or 390859-005,70-41143-T1	✓
EMU	375393-001, 70-40145-S2/375393-005, 70-40145-T2	--
FC I/O module A	364549-005,70-40616-T4 or 364549-009,70-40616-T5 (all RoHS)	--
FC I/O module B	364548-005,70-40615-T4 or 364548-009,70-40615-T5 (all RoHS)	--
2GB SFP FC copper cable assembly, 0.6M	372631-001	•
2GB SFP FC copper cable assembly, 2.0M	372630-001	•
4GB SFP FC copper cable assembly, 0.6M	17-05405-S2 (RoHS)	•
4GB SFP FC copper cable assembly, 2.0M	17-05405-S1 (RoHS)	•
2GB SFP Transceiver	229204-001 or 416729-001 (both RoHS)	•
4GB SFP Transceiver	416729-001 (RoHS)	•
Controller EVA8000	390855-001, 70-41138-S1/ 390855-005, 70-41138-T1	--
Controller EVA8100	390855-006, 70-41138-Y1 (both RoHS)	--
Controller EVA6000/4000	390856-001,70-41138-S2/ 390856-005,70-41138-T2	--
Controller EVA6100/4100	390856-006, 70-41138-Y2 (both RoHS)	--

Table 30 Hardware component CSR support *(continued)*

Description	Spare part number (non RoHS/RoHS)	CSR
Front panel bezel EVA8000	390853-001, 70-41140-S1/ 411632-005, 70-41140-S3	✓
Front panel bezel EVA8100	390854-001, 70-41140-S2/ 411632-006, 70-41140-S5	✓
Front panel bezel EVA4000/6000	411633-005, 70-41140-S4 (both RoHS)	✓
Front panel bezel EVA4100/6100	411633-006, 70-41140-S6 (both RoHS)	✓

For more information about CSR, contact your local service provider. For North America, see the CSR website:

http://www.hp.com/go/selfrepair

To determine the warranty service provided for this product, see the warranty information website:

http://www.hp.com/go/storagewarranty

To order a replacement part, contact an HP-authorized service provider or see the HP Parts Store online:

http://www.hp.com/buy/parts

Replacing the failed component

△ **CAUTION:** Components can be damaged by electrostatic discharge. Use proper anti-static protection.

- Always transport and store CRUs in an ESD protective enclosure.
- Do not remove the CRU from the ESD protective enclosure until you are ready to install it.
- Always use ESD precautions, such as a wrist strap, heel straps on conductive flooring, and an ESD protective smock when handling ESD sensitive equipment.
- Avoid touching the CRU connector pins, leads, or circuitry.
- Do not place ESD generating material such as paper or non anti-static (pink) plastic in an ESD protective enclosure with ESD sensitive equipment.

- HP recommends waiting until periods of low storage system activity to replace a component.
- When replacing components at the rear of the rack, cabling may obstruct access to the component. Carefully move any cables out of the way to avoid loosening any connections. In particular, avoid cable damage that may be caused by:
  - Kinking or bending.
  - Disconnecting cables without capping. If uncapped, cable performance may be impaired by contact with dust, metal or other surfaces.
  - Placing removed cables on the floor or other surfaces, where they may be walked on or otherwise compressed.
- Replacement procedures are provided with each component. You can also download the following replacement instructions from the Manuals page of the HP Business Support Center website. See "Related information" (page 101) for more information.
  - *HP Controller Blower Replacement Instructions*
  - *HP Controller Enclosure Cache Battery Replacement Instructions*

- *HP Controller Power Supply Replacement Instructions*
- *HP Disk Enclosure Power Supply/Blower Replacement Instructions*
- *HP Fibre Channel Disk Drive Replacement Instructions*
- *HP Operator Control Panel Replacement Instructions*

Returning the defective part

In the materials shipped with a replacement CSR part, HP specifies whether the defective component must be returned to HP. Where required, you must ship the defective component back to HP within a defined period of time, normally five (5) business days. The defective component must be returned with the associated documentation provided in the shipping material. Failure to return the defective component may result in HP billing you for the replacement. With a customer self repair, HP will pay all shipping and component return costs and determine the courier/carrier to be used.

6 Support and other resources

Contacting HP

For worldwide technical support information, see the HP support website:

http://www.hp.com/support

Before contacting HP, collect the following information:

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

Subscription service

HP recommends that you register your product at the Subscriber's Choice for Business website:

http://www.hp.com/go/e-updates

After registering, you will receive e-mail notification of product enhancements, new driver versions, firmware updates, and other product resources.

Documentation feedback

HP welcomes your feedback.

To make comments and suggestions about product documentation, please send a message to **storagedocsfeedback@hp.com**. All submissions become the property of HP.

Related information

Documents

You can find related documents from the Manuals page of the HP Business Support Center website:

http://www.hp.com/support/manuals

Click **Disk Storage Systems** or **Storage Software** under storage, and then select your product.

Websites

For additional information, see the following HP websites:

- http://www.hp.com
- http://www.hp.com/go/storage
- http://www.hp.com/service_locator
- http://www.hp.com/support/manuals
- http://www.software.hp.com

Document conventions and symbols

Table 31 Document conventions

Convention	Element
Blue text: Table 31 (page 102)	Cross-reference links and e-mail addresses
Blue, underlined text: http://www.hp.com	website addresses
Bold text	• Keys that are pressed • Text typed into a GUI element, such as a box • GUI elements that are clicked or selected, such as menu and list items, buttons, tabs, and check boxes
Italic text	Text emphasis
`Monospace` text	• File and directory names • System output • Code • Commands, their arguments, and argument values
`Monospace, italic` text	• Code variables • Command variables
`Monospace, bold` text	Emphasized monospace text

WARNING! Indicates that failure to follow directions could result in bodily harm or death.

CAUTION: Indicates that failure to follow directions could result in damage to equipment or data.

IMPORTANT: Provides clarifying information or specific instructions.

NOTE: Provides additional information.

TIP: Provides helpful hints and shortcuts.

Rack stability

WARNING! To reduce the risk of personal injury or damage to equipment:

- Extend leveling jacks to the floor.
- Ensure that the full weight of the rack rests on the leveling jacks.
- Install stabilizing feet on the rack.
- In multiple-rack installations, secure racks together.
- Extend only one rack component at a time. Racks may become unstable if more than one component is extended.

Customer self repair

HP customer self repair (CSR) programs allow you to repair your product. If a CSR part needs replacing, HP ships the part directly to you so that you can install it at your convenience. Some parts do not qualify for CSR. Your HP-authorized service provider will determine whether a repair can be accomplished by CSR.

For more information about CSR, contact your local service provider. For North America, see the CSR website:

http://www.hp.com/go/selfrepair

A Regulatory notices and specifications

This appendix includes regulatory notices and product specifications for the HP Enterprise Virtual Array family.

Regulatory notices

Federal Communications Commission (FCC) notice

Part 15 of the Federal Communications Commission (FCC) Rules and Regulations has established Radio Frequency (RF) emission limits to provide an interference-free radio frequency spectrum. Many electronic devices, including computers, generate RF energy incidental to their intended function and are, therefore, covered by these rules. These rules place computers and related peripheral devices into two classes, A and B, depending upon their intended installation. Class A devices are those that may reasonably be expected to be installed in a business or commercial environment. Class B devices are those that may reasonably be expected to be installed in a residential environment (for example, personal computers). The FCC requires devices in both classes to bear a label indicating the interference potential of the device as well as additional operating instructions for the user.

The rating label on the device shows the classification (A or B) of the equipment. Class B devices have an FCC logo or FCC ID on the label. Class A devices do not have an FCC logo or FCC ID on the label. After the class of the device is determined, see the corresponding statement in the following sections.

FCC Class A certification

This equipment generates, uses, and may emit radio frequency energy. The equipment has been type tested and found to comply with the limits for a Class A digital device pursuant to Part 15 of the FCC rules, which are designed to provide reasonable protection against such radio frequency interference.

Operation of this equipment in a residential area may cause interference, in which case the user at the user's own expense will be required to take whatever measures may be required to correct the interference.

Any modifications to this device—unless approved by the manufacturer—can void the user's authority to operate this equipment under Part 15 of the FCC rules.

NOTE: Additional information on the need to interconnect the device with shielded (data) cables or the need for special devices, such as ferrite beads on cables, is required if such means of interference suppression was used in the qualification test for the device. This information will vary from device to device and needs to be obtained from the HP EMC group.

Class A equipment

This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at personal expense.

Class B equipment

This equipment has been tested and found to comply with the limits for a Class B digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. However, there is no guarantee that

interference will not occur in a particular installation. If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Connect the equipment into an outlet on a circuit that is different from that to which the receiver is connected.
- Consult the dealer or an experienced radio or television technician for help.

Declaration of conformity for products marked with the FCC logo, United States only

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) this device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

For questions regarding your product, see http://thenew.hp.com.

For questions regarding this FCC declaration, contact:

- Hewlett-Packard Company Product Regulations Manager, 3000 Hanover St., Palo Alto, CA 94304
- Or call 1-650-857-1501

To identify this product, see the part, series, or model number found on the product.

Modifications

The FCC requires the user to be notified that any changes or modifications made to this device that are not expressly approved by Hewlett-Packard Company may void the user's authority to operate the equipment.

Cables

Connections to this device must be made with shielded cables with metallic RFI/EMI connector hoods in order to maintain compliance with FCC Rules and Regulations.

Laser device

All Hewlett-Packard systems equipped with a laser device comply with safety standards, including International Electrotechnical Commission (IEC) 825. With specific regard to the laser, the equipment complies with laser product performance standards set by government agencies as a Class 1 laser product. The product does not emit hazardous light; the beam is totally enclosed during all modes of customer operation and maintenance.

Laser safety warnings

Heed the following warning:

⚠ **WARNING!** **WARNING**: To reduce the risk of exposure to hazardous radiation:

• Do not try to open the laser device enclosure. There are no user-serviceable components inside.

• Do not operate controls, make adjustments, or perform procedures to the laser device other than those specified herein.

• Allow only HP authorized service technicians to repair the laser device.

Compliance with CDRH regulations

The Center for Devices and Radiological Health (CDRH) of the U.S. Food and Drug Administration implemented regulations for laser products on August 2, 1976. These regulations apply to laser products manufactured from August 1, 1976. Compliance is mandatory for products marketed in the United States.

Certification and classification information

This product contains a laser internal to the Optical Link Module (OLM) for connection to the Fibre communications port.

In the USA, the OLM is certified as a Class 1 laser product conforming to the requirements contained in the Department of Health and Human Services (DHHS) regulation 21 CFR, Subchapter J. The certification is indicated by a label on the plastic OLM housing.

Outside the USA, the OLM is certified as a Class 1 laser product conforming to the requirements contained in IEC 825-1:1993 and EN 60825-1:1994, including Amendment 11:1996.

The OLM includes the following certifications:

- UL Recognized Component (USA)
- CSA Certified Component (Canada)
- TUV Certified Component (European Union)
- CB Certificate (Worldwide)

Canadien notice (avis Canadien)

Class A equipment

This Class A digital apparatus meets all requirements of the Canadian Interference-Causing Equipment Regulations.

Cet appareil numérique de la classe A respecte toutes les exigences du Règlement sur le matériel brouilleur du Canada.

Class B equipment

This Class B digital apparatus meets all requirements of the Canadian Interference-Causing Equipment Regulations.

Cet appareil numérique de la classe B respecte toutes les exigences du Règlement sur le matériel brouilleur du Canada.

European union notice

Products with the CE Marking comply with both the EMC Directive (89/336/EEC) and the Low Voltage Directive (73/23/EEC) issued by the Commission of the European Community.

Compliance with these directives implies conformity to the following European Norms (the equivalent international standards are in parenthesis):

- EN55022 (CISPR 22) - Electromagnetic Interference
- EN55024 (IEC61000-4-2, 3, 4, 5, 6, 8, 11) - Electromagnetic Immunity
- EN61000-3-2 (IEC61000-3-2) - Power Line Harmonics
- EN61000-3-3 (IEC61000-3-3) - Power Line Flicker
- EN60950 (IEC950) - Product Safety

Notice for France

DECLARATION D'INSTALLATION ET DE MISE EN EXPLOITATION d'un matériel de traitement de l'information (ATI), classé A en fonction des niveaux de perturbations radioélectriques émis, définis dans la norme européenne EN 55022 concernant la Compatibilité Electromagnétique.

WEEE Recycling Notices

English notice

Disposal of waste equipment by users in private household in the European Union

This symbol on the product or on its packaging indicates that this product must not be disposed of with your other household waste. Instead, it is your responsibility to dispose of your waste equipment by handing it over to a designated collection point for recycling of waste electrical and electronic equipment. The separate collection and recycling of your waste equipment at the time of disposal will help to conserve natural resources and ensure that it is recycled in a manner that protects human health and the environment. For more information about where you can drop off your waste equipment for recycling, please contact your local city office, your household waste disposal service, or the shop where you purchased the product.

Dutch notice

Verwijdering van afgedankte apparatuur door privé-gebruikers in de Europese Unie

Dit symbool op het product of de verpakking geeft aan dat dit product niet mag worden gedeponeerd bij het normale huishoudelijke afval. U bent zelf verantwoordelijk voor het inleveren van uw afgedankte apparatuur bij een inzamelingspunt voor het recyclen van oude elektrische en elektronische apparatuur. Door uw oude apparatuur apart aan te bieden en te recyclen, kunnen natuurlijke bronnen worden behouden en kan het materiaal worden hergebruikt op een manier waarmee de volksgezondheid en het milieu worden beschermd. Neem contact op met uw gemeente, het afvalinzamelingsbedrijf of de winkel waar u het product hebt gekocht voor meer informatie over inzamelingspunten waar u oude apparatuur kunt aanbieden voor recycling.

Czechoslovakian notice

Likvidace zařízení soukromými domácími uživateli v Evropské unii

Tento symbol na produktu nebo balení označuje výrobek, který nesmí být vyhozen spolu s ostatním domácím odpadem. Povinností uživatele je předat takto označený odpad na předem určené sběrné místo pro recyklaci elektrických a elektronických zařízení. Okamžité třídění a recyklace odpadu pomůže uchovat přírodní prostředí a zajistí takový způsob recyklace, který ochrání zdraví a životní prostředí člověka. Další informace o možnostech odevzdání odpadu k recyklaci získáte na příslušném obecním nebo městském úřadě, od firmy zabývající se sběrem a svozem odpadu nebo v obchodě, kde jste produkt zakoupili.

Estonian notice

Seadmete jäätmete kõrvaldamine eramajapidamistes Euroopa Liidus

See tootel või selle pakendil olev sümbol näitab, et kõnealust toodet ei tohi koos teiste majapidamisjäätmetega kõrvaldada. Teie kohus on oma seadmete jäätmed kõrvaldada, viies need elektri- ja elektroonikaseadmete jäätmete ringlussevõtmiseks selleks ettenähtud kogumispunkti. Seadmete jäätmete eraldi kogumine ja ringlussevõtmine kõrvaldamise ajal aitab kaitsta loodusvarasid ning tagada, et ringlussevõtmine toimub viisil, mis kaitseb inimeste tervist ning keskkonda. Lisateabe saamiseks selle kohta, kuhu oma seadmete jäätmed ringlussevõtmiseks viia, võtke palun ühendust oma kohaliku linnakantselei, majapidamisjäätmete kõrvaldamise teenistuse või kauplusega, kust Te toote ostsite.

Finnish notice

Laitteiden hävittäminen kotitalouksissa Euroopan unionin alueella

Jos tuotteessa tai sen pakkauksessa on tämä merkki, tuotetta ei saa hävittää kotitalousjätteiden mukana. Tällöin hävitettävä laite on toimitettava sähkölaitteiden ja elektronisten laitteiden kierrätyspisteeseen. Hävitettävien laitteiden erillinen käsittely ja kierrätys auttavat säästämään

luonnonvaroja ja varmistamaan, että laite kierrätetään tavalla, joka estää terveyshaitat ja suojelee luontoa. Lisätietoja paikoista, joihin hävitettävät laitteet voi toimittaa kierrätettäväksi, saa ottamalla yhteyttä jätehuoltoon tai liikkeeseen, josta tuote on ostettu.

French notice

Élimination des appareils mis au rebut par les ménages dans l'Union européenne

Le symbole apposé sur ce produit ou sur son emballage indique que ce produit ne doit pas être jeté avec les déchets ménagers ordinaires. Il est de votre responsabilité de mettre au rebut vos appareils en les déposant dans les centres de collecte publique désignés pour le recyclage des équipements électriques et électroniques. La collecte et le recyclage de vos appareils mis au rebut indépendamment du reste des déchets contribue à la préservation des ressources naturelles et garantit que ces appareils seront recyclés dans le respect de la santé humaine et de l'environnement. Pour obtenir plus d'informations sur les centres de collecte et de recyclage des appareils mis au rebut, veuillez contacter les autorités locales de votre région, les services de collecte des ordures ménagères ou le magasin dans lequel vous avez acheté ce produit.

German notice

Entsorgung von Altgeräten aus privaten Haushalten in der EU

Das Symbol auf dem Produkt oder seiner Verpackung weist darauf hin, dass das Produkt nicht über den normalen Hausmüll entsorgt werden darf. Benutzer sind verpflichtet, die Altgeräte an einer Rücknahmestelle für Elektro- und Elektronik-Altgeräte abzugeben. Die getrennte Sammlung und ordnungsgemäße Entsorgung Ihrer Altgeräte trägt zur Erhaltung der natürlichen Ressourcen bei und garantiert eine Wiederverwertung, die die Gesundheit des Menschen und die Umwelt schützt. Informationen dazu, wo Sie Rücknahmestellen für Ihre Altgeräte finden, erhalten Sie bei Ihrer Stadtverwaltung, den örtlichen Müllentsorgungsbetrieben oder im Geschäft, in dem Sie das Gerät erworben haben.

Greek notice

Απόρριψη άχρηστου εξοπλισμού από χρήστες σε ιδιωτικά νοικοκυριά στην Ευρωπαϊκή Ένωση

Το σύμβολο αυτό στο προϊόν ή τη συσκευασία του υποδεικνύει ότι το συγκεκριμένο προϊόν δεν πρέπει να διατίθεται μαζί με τα άλλα οικιακά σας απορρίμματα. Αντίθετα, είναι δική σας ευθύνη να απορρίψετε τον άχρηστο εξοπλισμό σας παραδίδοντάς τον σε καθορισμένο σημείο συλλογής για την ανακύκλωση άχρηστου ηλεκτρικού και ηλεκτρονικού εξοπλισμού. Η ξεχωριστή συλλογή και ανακύκλωση του άχρηστου εξοπλισμού σας κατά την απόρριψη θα συμβάλει στη διατήρηση των φυσικών πόρων και θα διασφαλίσει ότι η ανακύκλωση γίνεται με τρόπο που προστατεύει την ανθρώπινη υγεία και το περιβάλλον. Για περισσότερες πληροφορίες σχετικά με το πού μπορείτε να παραδώσετε τον άχρηστο εξοπλισμό σας για ανακύκλωση, επικοινωνήστε με το αρμόδιο τοπικό γραφείο, την τοπική υπηρεσία διάθεσης οικιακών απορριμμάτων ή το κατάστημα όπου αγοράσατε το προϊόν.

Απόρριψη άχρηστου εξοπλισμού από χρήστες σε ιδιωτικά νοικοκυριά στην Ευρωπαϊκή Ένωση

Το σύμβολο αυτό στο προϊόν ή τη συσκευασία του υποδεικνύει ότι το συγκεκριμένο προϊόν δεν πρέπει να διατίθεται μαζί με τα άλλα οικιακά σας απορρίμματα. Αντίθετα, είναι δική σας ευθύνη να απορρίψετε τον άχρηστο εξοπλισμό σας παραδίδοντάς τον σε καθορισμένο σημείο συλλογής για την ανακύκλωση άχρηστου ηλεκτρικού και ηλεκτρονικού εξοπλισμού. Η ξεχωριστή συλλογή και ανακύκλωση του άχρηστου εξοπλισμού σας κατά την απόρριψη θα συμβάλει στη διατήρηση των φυσικών πόρων και θα διασφαλίσει ότι η ανακύκλωση γίνεται με τρόπο που προστατεύει την ανθρώπινη υγεία και το περιβάλλον. Για περισσότερες πληροφορίες σχετικά με το πού μπορείτε να παραδώσετε

τον άχρηστο εξοπλισμό σας για ανακύκλωση, επικοινωνήστε με το αρμόδιο τοπικό γραφείο, την τοπική υπηρεσία διάθεσης οικιακών απορριμμάτων ή το κατάστημα όπου αγοράσατε το προϊόν.

Hungarian notice

Készülékek magánháztartásban történő selejtezése az Európai Unió területén

A készüléken, illetve a készülék csomagolásán látható azonos szimbólum annak jelzésére szolgál, hogy a készülék a selejtezés során az egyéb háztartási hulladéktól eltérő módon kezelendő. A vásárló a hulladékká vált készüléket köteles a kijelölt gyűjtőhelyre szállítani az elektromos és elektronikai készülékek újrahasznosítása céljából. A hulladékká vált készülékek selejtezéskori begyűjtése és újrahasznosítása hozzájárul a természeti erőforrások megőrzéséhez, valamint biztosítja a selejtezett termékek környezetre és emberi egészségre nézve biztonságos feldolgozását. A begyűjtés pontos helyéről bővebb tájékoztatást a lakhelye szerint illetékes önkormányzattól, az illetékes szemételtakarító vállalattól, illetve a terméket elárusító helyen kaphat.

Italian notice

Smaltimento delle apparecchiature da parte di privati nel territorio dell'Unione Europea

Questo simbolo presente sul prodotto o sulla sua confezione indica che il prodotto non può essere smaltito insieme ai rifiuti domestici. È responsabilità dell'utente smaltire le apparecchiature consegnandole presso un punto di raccolta designato al riciclo e allo smaltimento di apparecchiature elettriche ed elettroniche. La raccolta differenziata e il corretto riciclo delle apparecchiature da smaltire permette di proteggere la salute degli individui e l'ecosistema. Per ulteriori informazioni relative ai punti di raccolta delle apparecchiature, contattare l'ente locale per lo smaltimento dei rifiuti, oppure il negozio presso il quale è stato acquistato il prodotto.

Latvian notice

Nolietotu iekārtu iznīcināšanas noteikumi lietotājiem Eiropas Savienības privātajās mājsaimniecībās

Šāds simbols uz izstrādājuma vai uz tā iesaiņojuma norāda, ka šo izstrādājumu nedrīkst izmest kopā ar citiem sadzīves atkritumiem. Jūs atbildat par to, lai nolietotās iekārtas tiktu nodotas speciāli iekārtotos punktos, kas paredzēti izmantoto elektrisko un elektronisko iekārtu savākšanai otrreizējai pārstrādei. Atsevišķa nolietoto iekārtu savākšana un otrreizējā pārstrāde palīdzēs saglabā dabas resursus un garantēs, ka šīs iekārtas tiks otrreizēji pārstrādātas tādā veidā, lai pasargātu vidi un cilvēku veselību. Lai uzzinātu, kur nolietotās iekārtas var izmest otrreizējai pārstrādei, jāvēršas savas dzīves vietas pašvaldībā, sadzīves atkritumu savākšanas dienestā vai veikalā, kurā izstrādājums tika nopirkts.

Lithuanian notice

Vartotojų iš privačių namų ūkių įrangos atliekų šalinimas Europos Sąjungoje

Šis simbolis ant gaminio arba jo pakuotės rodo, kad šio gaminio šalinti kartu su kitomis namų ūkio atliekomis negalima. Šalintinas įrangos atliekas privalote pristatyti į specialią surinkimo vietą elektros ir elektroninės įrangos atliekoms perdirbti. Atskirai surenkamos ir perdirbamos šalintinos įrangos atliekos padės saugoti gamtinius išteklius ir užtikrinti, kad jos bus perdirbtos tokiu būdu, kuris nekenkia žmonių sveikatai ir aplinkai. Jeigu norite sužinoti daugiau apie tai, kur galima pristatyti perdirbtinas įrangos atliekas, kreipkitės į savo seniūniją, namų ūkio atliekų šalinimo tarnybą arba parduotuvę, kurioje įsigijote gaminį.

Polish notice

Pozbywanie się zużytego sprzętu przez użytkowników w prywatnych gospodarstwach domowych w Unii Europejskiej

Ten symbol na produkcie lub jego opakowaniu oznacza, że produktu nie wolno wyrzucać do zwykłych pojemników na śmieci. Obowiązkiem użytkownika jest przekazanie zużytego sprzętu do wyznaczonego punktu zbiórki w celu recyklingu odpadów powstałych ze sprzętu elektrycznego i elektronicznego. Osobna zbiórka oraz recykling zużytego sprzętu pomogą w ochronie zasobów naturalnych i zapewnią ponowne wprowadzenie go do obiegu w sposób chroniący zdrowie człowieka i środowisko. Aby uzyskać więcej informacji o tym, gdzie można przekazać zużyty sprzęt do recyklingu, należy się skontaktować z urzędem miasta, zakładem gospodarki odpadami lub sklepem, w którym zakupiono produkt.

Portuguese notice

Descarte de Lixo Elétrico na Comunidade Européia

Este símbolo encontrado no produto ou na embalagem indica que o produto não deve ser descartado no lixo doméstico comum. É responsabilidade do cliente descartar o material usado (lixo elétrico), encaminhando-o para um ponto de coleta para reciclagem. A coleta e a reciclagem seletivas desse tipo de lixo ajudarão a conservar as reservas naturais; sendo assim, a reciclagem será feita de uma forma segura, protegendo o ambiente e a saúde das pessoas. Para obter mais informações sobre locais que reciclam esse tipo de material, entre em contato com o escritório da HP em sua cidade, com o serviço de coleta de lixo ou com a loja em que o produto foi adquirido.

Slovakian notice

Likvidácia vyradených zariadení v domácnostiach v Európskej únii

Symbol na výrobku alebo jeho balení označuje, že daný výrobok sa nesmie likvidovať s domovým odpadom. Povinnosťou spotrebiteľa je odovzdať vyradené zariadenie v zbernom mieste, ktoré je určené na recykláciu vyradených elektrických a elektronických zariadení. Separovaný zber a recyklácia vyradených zariadení prispieva k ochrane prírodných zdrojov a zabezpečuje, že recyklácia sa vykonáva spôsobom chrániacim ľudské zdravie a životné prostredie. Informácie o zberných miestach na recykláciu vyradených zariadení vám poskytne miestne zastupiteľstvo, spoločnosť zabezpečujúca odvoz domového odpadu alebo obchod, v ktorom ste si výrobok zakúpili.

Slovenian notice

Odstranjevanje odslužene opreme uporabnikov v zasebnih gospodinjstvih v Evropski uniji

Ta znak na izdelku ali njegovi embalaži pomeni, da izdelka ne smete odvreči med gospodinjske odpadke. Nasprotno, odslužено opremo morate predati na zbirališče, pooblaščeno za recikliranje odslužene električne in elektronske opreme. Ločeno zbiranje in recikliranje odslužene opreme prispeva k ohranjanju naravnih virov in zagotavlja recikliranje te opreme na zdravju in okolju neškodljiv način. Za podrobnejše informacije o tem, kam lahko odpeljete odslužено opremo na recikliranje, se obrnite na pristojni organ, komunalno službo ali trgovino, kjer ste izdelek kupili.

Spanish notice

Eliminación de residuos de equipos eléctricos y electrónicos por parte de usuarios particulares en la Unión Europea

Este símbolo en el producto o en su envase indica que no debe eliminarse junto con los desperdicios generales de la casa. Es responsabilidad del usuario eliminar los residuos de este tipo depositándolos en un "punto limpio" para el reciclado de residuos eléctricos y electrónicos.

La recogida y el reciclado selectivos de los residuos de aparatos eléctricos en el momento de su eliminación contribuirá a conservar los recursos naturales y a garantizar el reciclado de estos residuos de forma que se proteja el medio ambiente y la salud. Para obtener más información sobre los puntos de recogida de residuos eléctricos y electrónicos para reciclado, póngase en contacto con su ayuntamiento, con el servicio de eliminación de residuos domésticos o con el establecimiento en el que adquirió el producto.

Swedish notice

Bortskaffande av avfallsprodukter från användare i privathushåll inom Europeiska Unionen

Om den här symbolen visas på produkten eller förpackningen betyder det att produkten inte får slängas på samma ställe som hushållssopor. I stället är det ditt ansvar att bortskaffa avfallet genom att överlämna det till ett uppsamlingsställe avsett för återvinning av avfall från elektriska och elektroniska produkter. Separat insamling och återvinning av avfallet hjälper till att spara på våra naturresurser och gör att avfallet återvinns på ett sätt som skyddar människors hälsa och miljön. Kontakta ditt lokala kommunkontor, din närmsta återvinningsstation för hushållsavfall eller affären där du köpte produkten för att få mer information om var du kan lämna ditt avfall för återvinning.

Germany noise declaration

Schalldruckpegel Lp = 70 dB(A)

Am Arbeitsplatz (operator position)

Normaler Betrieb (normal operation)

Nach ISO 7779:1999 (Typprüfung)

Japanese notice

ご使用になっている装置にVCCIマークが付いていましたら、次の説明文をお読み下さい。

この装置は、情報処理装置等電波障害自主規制協議会（VCCI）の基準に基づくクラスB情報技術装置です。この装置は、家庭環境で使用することを目的としていますが、この装置がラジオやテレビジョン受信機に近接して使用されると、受信障害を引き起こすことがあります。
取扱説明書に従って正しい取り扱いをして下さい。

VCCIマークが付いていない場合には、次の点にご注意下さい。

この装置は、情報処理装置等電波障害自主規制協議会（VCCI）の基準に基づくクラスA情報技術装置です　この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

Harmonics conformance (Japan)

Taiwanese notice

警告使用者:
這是甲類的資訊產品，在居住的環境中使用時，可能會造成射頻干擾，在這種情況下，使用者會被要求採取某些適當的對策。

Japanese power cord notice

製品には、同梱された電源コードをお使い下さい。
同梱された電源コードは、他の製品では使用出来ません。

Country-specific certifications

HP tests electronic products for compliance with country-specific regulatory requirements, as an individual item or as part of an assembly. The product label (see Figure 39 (page 112)) specifies the regulations with which the product complies.

NOTE: Components without an individual product certification label are qualified as part of the next higher assembly (for example, enclosure, rack, or tower).

Figure 39 Typical enclosure certification label

CXO8157A

NOTE: The certification symbols on the label depend upon the certification level. For example, the FCC Class A certification symbol is not the same as the FCC Class B certification symbol.

Storage system specifications

This appendix defines the physical, environmental, and power specifications of the EVA 4x00/6x00/8x00 storage systems.

Physical specifications

This section describes the physical specifications of the drive enclosure and elements.

Table 32 (page 112) defines the dimensions and weights of the storage system components.

Table 32 Enterprise Virtual Array 4x00/6x00/8x00 Product Dimensions, Weight and Clearance

Physical Dimensions	Height in/cm	Width in/cm	Depth in/cm	Max Weight lb/kg	Req. Front Clearance in/cm	Req. Rear Clearance in/cm
Evacuees and EVA8x00 2C2D (42U rack)	78.75 (200.03)	23.7 (60.3)	40.2 (102.2)	537 (244.1)	30 (76.2)	30 (76.2)
EVA6x00 and EVA8x00 2C6D (42U rack)	78.75 (200.03)	23.7 (60.3)	40.2 (102.2)	854 (308.2)	30 (76.2)	30 (76.2)
EVA8x00 2C12D (42U rack)	78.75 (200.03)	23.7 (60.3)	40.2 (102.2)	1290 (586.4)	30 (76.2)	30 (76.2)
EVA4x00/6x00/8x00 Controller Assembly	7.0/17.78	17.6/44.70	27.5/69.85	120/54.55	N/A	N/A
M5314B/M5314C Drive Enclosure	5.25/13.34	19.0/42.26	20/50.8	71/32.21	N/A	N/A

Environmental specifications

To ensure optimum product operation, you must maintain the operational environmental specifications listed in Table 33 (page 112). The ambient temperature (the enclosure air intake or room temperature) is especially critical.

Table 33 Environmental specifications

Operating Temperature	50° to 95° F (10° to 35° C) - Reduce rating by 1° F for each 1000 ft. altitude (1.8° C/1,000 m)
Shipping Temperature	-40° to 150° F (-40° to 66° C)

Table 33 Environmental specifications *(continued)*

Humidity	10% to 90% non-condensing
Shipping Humidity	5% to 90% non-condensing
Altitude	Up to 8,000 ft. (2,400 m)
Air Quality	Not to exceed 500,000 particles per cubic foot of air at a size of 0.5 micron or larger

Power specifications

The input voltage is a function of the country-specific input voltage to Enterprise storage system rack power distribution units (PDUs). Table 34 (page 113) defines the AC input power available to the drive enclosure power supplies.

△ **CAUTION:** The AC power distribution within a rack ensures balance to each PDA and reduces the possibility of an overload condition. Changing the cabling to or from a PDM could cause an overload condition.

Table 34 Enterprise storage system AC input line voltages

Specification	Minimal	Nominal	Maximum
	60 Hz service		
AC Line Voltage	57 Hz	60 Hz	63 Hz
AC Line Voltage—Japan	180 VAC	202 VAC	220 VAC
AC Line Voltage–North America	180 VAC	208 VAC	220 VAC
AC Line Voltage–Europe	208 VAC	240 VAC	254 VAC
	50 Hz service		
AC Line Frequency	47 Hz	50 Hz	53 Hz
AC Line Voltage–Japan	180 VAC	202 VAC	220 VAC
AC Line Voltage–North America	190 VAC	220 VAC	235 VAC
AC Line Voltage–North America	200 VAC	230 VAC	244 VAC
AC Line Voltage–Europe	208 VAC	240 VAC	254 VAC

Table 35 Power Data (North America/Europe/Japan) maximum configuration

AC plug type (quantity 2)	North America – 3 wire NEMA No. L6-30P, 30 amp (208 to 240V 50–60Hz 30A) Europe – 3 wire, 2 pole IEC 309, 30 amp, (220 to 240V 50Hz 32A)
Number of phases	Single
Rated current	17A @ 200V-240V AC, 60Hz total, 4.25 A per power cord
Nominal Line Voltage	North America – 208 or 230V Europe – 230V Japan – 206V
Range Line Voltage	187 to 256V
Line Frequency	North America 60Hz, Europe 50Hz, Japan 50 or 60 Hz

Table 36 EVA4x00 power specifications — 208 Volts

Specification		2C1D	2C2D	2C3D	2C4D
Typical[1]	Total System Wattage	638	1013	1390	1767
	Total System BTU/hour	1729	3014	4300	5585
	Input Current (A) - Typical per line	1.6	2.6	3.5	4.4
	In Rush Current (A)	98	132	170	220
Failover Mode	Input Current (A) - Maximum per line	2.7	4.3	5.9	7.5

[1] Typical is described as a system in normal steady state operation. (i.e., both PDUs operating normally, the array reading/writing to disk drives in a production environment).
This data represents fully populated drive enclosures with 15K RPM disk drives. Other drive types may vary slightly. For example, if you are using 10K RPM drives, the power specifications will be approximately 20% less than the 15K RPM drives.

Table 37 EVA4x00 power specifications — 230 Volts

Specification		2C1D	2C2D	2C3D	2C4D
Typical[1]	Total System Wattage	638	1013	1390	1767
	Total System BTU/hour	1729	3014	4300	5585
	Input Current (A) - Typical per line	1.5	2.4	3.3	4.2
	In Rush Current (A)	104	147	190	244
Failover Mode	Input Current (A) - Maximum per line	2.6	4.1	5.5	7.1

[1] Typical is described as a system in normal steady state operation. (i.e., both PDUs operating normally, the array reading/writing to disk drives in a production environment).
This data represents fully populated drive enclosures with 15K RPM disk drives. Other drive types may vary slightly. For example, if you are using 10K RPM drives, the power specifications will be approximately 20% less than the 15K RPM drives.

Table 38 EVA6x00 power specifications — 208 Volts

Specification		2C4D	2C5D	2C6D	2C7D	2C8D
Typical[1]	Total System Wattage	1837	2214	2590	2967	3344
	Total System BTU/hour	6268	7553	8838	10124	11409
	Input Current (A) - Typical per line	4.6	5.5	6.5	7.4	8.4
	In Rush Current (A)	220	250	280	321	363
Failover Mode	Input Current (A) - Maximum per line	7.8	9.4	11.0	12.6	14.2

[1] Typical is described as a system in normal steady state operation. (i.e., both PDUs operating normally, the array reading/writing to disk drives in a production environment).

This data represents fully populated drive enclosures with 15K RPM disk drives. Other drive types may vary slightly. For example, if you are using 10K RPM drives, the power specifications will be approximately 20% less than the 15K RPM drives.

Table 39 EVA6x00 power specifications — 230 Volts

Specification		2C4D	2C5D	2C6D	2C7D	2C8D
Typical[1]	Total System Wattage	1837	2214	2590	2967	3344
	Total System BTU/hour	6268	7553	8838	10124	11409
	Input Current (A) - Typical per line	4.3	5.2	6.1	7.0	7.9
	In Rush Current (A)	244	272	311	357	403
Failover Mode	Input Current (A) - Maximum per line	7.3	8.8	10.0	11.9	13.3

[1] Typical is described as a system in normal steady state operation. (i.e., both PDUs operating normally, the array reading/writing to disk drives in a production environment).
This data represents fully populated drive enclosures with 15K RPM disk drives. Other drive types may vary slightly. For example, if you are using 10K RPM drives, the power specifications will be approximately 20% less than the 15K RPM drives.

Table 40 EVA8x00 power specifications — 208 Volts

Specification		2C2D	2C6D	2C8D	2C10D	2C12D
Typical[1]	Total System Wattage	1153	2660	3414	4167	4920
	Total System BTU/hour	3936	9077	11648	14218	16789
	Input Current (A) - Typical per line	2.9	6.7	8.5	10.4	12.3
	In Rush Current (A)	132	280	363	451	528
Failover Mode	Input Current (A) - Maximum per line	4.9	11.3	14.5	17.7	20.8

[1] Typical is described as a system in normal steady state operation. (i.e., both PDUs operating normally, the array reading/writing to disk drives in a production environment).
This data represents fully populated drive enclosures with 15K RPM disk drives. Other drive types may vary slightly. For example, if you are using 10K RPM drives, the power specifications will be approximately 20% less than the 15K RPM drives.

Table 41 EVA8x00 power specifications — 230 Volts

Specification		2C2D	2C6D	2C8D	2C10D	2C12D
Typical[1]	Total System Wattage	1153	2660	3414	4167	4920
	Total System BTU/hour	3936	9077	11648	14218	16789
	Input Current (A) - Typical per line	2.7	6.3	8.1	9.8	11.6
	In Rush Current (A)	147	311	403	500	586
Failover Mode	Input Current (A) - Maximum per line	4.6	10.6	13.7	16.7	19.7

[1] Typical is described as a system in normal steady state operation. (i.e., both PDUs operating normally, the array reading/writing to disk drives in a production environment).
This data represents fully populated drive enclosures with 15K RPM disk drives. Other drive types may vary slightly. For example, if you are using 10K RPM drives, the power specifications will be approximately 20% less than the 15K RPM drives.

B EMU-generated condition reports

This section provides a description of the EMU generated condition reports that contain the following information:

- Element type (*et*), a hexadecimal number in the range *01* through *FF*.
- Element number (*en*), a decimal number in the range *00* through *99* that identifies the specific element with a problem.
- Error code (*ec*), a decimal number in the range *00* through *99* that defines a specific problem.
- The recommended corrective action.

NOTE: The conventions used to differentiate between the elements of the condition report are:

- Element type—period after each character
- Element number—period after the second character
- Error code—no periods

The EMU can send error messages to the controller for transmission to HP P6000 Command View. The messages displayed are specific to HP P6000 Command View and are not within the scope of this publication.

The I/O modules have the built-in intelligence to:

- Observe fibre channel events
- Bypass drive ports based on events
- Perform drive port testing and monitoring to prevent poor-performing drives from participating in the loop
- Communicate fibre channel events to the controllers

This appendix explains the condition report format, correcting problems, and how to identify element types. The error codes are arranged in element type sequence (that is, *0.1., 0.2., 0.3., etc.*).

Condition report format

When the EMU alphanumeric display is Er, three additional displays identify the possible cause of the problem: the element type, the specific element, and the error code, which defines the possible cause of the problem.

- The first-level display identifies the type of element affected with two alphanumeric characters separated by periods such as *0.1., 0.2., 1.3., F.F.*, and so forth. A disk drive problem would display an element type number of *0.1*.
- The second-level display identifies the element affected with a two-digit, decimal number followed by a period. For example, when a bay 6 drive error occurs, the element number display is *06.*; a display of *14.* indicates a bay 14 problem.
- The third-level display identifies a specific problem, the error code with a two-digit, decimal number. For example, should the problem be either the installation of an incorrectly configured drive or one that cannot operate at the loop link rate, the display is *01*.

Correcting errors

Correcting an error may require you to perform a specific set of actions. In some cases, the only available corrective action is to replace the element.

Table 42 (page 118) lists the element type codes assigned to the drive enclosure elements.

Table 42 Assigned element type codes

Code	Element
0.1.	Disk Drives
0.2.	Power Supplies
0.3.	Blowers
0.4.	Temperature Sensors
0.6	Audible Alarm [1]
0.7.	EMU
0.C.	Controller OCP LCD[1]
0.F.	Transceivers
1.0.	Language[1]
1.1.	Communication Port
1.2.	Voltage Sensors
1.	Current Sensors
8.0	Drive Enclosure[1]
8.2.	Drive Enclosure Backplane
8.7.	I/O Modules

[1] Does not generate a condition report. However, for any error, you should record the error code. Then, implement the recommended corrective action.

Drive conditions

The format of a disk drive condition report is *0.1.en.ec*, where:

- *0.1.* is the disk drive element type number
- *en.* is the two-character disk drive element
- *ec* is the error code

A direct correlation exists between the disk drive element number and the bay number. However, no direct correlation exists between the disk drive bay number and the device Fibre Channel drive enclosure physical address. The Fibre Channel drive enclosure physical address is assigned by negotiation during system initialization.

The following sections define the disk drive error codes.

0.1.en.01 CRITICAL condition—Drive configuration or drive link rate

As each drive spins up and comes on-line, the EMU determines if the drive is Fibre Channel compatible and can operate at the link rate (1 Gbps or 2 Gbps) established by the I/O module. If either of these conditions are not met, the EMU issues the condition report *0.1.en.01*.

The corrective actions for these conditions are:

- When the drive is not Fibre Channel-compatible you must install a Fibre Channel compatible drive or a drive blank.
- When the drive is Fibre Channel-compatible, the EMU compares the drive link rate with the I/O module link rate, the loop link rate.

If the EMU cannot determine the drive link rate, the EMU activates the drive bypass function for one minute. During this time the EMU continually checks the drive to determine the link rate.

- If the EMU determines the drive cannot operate at the Fibre Channel link rate set by the I/O module, the drive bypass function ends and the drive is placed on the loop. This does not generate a condition report.
- The EMU issues the condition report *0.1.en.01* when the drive link rate is incompatible with Fibre Channel link rate.
- When the EMU cannot determine the drive link rate during the one-minute drive bypass time, the EMU places the drive on the loop. This process allows the drive to negotiate for an address.
  - If negotiation indicates the link rates are compatible, the EMU rechecks the drive link rate to verify compatibility.
  - If negotiation indicates the link rates are incompatible, an error condition exists and drive loop data transfers stop.

This condition report remains active until the problem is corrected. The problem affects disk drive en. Therefore, correction to prevent the possible failure of other elements is not required.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Replace the defective drive with:
   - A Fibre Channel-compatible drive.
   - A Fibre Channel drive capable of operating at a link rate supported by I/O modules and transceivers.
3. Observe the EMU to ensure the error is corrected.
4. If unable to correct the problem, contact your authorized service representative.

0.1.en.02 INFORMATION condition—Drive missing

The drive is improperly installed or missing. Either option could affect the enclosure air flow and cause an over temperature condition for another element.

- This error remains active for one minute, or until the problem is corrected, whichever occurs first.
- Immediate correction is not required. However, correction cannot be delayed indefinitely.

Complete the following procedure to correct this problem.

1. Record all six characters of the condition report.
2. Remove and install the drive to ensure that it is properly installed.
3. Observe the EMU to ensure the error is corrected.
4. If removing and installing the drive did not correct the problem, install a replacement drive or a drive blank.
5. Observe the EMU to ensure the error is corrected.
6. If unable to correct the problem, contact your authorized service representative.

0.1.en.03 INFORMATION condition—Drive software lock active

Some enclosures have a software-activated lock that prevents physically removing a drive while this feature is active. This feature can be activated even when an enclosure does not have a physical lock. Removing a drive when this feature is active generates a condition report. This error remains active for 15 seconds.

No action is required to correct this condition.

0.1.en.04 CRITICAL condition—Loop a drive link rate incorrect

The drive is capable of operating at the loop link rate but is running at a different rate. For example, the drive is operating at 1 Gbps, and the loop is operating at 2 Gbps. Only when the drive is operating at the Fibre Channel link rate established by the I/O module can this drive transfer data.

This error remains active until the problem is corrected.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Remove and replace the drive in the enclosure.
3. Observe the drive status indicators to ensure the drive is operational.
4. Observe the EMU to ensure the error is corrected.
5. If removing and replacing the drive did not correct the problem, replace the drive.
6. Observe the drive status indicators to ensure the drive is operational.
7. Observe the EMU to ensure the error is corrected.
8. If unable to correct the problem, contact your authorized service representative.

0.1.en.05 CRITICAL condition—Loop b drive link rate incorrect

The drive is capable of operating at the loop link rate but is running at a different rate. For example, the drive is operating at 1 Gbps, and the loop is operating at 2 Gbps. Only when the drive is operating at the Fibre Channel link rate established by the I/O module can this drive transfer data.

This error remains active until the problem is corrected.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Remove and replace the drive in the enclosure.
3. Observe the drive status indicators to ensure the drive is operational.
4. Observe the EMU to ensure the error is corrected.
5. If removing and replacing the drive did not correct the problem, replace the drive.
6. Observe the drive status indicators to ensure the drive is operational.
7. Observe the EMU to ensure the error is corrected.
8. If unable to correct the problem, contact your authorized service representative.

Power supply conditions

The format of a power supply condition report is *0.2.en.ec*, where:

- *0.2.* is the power supply element type number
- *en.* is the two-character power supply element number
- *ec* is the error code

Figure 40 (page 120) shows the location of power supply 1 and power supply 2.

Figure 40 Power supply element numbering

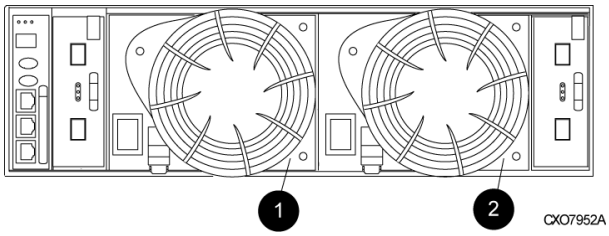

The following sections define the power supply condition reports.

0.2.en.01 NONCRITICAL Condition—Power supply AC input missing

The loss of the AC input to a power supply makes the remaining power supply a single point of failure.

This condition report remains active until AC power is applied to the power supply.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.

2. Ensure that there is AC power to the rack PDU, and from the PDU to the PDMs, and that the PDU and PDM circuit breakers are not reset.

   If there is no AC power to the PDU, contact building facilities management.

   Verify that the power supply AC power cord is properly connected.
3. If AC is present, and the rack power distribution circuitry is functioning properly, the power supply indicator should be on.
4. Observe the EMU to ensure the error is corrected.
5. Contact your authorized service representative.

0.2.en.02 UNRECOVERABLE condition—Power supply missing

This condition report indicates a power supply is not installed or installed incorrectly. Both of these conditions affect air flow within the enclosure and can cause an over-temperature condition. Enclosure shutdown is imminent.

The operational power supply will automatically shut down after seven minutes, thereby disabling the enclosure. This condition report remains active until either the problem is corrected, or the operational power supply shuts down, whichever occurs first.

To correct this problem, record all six characters of the condition report, then contact your authorized service representative.

△ **CAUTION:** Removing power from an enclosure may cause the loss or corruption of data. To avoid this condition, shut down the system using HP P6000 Command View. An automatic shutdown and possible data corruption may result if the power supply is removed before a replacement is available.

0.2.en.03 CRITICAL condition—Power supply load unbalanced

This condition report indicates that a component within a power supply may have failed. This can make the remaining power supply a single point of failure.

This condition report remains active until corrected.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Ensure that the blower on the power supply is functioning properly. If not, correct the blower condition and wait one minute.
3. Contact your authorized service representative.

Blower conditions

The format of a blower condition report is *0.3.en.ec*, where:

- *0.3.* is the blower element type number
- *en.* is the two-character blower element number
- *ec* is the error code

As shown in Figure 41 (page 122), blower 1 is in location 1 and blower 2 is in location 2.

Figure 41 Blower element numbering

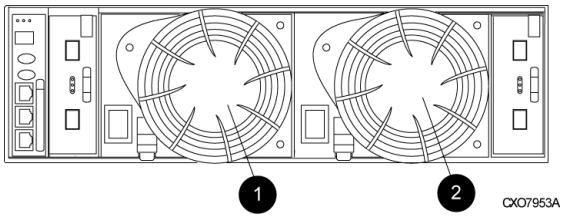

△ **CAUTION:** A single blower operating at high speed can provide sufficient air flow to cool an enclosure and the elements for up to 100 hours. However, operating an enclosure at temperatures approaching an overheating threshold can damage elements and may reduce the MTBF of a specific element. Immediate replacement of the defective blower is required.

The following sections define the power supply condition reports.

0.3.en.01 NONCRITICAL condition—Blower speed

A blower is operating at a speed outside of the EMU specified range, possibly because of a bearing problem. This can affect enclosure cooling and cause an element to fail. This condition report remains active until the problem is corrected.

This error does not normally require immediate correction. However, an error of this type could contribute to an element overheating.

HP recommends replacing the blower as soon as possible.

To correct this problem, record all six characters of the condition report, then contact your HP authorized service representative.

0.3.en.02 CRITICAL condition—Blower speed

A blower is operating at a speed that is significantly outside the EMU specified range, possibly because of a bearing problem. This can cause the loss of cooling and cause an element to fail. The error remains active until the problem is corrected.

HP recommends replacing the blower as soon as possible.

To correct this problem, record all six characters of the condition report, then contact your authorized service representative.

0.3.en.03 UNRECOVERABLE condition—Blower failure

A blower has stopped. The operational blower now operates at high speed and is a single point of failure. This condition report remains active until the problem is corrected.

To correct this problem, record all six characters of the condition report, then contact your authorized service representative.

0.3.en.04 UNRECOVERABLE condition—Blower internal

A power supply reported an internal blower error that could affect enclosure cooling and cause an element to fail. HP recommends correcting the problem before the blower fails. This condition report remains active until the problem is corrected.

To correct this problem, record all six characters of the condition report, then contact your HP authorized service representative.

0.3.en.05 NONCRITICAL condition—Blower missing

A blower has been removed or is improperly installed. Even though the blower flaps close to maintain the proper air flow, the reduced cooling capability can cause overheating, causing an element to fail. This condition report remains active until the problem is corrected.

To correct this problem, record all six characters of the condition report, then contact your HP authorized service representative.

0.3.en.06 UNRECOVERABLE condition—No blowers installed

NOTE: IMPORTANT

When this condition exists there will be two error messages.

The first message will be *0.3.en.05* and will identify the first blower.

The second message will be *0.3.en.06* and will identify the second blower.

The EMU cannot detect any installed blowers. *Shutdown is imminent!* The EMU will shut down the enclosure in *seven minutes* unless you correct the problem. This condition report remains active until you correct the problem or the EMU shuts down the power supplies, whichever occurs first.

Complete the following procedure to correct this problem.

1. Record all six characters of the condition report.
2. Use the controller shutdown procedure to shut down the controllers.
3. Contact your authorized service representative.

CAUTION: An automatic shutdown and possible data corruption may result if the blower is removed before a replacement is available.

Temperature conditions

The format of a temperature condition report is *0.4.en.ec*, where:

- *0.4.* is the temperature sensor element type
- *en.* is the two-character temperature sensor element
- *ec* is the error code

See Table 43 (page 123) to determine the location of each temperature sensor.

Table 43 Temperature sensor element numbering

Sensor	Sensor location	Sensor	Sensor location
01.	Power Supply 1 Exhaust	10.	Drive Bay 7
02.	Power Supply 2 Exhaust	11.	Drive Bay 8
03.	EMU	12.	Drive Bay 9
04.	Drive Bay 1	13.	Drive Bay 10
05.	Drive Bay 2	14.	Drive Bay 11
06.	Drive Bay 3	15.	Drive Bay 12
07.	Drive Bay 4	16.	Drive Bay 13
08.	Drive Bay 5	17.	Drive Bay 14
09.	Drive Bay 6		

The following sections list the temperature condition reports and the default temperature thresholds. Use HP P6000 Command View to view the temperature sensor ranges for the disk drives, EMU, and power supplies.

0.4.en.01 NONCRITICAL condition—High temperature

This condition report indicates that an element temperature is approaching, but has not reached, the high temperature CRITICAL threshold. Continued operation under these conditions may result in a CRITICAL condition. This condition report remains active until the problem is corrected.

Complete the following procedure to correct this problem.

1. Record all six characters of the condition report.
2. Ensure that all elements are properly installed to maintain proper air flow.

3. Ensure that nothing is obstructing the air flow at either the front of the enclosure or the rear of the blower.
4. Ensure that both blowers are operating properly (the indicators are on) and neither blower is operating at high speed.
5. Verify that the ambient temperature range is +10° C to +35° C (+50° F to +95° F). Correct the ambient conditions.
6. Observe the EMU to ensure the error is corrected.
7. If unable to correct the problem, contact your authorized service representative.

0.4.en.02 CRITICAL condition—High temperature

This condition report indicates that an element temperature is above the high temperature CRITICAL threshold. Continued operation under these conditions may result in element failure and may reduce an element MTBF. This condition report remains active until the problem is corrected.

Complete the following procedure to correct this problem.

1. Record all six characters of the condition report.
2. Ensure that all elements are properly installed to maintain proper air flow.
3. Ensure that nothing is obstructing the air flow at either the front of the enclosure or the rear of the blower.
4. Ensure that both blowers are operating properly (the indicators are on) and neither blower is operating at high speed.
5. Verify that the ambient temperature range is +10° C to +35° C (+50° F to +95° F). Adjust as necessary.
6. Observe the EMU to ensure the error is corrected.
7. If unable to correct the problem, contact your authorized service representative.

0.4.en.03 NONCRITICAL condition—Low temperature

This condition report indicates that an element temperature is approaching, but has not reached, the low temperature CRITICAL threshold. Continued operation under these conditions may result in a CRITICAL condition. This condition report remains active until the problem is corrected.

Complete the following procedure to correct this problem.

1. Record all six characters of the condition report.
2. Verify that the ambient temperature range is +10° C to +35° C (+50° F to +95° F). Adjust as necessary.
3. Observe the EMU to ensure the error is corrected.
4. If the ambient temperature is correct and the problem persists, contact your Authorized Service Representative.

0.4.en.04 CRITICAL condition—Low temperature

This condition report indicates that an element temperature has reached the low temperature CRITICAL threshold. HP recommends correcting this error to prevent affecting other elements. This condition report remains active until the problem is corrected.

Complete the following procedure to correct this problem.

1. Record all six characters of the condition report.
2. Verify that the ambient temperature range is +10° C to +35° C (+50° F to +95° F). Adjust as necessary.
3. Observe the EMU to ensure the error is corrected.
4. If the ambient temperature is correct and the problem persists, contact your authorized service representative.

0.4.en.05 UNRECOVERABLE condition—High temperature

This condition report indicates that the EMU has evaluated the temperature of the three temperature groups (EMU, disk drives, and power supplies), and determined that the average temperature of two of the three groups exceeds the critical level (use HP P6000 Command View to view the

temperature thresholds). Under these conditions the EMU starts a timer that will automatically shut down the enclosure in seven minutes unless you correct the problem. *Enclosure shutdown is imminent!*

CAUTION: An automatic shutdown and possible data corruption may result if the procedure below is not performed *immediately*.

Complete the following procedure to correct this problem.

1. Ensure that all disk drives, I/O modules, and power supply elements are fully seated.
2. Ensure that all blowers are operating properly.
3. Verify that the ambient temperature range is +10° C to +35° C (+50° F to +95° F). Adjust as necessary.
4. If steps 1, 2 or 3 did not reveal a problem, use HP P6000 Command View to request the HSV controller to shut down the drive enclosure. Completing this action will halt the drive enclosure data transfers.
5. Contact your authorized service representative and request assistance.

EMU conditions

The format of an EMU condition report is *0.7.01.ec*, where:

- *0.7.* is the EMU element type number
- *01.* is the two-character EMU element number
- *ec* is the error code

NOTE: There is only one EMU in a drive enclosure. Therefore, the element number is always 01.

Resetting the EMU

In some cases, the only corrective action for an EMU error is to replace the EMU. Call your authorized service representative if this action is required. Another option is to reset the EMU using the following procedure.

1. Firmly grasp the EMU mounting handle and pull the EMU partially out of the enclosure.

   NOTE: You do not need to remove the EMU from the enclosure or disconnect the cables. You must avoid putting any strain on the cables or connectors.

2. Wait 30 seconds, and then push the EMU in and fully seat the element in the backplane. The EMU should display any enclosure condition report within two minutes.

07.01.01 CRITICAL condition—EMU internal clock

There is an internal EMU clock error that will remain active until the problem is corrected.

Complete the following procedure to correct this problem.

1. Record all six characters of the condition report.
2. Reset the EMU.
3. If resetting the EMU did not correct the problem, replace the EMU.
4. Observe the EMU to ensure the error is corrected.
5. If unable to correct the problem, contact your HP authorized service representative.

07.01.02 UNRECOVERABLE condition—EMU interrupted

The Inter-IC (I2C) bus is not processing data and the EMU is unable to monitor or report the status of the elements or enclosures. *IMMEDIATE* corrective action is required to ensure proper enclosure operation. This condition report remains active until the problem is corrected.

Complete the following procedure *NOW* to correct this problem.

1. Record all six characters of the condition report.
2. Reset the EMU.
3. Observe the EMU to ensure the error is corrected.

4. If resetting the EMU did not correct the problem, replace the EMU.
5. If unable to correct the problem, contact your HP authorized service representative.

0.7.01.03 UNRECOVERABLE Condition—Power supply shutdown

This message only appears in HP P6000 Command View to report a power supply has already shut down. This message can be the result of the controller shutdown command or an EMU or power supply initiated power shutdown.

This message cannot be displayed until after restoration of power. Therefore, there is no corrective action required.

0.7.01.04 INFORMATION condition—EMU internal data

The EMU is unable to collect data for the SCSI-3 Engineering Services (SES) page. This condition report remains active for 15 seconds. The condition report affects only internal EMU operations. There is no degradation of enclosure operations.

The EMU initiates automatic recovery procedures.

If the problem is not automatically corrected after one minute, contact your HP authorized service representative.

0.7.01.05 UNRECOVERABLE condition—Backplane NVRAM

NOTE: IMPORTANT

Backplane NVRAM errors usually occur during manufacture. At this time they are identified and corrected. They rarely occur during normal operation.

When a backplane NVRAM is not programmed or cannot be read by the EMU, there is no communication with the disk drives. This condition report remains active until the problem is corrected.

Complete the following procedure to correct this problem.

1. Record all six characters of the condition report.
2. Reset the EMU.
3. Observe the EMU to ensure the error is corrected.
4. If resetting the EMU did not correct the problem, contact your HP authorized service representative.

0.7.01.10 NONCRITICAL condition—NVRAM invalid read data

The data read from the EMU NVRAM is invalid. This error initiates an automatic recovery process. This condition report remains active until the problem is corrected.

If the automatic recovery process does not correct the problem, complete the following procedure.

1. Record all six characters of the condition report.
2. Reset the EMU.
3. Observe the EMU to ensure the error is corrected.
4. If resetting the enclosure did not correct the problem, contact your HP authorized service representative.

0.7.01.11 NONCRITICAL condition—EMU NVRAM write failure

The EMU cannot write data to the NVRAM. This condition report remains active until the problem is corrected.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Reset the EMU.
3. Observe the EMU to ensure the error is corrected.
4. If resetting the enclosure did not correct the problem, contact your HP authorized service representative.

0.7.01.12 NONCRITICAL condition—EMU cannot read NVRAM data

The EMU is unable to read data from the NVRAM. This condition report remains active until the problem is corrected.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Reset the EMU.
3. Observe the EMU to ensure the error is corrected.
4. If resetting the enclosure did not correct the problem, contact your HP authorized service representative.

0.7.01.13 UNRECOVERABLE condition—EMU load failure

The EMU Field Programmable Gate Array (FPGA) that controls the ESI bus failed to load information required for EMU operation. This condition report remains active until the problem is corrected.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Reset the EMU.
3. Observe the EMU to ensure the error is corrected.
4. If resetting the enclosure did not correct the problem, contact your HP authorized service representative.

0.7.01.14 NONCRITICAL condition—EMU enclosure address

Either the enclosure address is incorrect or the enclosure has no address. Possible causes include a defective enclosure address bus cable, an incorrectly connected cable, or a defective enclosure address bus enclosure ID expansion cable. This condition report remains active until the problem is corrected.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Remove and reconnect the cable between the address bus enclosure ID expansion cable and the EMU.

   NOTE: The EMU display may not display a change in condition for up to 30 seconds.

3. Observe the EMU to ensure the error is corrected.
4. If the problem is not corrected, remove and reinstall the lower and upper terminators, and all the enclosure ID expansion cable-to-enclosure ID expansion cables.
5. Observe the EMU to ensure the error is corrected.
6. Reset the EMU, then observe the EMU to ensure the error is corrected.
7. If resetting the EMU did not correct the problem, contact your HP authorized service representative.

0.7.01.15 UNRECOVERABLE condition—EMU hardware failure

The EMU has detected an internal hardware problem. This condition report remains active until the problem is corrected.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Reset the EMU.
3. Observe the EMU to ensure the error does not recur within the first minute.
4. If the error does recur, contact your HP authorized service representative. The EMU is inoperative and must be replaced as soon as possible.

0.7.01.16 INFORMATION condition—EMU internal ESI data corrupted

The EMU ESI data is corrupted. This condition does not affect any other element and no action is required.

0.7.01.17 UNRECOVERABLE condition—Power shutdown failure

The power supply did not respond to a controller, EMU, or power supply shut down command. Shutting down the supply is required to prevent overheating.

Complete the following procedure to correct the problem:

1. Record all six characters of the condition report.
2. Move the power cord bail lock 1, Figure 42 (page 128), to the left.
3. Disconnect the AC power cord 2 from the supply.

Figure 42 Disconnecting AC power

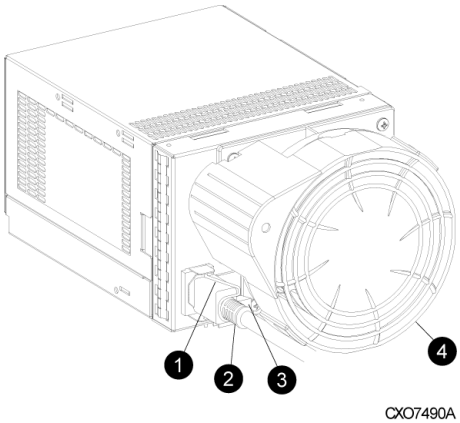

0.7.01.18 UNRECOVERABLE condition—EMU hardware failure

The EMU has detected an internal hardware problem. This condition report remains active until the problem is corrected.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Reset the EMU.
3. Observe the EMU to ensure the error does not recur within the first minute.
4. If the error does recur, contact your HP authorized service representative. The EMU is inoperative and must be replaced as soon as possible.

0.7.01.19 UNRECOVERABLE condition—EMU ESI driver failure

The EMU has detected an internal hardware problem. This condition report remains active until the problem is corrected.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Reset the EMU.
3. Observe the EMU to ensure the error does not recur within the first minute.
4. If the error does recur, contact your HP authorized service representative. The EMU is inoperative and must be replaced as soon as possible.

Transceiver conditions

The format of a transceiver condition report is *0.F.en.ec*, where:

- *0.F.* is the transceiver element type number
- *en.* is the two-character transceiver element number (see Figure 43 (page 129))
- *ec* is the error code

Figure 43 Transceiver element numbering

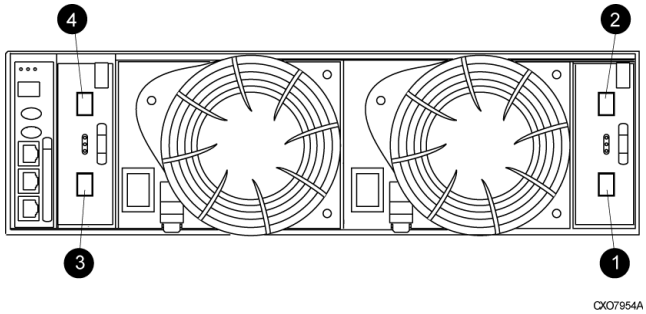

1. Transceiver 01
2. Transceiver 02
3. Transceiver 03
4. Transceiver 04

0.F.en.01 CRITICAL condition—Transceiver incompatibility

The transceivers on this link are not the same type or they are incompatible with the I/O module. This error prevents the controller from establishing a link with the enclosure disk drives and eliminates the enclosure dual-loop capability. This error remains active until the problem is corrected.

To correct this problem, record all six characters of the condition report, then contact your HP authorized service representative.

0.F.en.02 CRITICAL condition—Transceiver data signal lost

This symptom can occur when a controller has been powered off or a cable has been removed from the transceiver. The transceiver can no longer detect a data signal. This error prevents the controller from transferring data on a loop and eliminates the enclosure dual-loop capability. This error remains active until the problem is fixed.

To correct this problem, record all six characters of the condition report, then contact your HP authorized service representative.

0.F.en.03 CRITICAL condition—Transceiver fibre channel drive enclosure bus fault

The system has detected a Fibre Channel drive enclosure bus fault involving a transceiver. This error prevents the controller from transferring data on a loop and eliminates the enclosure dual-loop capability.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Check all the transceivers and cables to ensure they are properly connected.
3. Check all the transceivers on the loop to ensure they are compatible with the I/O module.
4. If the problem persists, contact your HP authorized service representative.

0.F.en.04 CRITICAL condition—Transceiver removed

The EMU detects that a transceiver has been removed. This error remains active until the problem is fixed.

The error can be cleared by doing one of the following:

1. Install a new transceiver,

   or

   Reconfigure the system by moving from a loop topology to one with Vixel switches. This change makes the transceiver is unnecessary.
2. Clear the error by resetting the EMU or by removing and then re-installing the I/O module.

0.F.en.05 CRITICAL condition—Invalid fibre channel character

This symptom can occur under the following conditions:

- The incoming data stream is corrupted.
- A cable is not completely connected.
- The signal is degraded.

This error prevents the controller from transferring data on a loop and eliminates the enclosure dual-loop capability. This error remains active until the problem is fixed.

To correct this problem, record all six characters of the condition report, then contact your HP authorized service representative.

CAN bus communication port conditions

The format of a CAN bus communication port report is *1.1.03.ec* , where:

- *1.1.* is the communication port element type
- *03.* is the two-character CAN bus element number
- *ec* is the error code

NOTE: The only communication port for which conditions are reported is the CAN bus. Therefore, the element number is always 03.

Resetting the EMU

In some cases, the only corrective action for an EMU error is to replace the EMU. Call your authorized service representative if this action is required. Another option is to reset the EMU using the following procedure:

1. Firmly grasp the EMU mounting handle and pull the EMU partially out of the enclosure.

   NOTE: You do not need to remove the EMU from the enclosure or disconnect the cables. You must avoid putting any strain on the cables or connectors.

2. Wait 30 seconds, and then push the EMU in and fully seat the element in the backplane. The EMU should display any enclosure condition report within two minutes.

1.1.03.01 NONCRITICAL condition—Communication error

This condition report indicates that the EMU is unable to communicate over the CAN bus. Continued operation under these conditions may result in the failure to restore loop functionality when there is a disk drive disrupting the loop. This error initiates an automatic recovery process. This condition report remains active until the automatic recovery process is complete or until the EMU is reset.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Reset the EMU.
3. Observe the EMU to ensure the error does not recur within the first minute.
4. If the error does recur, contact your HP authorized service representative. The EMU is inoperative and must be replaced as soon as possible.

1.1.03.02 INFORMATION condition—Recovery completed

This condition report notes completion of the automatic recovery initiated by the occurrence of the 1.1.03.01 condition. This condition report remains active until one of the following occurs:

- 90 seconds elapses
- The CURRENT ALARM QUEUE is read via SES
- The RECENT ALARM LOG is read via SES

No action is required.

1.1.03.03 INFORMATION condition—Overrun recovery

This condition report notes automatic recovery initiated by the occurrence of too many data overruns with respect to received messages on the CAN bus. This condition report remains active until one of the following occurs:

- 90 seconds elapses
- The CURRENT ALARM QUEUE is read via SES
- The RECENT ALARM LOG is read via SES

No action is required.

Voltage sensor and current sensor conditions

The format of these sensor condition reports is *1.2.en.ec* for a voltage sensor, and *1.3.en.ec* for a current sensor, where:

- *1.2.* is the voltage sensor element type
- *1.3.* is the current sensor element type number
- *en.* is the sensor element number
- *ec* is the error code

Table 44 (page 131) lists the location of the power supply voltage and current sensors.

Table 44 Voltage and current sensor locations

Sensor	Sensor Element Location
01.	Power Supply 1 +5 VDC
02.	Power Supply 1 +12 VDC
03.	Power Supply 2 +5 VDC
04.	Power Supply 2 +12 VDC

Use HP P6000 Command View to view the voltage and current error thresholds for both +5 VDC and +12 VDC power supplies.

1.2.en.01 NONCRITICAL condition—High voltage

This condition report indicates that an element voltage is approaching, but has not reached, the high voltage CRITICAL threshold. Continued operation under these conditions may result in a CRITICAL condition. This condition report remains active until the problem is corrected.

To correct this problem, record all six characters of the condition report, then contact your HP-authorized service representative.

1.2.en.02 CRITICAL condition—High voltage

This condition report indicates that an element voltage has reached the high voltage CRITICAL threshold. This report remains active until the problem is corrected.

To correct this problem, record all six characters of the condition report, then contact your HP-authorized service representative.

1.2.en.03 NONCRITICAL condition—Low voltage

This condition report indicates that an element voltage is approaching, but has not reached, the low voltage CRITICAL threshold. Continued operation under these conditions may result in a CRITICAL condition. This condition report remains active until the problem is corrected.

To correct this problem, record all six characters of the condition report, then contact your HP-authorized service representative.

1.2.en.04 CRITICAL condition—Low voltage

This condition report indicates that an element voltage has reached the low voltage CRITICAL threshold. This condition report remains active until the problem is corrected.

To correct this problem, record all six characters of the condition report, then contact your HP-authorized service representative.

1.3.en.01 NONCRITICAL condition—High current

This condition report indicates that an element current is approaching, but has not reached, the high current CRITICAL threshold. Continued operation under these conditions may result in a CRITICAL condition. This condition report remains active until the problem is corrected.

To correct this problem, record all six characters of the condition report, then contact your HP-authorized service representative.

1.3.en.02 CRITICAL condition—High current

This condition report indicates that an element current has reached the high current CRITICAL threshold. This condition report remains active until the problem is corrected.

To correct this problem, record all six characters of the condition report, then contact your HP-authorized service representative.

Backplane conditions

IMPORTANT: Backplane NVRAM errors usually occur during manufacture. At this time they are identified and corrected. They rarely occur during normal operation.

The format of a backplane condition report is *8.2.01.ec*, where:

- *8.2.* is the backplane element type number
- *01.* is the two-character backplane element number
- *ec* is the error code

The only corrective action available for this error is to replace the drive enclosure.

8.2.01.10 NONCRITICAL condition—Backplane NVRAM read

An invalid NVRAM read occurred and an automatic recovery process has begun. This condition report is active for 15 seconds.

If the automatic recovery process does not correct the problem, record all six characters of the condition report, then contact your HP-authorized service representative.

8.2.01.11 NONCRITICAL condition—Backplane NVRAM write failure

The system is unable to write data to the NVRAM. This problem prevents communication between elements in the enclosure. This condition report is active for 15 seconds.

To correct this problem, record all six characters of the condition report, then contact your HP-authorized service representative.

8.2.01.12 NONCRITICAL condition—Backplane NVRAM read failure

The system is unable to read data from the NVRAM. This problem prevents communication between elements in the enclosure. This condition report is active for 15 seconds.

To correct this problem, record all six characters of the condition report, then contact your HP-authorized service representative.

8.2.01.13 NONCRITICAL condition—Backplane WWN is blank

The system is unable to read valid data from the NVRAM. This report is active until corrected. This condition can result in incorrect device location data being displayed.

To correct this problem, record all six characters of the condition report, then contact your HP-authorized service representative.

I/O Module conditions

The format of an I/O module condition report is *8.7.en.ec*, where:

- *8.7.* is the I/O module element type number
- *en.* is the two-character I/O module element number (see Figure 44 (page 133))
- *ec* is the error code

Figure 44 I/O module element numbering

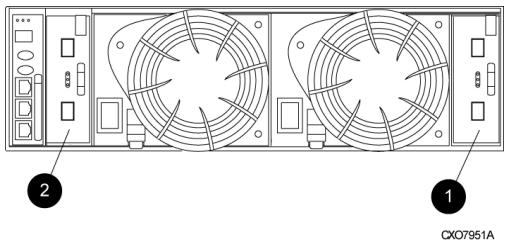

1. I/O Module A (01)
2. I/O Module B (02)

Correction of an I/O module problem normally requires replacing the module. The following sections define the I/O module problem by I/O module location.

8.7.en.01 CRITICAL condition—I/O module unsupported

The I/O module Fibre Channel link speed is not supported by the backplane. This error prevents the controller from establishing a link with enclosure drives and eliminates the enclosure dual-loop capability. This condition report remains active until the problem is corrected.

To correct this problem, record all six characters of the condition report, then contact your HP-authorized service representative.

8.7.en.02 CRITICAL condition—I/O module communication

The I/O module is unable to communicate with the EMU.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Contact your HP-authorized service representative.

IMPORTANT: Multiple erroneous error messages indicating I2C bus errors, such as NVRAM errors, blowers missing, and so forth, could indicate an EMU problem.

8.7.en.10 NONCRITICAL condition—I/O module NVRAM read

An invalid NVRAM read occurred and automatic recovery was initiated.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Observe the I/O module status indicators for an operational display.
3. Contact your HP-authorized service representative.

8.7.en.11 NONCRITICAL condition—I/O module NVRAM write

The system is unable to write data to the I/O module NVRAM.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Observe the I/O module status indicators for an operational display.

3. Contact your HP-authorized service representative.

8.7.en.12 NONCRITICAL condition—I/O Module NVRAM read failure

The system is unable to read data from the I/O module NVRAM.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Contact your HP-authorized service representative.

8.7.en.13 NONCRITICAL condition—I/O module removed

The system detects that an I/O module has been removed.

To correct the problem, install an I/O module.

Host conditions

The EMU has the capability of displaying host controller defined condition reports on the EMU alphanumeric display.

The format of a host condition report is *F.F.en.ec*, where:

- *F.F.* is the host element type number
- *en.* is the two-character host element number
- *ec* is the error code

The host controller can display host controller defined error codes on the EMU alphanumeric display

C Controller fault management

This appendix describes how the controller displays events and termination event information. Termination event information is displayed on the LCD. HP P6000 Command View enables you to view controller events. This appendix also discusses how to identify and correct problems.

Once you create a storage system, an error condition message has priority over other controller displays.

HP P6000 Command View provides detailed descriptions of the storage system error conditions, or faults. The Fault Management displays provide similar information on the LCD, but not in as much detail. Whenever possible, see HP P6000 Command View for fault information.

Using HP P6000 Command View

HP P6000 Command View provides detailed information about each event affecting system operation in either a Termination Event display or an Event display. These displays are similar, but not identical.

GUI termination event display

A problem that generates the Termination Event display prevents the system from performing a specific function or process. You can use the information in this display (see Figure 45 (page 135)) to diagnose and correct the problem.

NOTE: The major differences between the Termination Event display and the Event display are:

- The Termination Event display includes a Code Flag field; it does not include the EIP Type field.
- The Event display includes an EIP type field; it does not include a Code Flag field.
- The Event display includes a Corrective Action Code field.

Figure 45 GUI termination event display

Date	Time	SWCID	Evt No	Code Flag	Description

The fields in the Termination Event display include:

- Date—The date the event occurred.
- Time—The time the event occurred.
- SWCID—Software Identification Code. A hexadecimal number in the range 0–FF that identifies the controller software component reporting the event.
- Evt No—Event Number. A hexadecimal number in the range 0–FF that is the software component identification number.
- Code Flag—An internal code that includes a combination of other flags.
- Description—The condition that generated the event. This field may contain information about an individual field's content and validity.

GUI event display

A problem that generates the Event display reduces the system capabilities. You can use the information in this display (see Figure 46 (page 136)) to diagnose and correct problems.

NOTE: The major differences between the Event Display and the Termination Event display are:

- The Event display includes an EIP type field; it does not include a Code Flag field.
- The Event display includes a Corrective Action Code (CAC) field.
- The Termination Event display includes a Code Flag field; it does not include the EIP Type field.

Figure 46 Typical HP P6000 Command View Event display

Date	Time	SWCID	Evt No	CAC	EIP Type	Description

The Event display provides the following information:

- Date—The date the event occurred.
- Time—The time the even occurred.
- SWCID—Software Identification Code. A number in the range 1–256 that identifies the internal firmware module affected.
- Evt No—Event Number. A hexadecimal number in the range 0–FF that is the software component identification number.
- CAC—Corrective Action Code. A specific action to correct the problem.
- EIP Type—Event Information Packet Type. A hexadecimal character that defines the event information format.
- Description—The problem that generated the event.

Fault management displays

When you do not have access to the GUI, you can display and analyze termination codes (TCs) on the OCP LCD display. You can then use the event text code document, as described in the section titled "Interpreting Fault Management Information" to determine and implement corrective action. You can also provide this information to the authorized service representative should you require additional support. This lets the service representative identify the tools and components required to correct the condition in the shortest possible time.

When the fault management display is active, you can either display the last fault or display detailed information about the last 32 faults reported.

Displaying Last Fault Information

Complete the following procedure to display Last Fault information

1. When the Fault Management display is active, press ▶ to select the Last Fault menu.
2. Press ▶ to display the last fault information.

   The first line of the TC display contains the eight-character TC error code and the two-character IDX (index) code. The IDX is a reference to the location in the TC array that contains this error. The second line of the TC display identifies the affected parameter with a two-character parameter number (0–30), the eight-character parameter code affected, and the parameter code number.
3. Press ◀ to return to the Last Fault menu.

Displaying Detailed Information

The Detail View menu lets you examine detailed fault information stored in the Last Termination Event Array (LTEA). This array stores information for the last 32 termination events.

Complete the following procedure to display the LTEA information about any of the last 32 termination events:

1. When the Fault Management display is active (flashing), press ▼to select the Detail View menu.

   The LTEA selection menu is active (LTEA 0 is displayed).
2. Press ▼ or ▲ to increment to a specific error.
3. Press ▶ to observe data about the selected error.

Interpreting fault management information

Each version of HP P6000 Command View includes an ASCII text file that defines all the codes that the authorized service representative can view either on the GUI or on the OCP.

IMPORTANT: This information is for the exclusive use of the authorized service representative.

The file name identifies the controller model, file type, XCS baselevel id, and XCS version. For example, the file name *hsv210_event_cr08d3_5020.txt* provides the following information:

- `hsv210_`—The EVA controller model number
- `event_`—The type of information in the file
- `w010605_`—The base level build string (the file creation date).
  - `01`—The creation year
  - `06`—The creation month
  - `05`—The creation date
- `5020`—The XCS version

Table 45 (page 137) describes types of information available in this file.

Table 45 Controller event text description file

Information type	Description
Event Code	This hexadecimal code identifies the reported event type.
Termination Code (TC)	This hexadecimal code specifies the condition that generated the termination code. It might also define either a system or user initiated corrective action.
Coupled Crash Control Codes	This single digit, decimal character defines the requirement for the other controller to initiate a coupled crash control.0. Other controller SHOULD NOT complete a coupled crash.1. Other controller SHOULD complete a coupled crash.
Dump/Restart Control Codes	This single decimal character (0, 1, 3) defines the requirement to:0. Perform a crash dump and then restart the controller.1. DO NOT perform a crash dump; just restart the controller.3. DO NOT perform a crash dump; DO NOT restart the controller.
Corrective Action Codes (CAC)	These hexadecimal codes supplement the Termination Code information to identify the faulty element and the recommended corrective action.
Software Component ID Codes (SWCID)	These decimal codes identify software associated with the event.
Event Information Packets (EIP)	These codes specify the packet organization for specific type events.

D Non-standard rack specifications

The appendix provides information on the requirements when installing the EVA4x00/6x00/8x00 in a non-standard rack. All the requirements must be met to ensure proper operation of the storage system.

Rack specifications

Internal component envelope

EVA component mounting brackets require space to be mounted behind the vertical mounting rails. Room for the mounting of the brackets includes the width of the mounting rails and needed room for any mounting hardware, such as screws, clip nuts, etc. Figure 47 (page 138) shows the dimensions required for the mounting space for the EVA product line. It does not show required space for additional HP components such as servers.

Figure 47 Mounting space dimensions

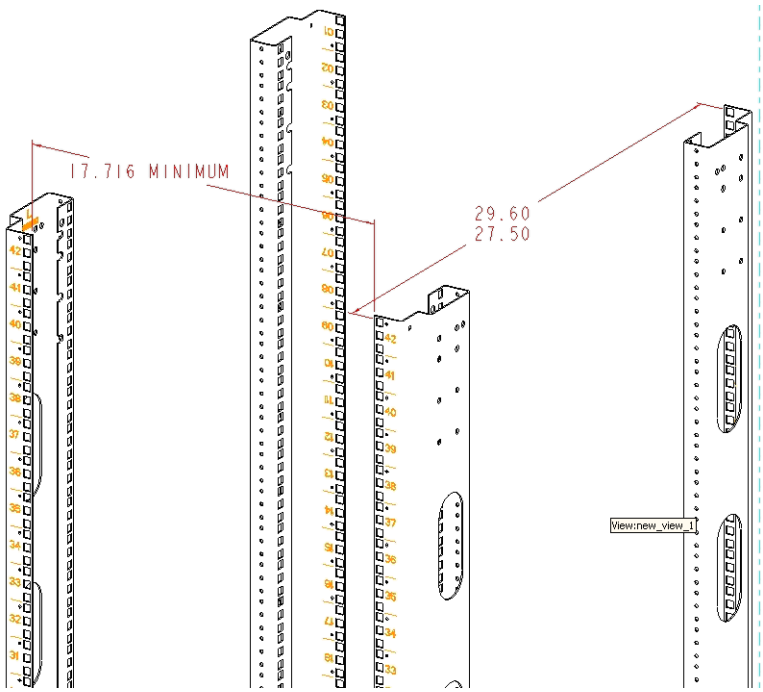

EIA310-D standards

The rack must meet the Electronic Industries Association, (EIA), Standard 310-D, Cabinets, Racks and Associated Equipment. The standard defines rack mount spacing and component dimensions specified in U units.

Copies of the standard are available for purchase at http://www.eia.org/.

EVA cabinet measures and tolerances

EVA component rack mount brackets are designed to fit cabinets with mounting rails set at depths from 27.5 inches to 29.6 inches, inside rails to inside rails.

Weights, dimensions and component CG measurements

Cabinet CG dimensions are reported as measured from the inside bottom of the cabinet (Z), the leading edge of the vertical mounting rails (Y), and the centerline of the cabinet mounting space (X). Component CG measurements are measured from the bottom of the U space the component is to occupy (Z), the mounting surface of the mounting flanges (Y), and the centerline of the component (X). Table 46 (page 139) lists the CG dimensions for the EVA components.

Determining the CG of a configuration may be necessary for safety considerations. CG considerations for CG calculations do not include cables, PDU's and other peripheral components. Some consideration should be made to allow for some margin of safety when estimating configuration CG.

Estimating the configuration CG requires measuring the CG of the cabinet the product will be installed in. Use the following formula:

$\Sigma d_{component} W = d_{system\ cg} W$

where $d_{component}$= the distance of interest and W = Weight

The distance of a component is its CG's distance from the inside base of the cabinet. For example, if a loaded disk enclosure is to be installed into the cabinet with its bottom at 10U, the distance for the enclosure would be (10*1.75)+2.7 inches.

Table 46 Component data

Component Data					
	U height[1]	**Weight (Lb)**	**X (in)**	**Y (in)**	**Z (in)**
HP 10K cabinet CG		233	-0.108	25.75	14.21
Filler panel , 3U	3	1.4	0	2.625	0
Fully loaded drive enclosure	3	74	-0.288	2.7	7.95
FC loop pair	1	24.6	-0.025	0.365	11.89
Filler panel, 1U	1	0.47	0	0.875	0
XL Controller Pair	4	120	-0.094	2.53	10.64

[1] 1U = 1.75 inches

Airflow and Recirculation

Component Airflow Requirements

Component airflow must be directed from the front of the cabinet to the rear. Components vented to discharge airflow from the sides must discharge to the rear of the cabinet.

Rack Airflow Requirements

The following requirements must be met to ensure adequate airflow and to prevent damage to the equipment:

- If the rack includes closing front and rear doors, allow 830 square inches (5,350 sq cm) of hole evenly distributed from top to bottom to permit adequate airflow (equivalent to the required 64 percent open area for ventilation).
- For side vented components, the clearance between the installed rack component and the side panels of the rack must be a minimum of 2.75 inches (7 cm).
- Always use blanking panels to fill all empty front panel U-spaces in the rack. This ensures proper airflow. Using a rack without blanking panels results in improper cooling that can lead to thermal damage.

Configuration Standards

EVA configurations are designed considering cable length, configuration CG, serviceability and accessibility, and to allow for easy expansion of the system. If at all possible, it is best to configure non HP cabinets in a like manner.

Environmental and operating specifications

This section identifies the product environmental and operating specifications.

NOTE: Further testing is required to update the information in Tables 45-47. Once testing is complete, these tables will be updated in a future release.

Power requirements

The following tables list the wattage and BTU/hour power requirements for the three supported operating voltages.

NOTE: Failover amperage can be estimated at approximately 90% of operational amperage listed.

Table 47 208V Wattage and BTU/Hour

Enclosures	EVA4x00				EVA6x00				EVA8x00			
	Amps	VA	Watts	BTU/h	Amps	VA	Watts	BTU/h	Amps	VA	Watts	BTU/h
12									24.5	5104	4920	16789
11									22.0	4578	4414	15060
10									20.1	4188	4037	13775
9									18.3	3797	3660	12489
8					16.0	3334	3214	10965	16.4	3406	3284	11204
7					14.1	2943	2837	9680	14.5	3015	2907	9919
6					12.3	2552	2460	8394	12.6	2625	2530	8633
5					10.4	2161	2083	7109	10.7	2234	2153	7348
4	8.2	1698	1637	5585	8.5	1770	1707	5824	8.9	1843	1777	6063
3	6.3	1307	1260	4300	6.6	1380	1330	4538	7.0	1452	1400	4777
2	4.4	916	883	3014	4.8	989	953	3253	5.1	1062	1023	3492
1	2.5	526	507	1729	2.9	598	577	1968	3.2	671	647	2207

Table 48 230V Wattage and BTU/Hour

Enclosures	EVA4x00				EVA6x00				EVA8x00			
	Amps	VA	Watts	BTU/h	Amps	VA	Watts	BTU/h	Amps	VA	Watts	BTU/h
12									21.6	4969	4790	16345
11									19.9	4578	4414	15060
10									18.2	4188	4037	13775
9									16.5	3797	3660	12489
8					14.5	3334	3214	10965	14.8	3406	3284	11204
7					12.8	2943	2837	9680	13.1	3015	2907	9919
6					11.1	2552	2460	8394	11.4	2625	2530	8633
5					9.4	2161	2083	7109	9.7	2234	2153	7348
4	7.4	1698	1637	5585	7.7	1770	1707	5824	8.0	1843	1777	6063
3	5.7	1307	1260	4300	6.0	1380	1330	4538	6.3	1452	1400	477
2	4.0	916	883	3014	4.3	989	953	3253	4.6	1062	1023	3492
1	2.3	526	507	1729	2.6	598	577	1968	2.9	671	647	2207

Table 49 100V Wattage and BTU/Hour

Enclosures	EVA4x00				EVA6x00				EVA8x00			
	Amps	VA	Watts	BTU/h	Amps	VA	Watts	BTU/h	Amps	VA	Watts	BTU/h
8					35.5	3545	3474	11855	EVA8x00 not supported			
7					31.5	3145	3082	10518				
6					27.5	2746	2691	9181				
5					23.5	2346	2299	7845				
4	18.7	1875	1837	6269	19.5	1946	1907	6508				
3	14.8	1475	1446	4933	15.5	1546	1516	5171				
2	10.8	1075	1054	3596	11.5	1147	1124	3835				
1	6.8	676	662	2259	7.5	747	732	2498				

UPS Selection

This section provides information that can be used when selecting a UPS for use with the EVA. The four HP UPS products listed in Table 50 (page 141) are available for use with the EVA and are included in this comparison. Table 51 (page 141) identifies the amount of time each UPS can sustain power under varying loads and with various UPS ERM (Extended Runtime Module) options.

The load imposed on the UPS for different disk enclosure configurations are listed in Table 52 (page 142), Table 53 (page 142), and Table 54 (page 143).

NOTE: The specified power requirements reflect fully loaded enclosures (14 disks) .

Table 50 HP UPS models and capacities

UPS Model	Capacity (in watts)
R1500	1340
R3000	2700
R5500	4500
R12000	12000

Table 51 UPS operating time limits

Load (percent)	Minutes of operation		
	With standby battery	With 1 ERM	With 2 ERMs
R1500			
100	5	23	49
80	6	32	63
50	13	57	161
20	34	146	290
R3000			
100	5	20	
80	6.5	30	
50	12	45	
20	40	120	

Table 51 UPS operating time limits *(continued)*

Load (percent)	Minutes of operation		
	With standby battery	With 1 ERM	With 2 ERMs
R5500			
100	7	24	46
80	9	31	60
50	19	61	106
20	59	169	303
R12000			
100	5	11	18
80	7	15	24
50	14	28	41
20	43	69	101

Table 52 EVA 8x00 UPS loading

Enclosures	Watts	% of UPS capacity	
		R5500	R12000
12	4920		41.0
11	4414	98.1	36.8
10	4037	89.7	33.6
9	3660	81.3	30.5
8	3284	73.0	27.4
7	2907	64.6	24.2
6	2530	56.2	21.1
5	2153	47.9	17.9
4	1777	39.5	14.8
3	1400	31.1	11.7
2	1023	22.7	8.5
1	647	14.4	5.4

Table 53 EVA 6x00 UPS loading

Enclosures	Watts	% of UPS capacity		
		R3000	R5500	R12000
8	3214		71.4	26.8
7	2837		63.0	23.6
6	2460	91.1	54.6	20.5
5	2083	77.2	46.2	17.3
4	1707	63.2	37.9	14.2
3	1330	49.3	29.5	11.1

Table 53 EVA 6x00 UPS loading *(continued)*

Enclosures	Watts	% of UPS capacity		
		R3000	R5500	R12000
2	953	35.3	21.2	7.9
1	577	21.4	12.8	4.8

Table 54 EVA 4x00 UPS loading

Enclosures	Watts	% of UPS capacity	
		R1500	R3000
4	1637		60.6
3	1260	94.0	46.6
2	883	65.9	32.7
1	507	37.9	18.7

Environmental specifications

Table 55 Environmental specifications

Operating temperature	50° to 95° F (10° to 35° C) - Reduce rating by 1° F for each 1000 ft. altitude (1.8° C/1,000 m)
Shipping temperature	-40° to 150° F (-40° to 66° C)
Operating humidity	10% to 90% non-condensing
Shipping humidity	5% to 90% non-condensing
Altitude	Up to 8,000 ft. (2,400 m)
Air quality	Not to exceed 500,000 particles per cubic foot of air at a size of 0.5 micron or larger

Shock and vibration specifications

Table 56 (page 144) lists the product operating shock and vibration specifications. This information applies to products weighing 45 Kg (100 lbs) or less.

NOTE: HP EVA products are designed and tested to withstand the operational shock and vibration limits specified in Table 56 (page 144). Transmission of site vibrations through non-HP racks exceeding these limits could cause operational failures of the system components.

Table 56 Operating Shock/Vibration

Shock test with half sine pulses of 10 G magnitude and 10 ms duration applied in all three axes (both positive and negative directions).
Sine sweep vibration from 5 Hz to 500 Hz to 5 Hz at 0.1 G peak, with 0.020" displacement limitation below 10 Hz. Sweep rate of 1 octave/minute. Test performed in all three axes.
Random vibration at 0.25 G rms level with uniform spectrum in the frequency range of 10 to 500 Hz. Test performed for two minutes each in all three axes.
Drives and other items exercised and monitored running appropriate exerciser (UIOX, P-Suite, etc.) with appropriate operating system and hardware.

E Single Path Implementation

This appendix provides guidance for connecting servers with a single path host bus adapter (HBA) to the Enterprise Virtual Array (EVA) storage system with no multi-path software installed. A single path HBA is defined as an HBA that has a single path to its LUNs. These LUNs are not shared by any other HBA in the server or in the SAN.

The failure scenarios demonstrate behavior when recommended configurations are employed, as well as expected failover behavior if guidelines are not met. To implement single adapter servers into a multi-path EVA environment, configurations should follow these recommendations.

NOTE: The purpose of single HBA configurations for non-mission critical storage access is to control costs. This appendix describes the configurations, limitations, and failover characteristics of single HBA servers under different operating systems. Much of the description herein are based upon a single HBA configuration resulting in a single path to the device, but such is not the case with OpenVMS and Tru64 UNIX.

HP OpenVMS and Tru64 UNIX have native multi-path features by default.

With OpenVMS and Tru64 UNIX, a single HBA configuration will result in two paths to the device by virtue of having connections to both EVA controllers. Single HBA configurations are not single path configurations with these operating systems.

In addition, cluster configurations of both OpenVMS and Tru64 UNIX provide enhanced availability and security. To achieve availability within cluster configurations, each member should be configured with its own HBA(s) and connectivity to shared LUNs. Cluster configuration will not be discussed further within this appendix as the enhanced availability requires both additional server hardware and HBAs which is contrary to controlling configuration costs for non-mission critical applications. For further information on cluster configurations and attributes, see the appropriate operating system guides and the SAN design guide.

NOTE: HP continually makes additions to its storage solution product line. For more information about the HP Fibre Channel product line, the latest drivers, and technical tips, and to view other documentation, see the HP website at

http://www.hp.com/country/us/eng/prodserv/storage.html

High-level solution overview

EVA was designed for highly dynamic enterprise environments requiring high data availability, fault tolerance, and high performance; thus, the EVA controller runs only in multi-path failover mode. Multi-path failover mode ensures the proper level of fault tolerance for the enterprise with mission-critical application environments. However, this appendix addresses the need for non-mission-critical applications to gain access to the EVA system running mission-critical production applications.

The non-mission-critical applications gain access to the EVA from a single path HBA server without running a multi-path driver. When a single path HBA server uses the supported configurations, a fault in the single path HBA server does not result in a fault in the other servers.

Benefits at a glance

The EVA is a high-performance array controller utilizing the benefits of virtualization. Virtualization within the storage system is ideal for environments needing high performance, high data availability, fault tolerance, efficient storage management, data replication, and cluster support. However, enterprise-level data centers incorporate non-mission-critical applications as well as applications that require high availability.

Single-path capability adds flexibility to budget allocation. There is a per-path savings as the additional cost of HBAs and multi-path software is removed from non-mission–critical application requirements. These servers can still gain access to the EVA by using single path HBAs without multi-path software. This reduces the costs at the server and infrastructure level.

Installation requirements

- The host must be placed in a zone with any EVA worldwide IDs (WWIDs) that access storage devices presented by the hierarchical storage virtualization (HSV) controllers to the single path HBA host. The preferred method is to use HBA and HSV WWIDs in the zone configurations.
- On HP-UX, Solaris, Microsoft Windows Server 2003 (32-bit), , Linux and IBM AIX operating systems, the zones consist of the single path HBA systems and one HSV controller port.
- On OpenVMS and Tru64 UNIX operating systems, the zones consist of the single HBA systems and two HSV controller ports. This will result in a configuration where there are two paths per device, or multiple paths.

Recommended mitigations

EVA is designed for the mission-critical enterprise environment. When used with multi-path software, high data availability and fault tolerance are achieved. In single path HBA server configurations, neither multi-path software nor redundant I/O paths are present. Server-based operating systems are not designed to inherently recover from unexpected failure events in the I/O path (for example, loss of connectivity between the server and the data storage). It is expected that most operating systems will experience undesirable behavior when configured in non-high-availability configurations.

Because of the risks of using servers with a single path HBA, HP recommends the following actions:

- Use servers with a single path HBA that are not mission-critical or highly available.
- Perform frequent backups of the single path server and its storage.

Supported configurations

All examples detail a small homogeneous Storage Area Network (SAN) for ease of explanation. Mixing of dual and single path HBA systems in a heterogeneous SAN is supported. In addition to this document, reference and adhere to the SAN Design Reference Guide for heterogeneous SANs, located at:

http://h18006.www1.hp.com/products/storageworks/san/documentation.html

General configuration components

All configurations require the following components:

- Enterprise VCS software
- HBAs
- Fibre Channel switches

Connecting a single path HBA server to a switch in a fabric zone

Each host must attach to one switch (fabric) using standard Fibre Channel cables. Each host has its single path HBA connected through switches on a SAN to one port of an EVA.

Because a single path HBA server has no software to manage the connection and ensure that only one controller port is visible to the HBA, the fabric containing the single path HBA server, SAN switch, and EVA controller must be zoned. Configuring the single path by switch zoning and the LUNs by Selective Storage Presentation (SSP) allows for multiple single path HBAs to reside in the same server. A single path HBA server with OpenVMS or Tru64 UNIX operating system should be zoned with two EVA controllers. See the HP SAN Design Reference Guide at the following HP website for additional information about zoning:

http://h18006.www1.hp.com/products/storageworks/san/documentation.html

To connect a single path HBA server to a SAN switch:

1. Plug one end of the Fibre Channel cable into the HBA on the server.
2. Plug the other end of the cable into the switch.

Figure 48 (page 147) and Figure 49 (page 147) represent configurations containing both single path HBA server and dual HBA server, as well as a SAN appliance, connected to redundant SAN

switches and EVA controllers. Whereas the dual HBA server has multi-path software that manages the two HBAs and their connections to the switch (with the exception of OpenVMS and Tru64 UNIX servers), the single path HBA has no software to perform this function. The dashed line in the figure represents the fabric zone that must be established for the single path HBA server. Note that in Figure 49 (page 147), servers with OpenVMS or Tru64 UNIX operating system should be zoned with two controllers.

Figure 48 Single path HBA server without OpenVMS or Tru64 UNIX

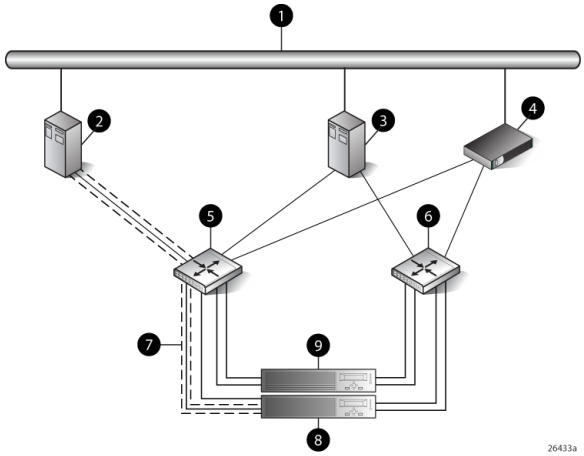

1 Network interconnection
2 Single HBA server
3 Dual HBA server
4 Management server
5 SAN switch 1
6 SAN switch 2
7 Fabric zone
8 Controller A
9 Controller B

Figure 49 Single path HBA server with OpenVMS or Tru64 UNIX

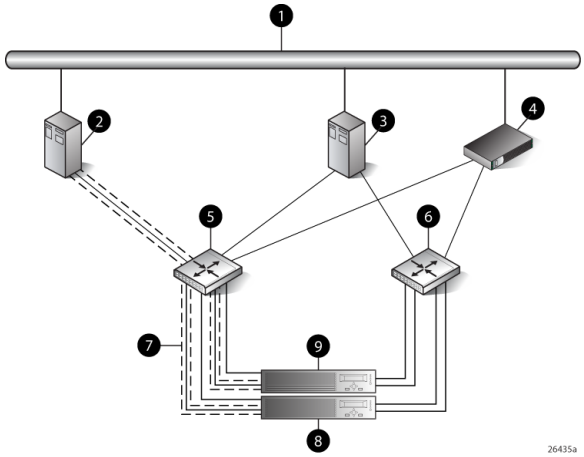

1 Network interconnection
2 Single HBA server
3 Dual HBA server
4 Management server
5 SAN switch 1
6 SAN switch 2
7 Fabric zone
8 Controller A
9 Controller B

HP-UX configuration

Requirements

- Proper switch zoning must be used to ensure each single path HBA has an exclusive path to its LUNs.
- Single path HBA server can be in the same fabric as servers with multiple HBAs.
- Single path HBA server cannot share LUNs with any other HBAs.
- In the use of snapshots and snapclones, the source virtual disk and all associated snapshots and snapclones must be presented to the single path hosts that are zoned with the same controller. In the case of snapclones, after the cloning process has completed and the clone becomes an ordinary virtual disk, you may present that virtual disk as you would any other ordinary virtual disk.

HBA configuration

- Host 1 is a single path HBA host.
- Host 2 is a multiple HBA host with multi-pathing software.

See Figure 50 (page 148).

Risks

- Disabled jobs hang and cannot umount disks.
- Path or controller failure may results in loss of data accessibility and loss of host data that has not been written to storage.

NOTE: For additional risks, see Table 57 (page 158).

Limitations

- HP P6000 Continuous Access is not supported with single-path configurations.
- Single path HBA server is not part of a cluster.
- Booting from the SAN is not supported.

Figure 50 HP-UX configuration

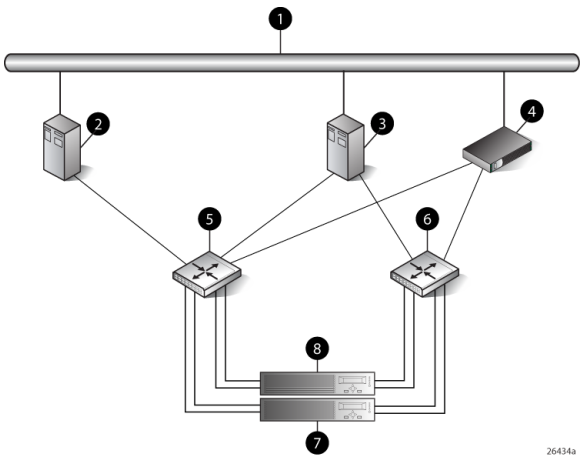

1 Network interconnection
2 Host 1
3 Host 2
4 Management server
5 SAN switch 1
6 SAN switch 2
7 Controller A
8 Controller B

Windows Server (32-bit) configuration

Requirements

- Switch zoning or controller level SSP must be used to ensure each single path HBA has an exclusive path to its LUNs.
- Single path HBA server can be in the same fabric as servers with multiple HBAs.
- Single path HBA server cannot share LUNs with any other HBAs.
- In the use of snapshots and snapclones, the source virtual disk and all associated snapshots and snapclones must be presented to the single path hosts that are zoned with the same controller. In the case of snapclones, after the cloning process has completed and the clone becomes an ordinary virtual disk, you may present that virtual disk as you would any other ordinary virtual disk.

HBA configuration

- Host 1 is a single path HBA host.
- Host 2 is a multiple HBA host with multi-pathing software.

See Figure 51 (page 149).

Risks

- Single path failure will result in loss of connection with the storage system.
- Single path failure may cause the server to reboot.
- Controller shutdown puts controller in a failed state that results in loss of data accessibility and loss of host data that has not been written to storage.

NOTE: For additional risks, see Table 58 (page 159).

Limitations

- HP P6000 Continuous Access is not supported with single path configurations.
- Single path HBA server is not part of a cluster.
- Booting from the SAN is not supported on single path HBA servers.

Figure 51 Windows Server (32-bit) configuration

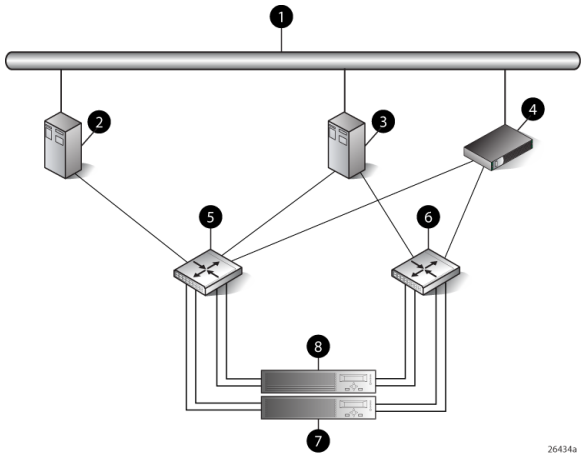

1 Network interconnection
2 Host 1
5 SAN switch 1
6 SAN switch 2

3 Host 2	7 Controller A
4 Management server	8 Controller B

Windows Server (64-bit) configuration

Requirements

- Switch zoning or controller level SSP must be used to ensure each single path HBA has an exclusive path to its LUNs.
- Single path HBA server can be in the same fabric as servers with multiple HBAs.
- Single path HBA server cannot share LUNs with any other HBAs.

HBA configuration

- Hosts 1 and 2 are single path HBA hosts.
- Host 3 is a multiple HBA host with multi-pathing software.

See Figure 52 (page 151).

NOTE: Single path HBA servers running the Windows Server 2003 (x64) operating system will support multiple single path HBAs in the same server. This is accomplished through a combination of switch zoning and controller level SSP. Any single path HBA server will support up to four single path HBAs.

Risks

- Single path failure will result in loss of connection with the storage system.
- Single path failure may cause the server to reboot.
- Controller shutdown puts controller in a failed state that results in loss of data accessibility and loss of host data that has not been written to storage.

NOTE: For additional risks, see Table 58 (page 159).

Limitations

- HP P6000 Continuous Access is not supported with single path configurations.
- Single path HBA server is not part of a cluster.
- Booting from the SAN is not supported on single path HBA servers.

Figure 52 Windows Server (64-bit) configuration

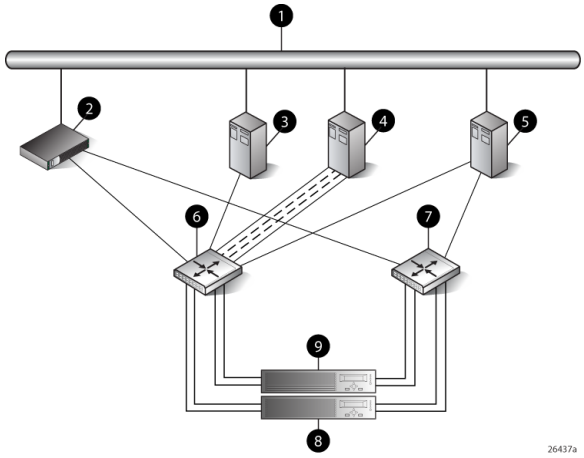

1 Network interconnection
2 Management server
3 Host 1
4 Host 2
5 Host 3
6 SAN switch 1
7 SAN switch 2
8 Controller A
9 Controller B

Oracle Solaris configuration

Requirements

- Switch zoning or controller level SSP must be used to ensure each single path HBA has an exclusive path to its LUNs.
- Single path HBA server can be in the same fabric as servers with multiple HBAs.
- Single path HBA server cannot share LUNs with any other HBAs.
- In the use of snapshots and snapclones, the source virtual disk and all associated snapshots and snapclones must be presented to the single path hosts that are zoned with the same controller. In the case of snapclones, after the cloning process has completed and the clone becomes an ordinary virtual disk, you may present that virtual disk as you would any other ordinary virtual disk.
- HBA must be properly configured to work in a single HBA server configuration. The user is required to:
  - Download and extract the contents of the TAR file.

HBA configuration

- Host 1 is a single path HBA host.
- Host 2 is a multiple HBA host with multi-pathing software.

See Figure 53 (page 152).

Risks

- Single path failure may result in loss of data accessibility and loss of host data that has not been written to storage.
- Controller shutdown results in loss of data accessibility and loss of host data that has not been written to storage.

NOTE: For additional risks, see Table 59 (page 159).

Limitations

- HP P6000 Continuous Access is not supported with single path configurations.
- Single path HBA server is not part of a cluster.
- Booting from the SAN is not supported.

Figure 53 Oracle Solaris configuration

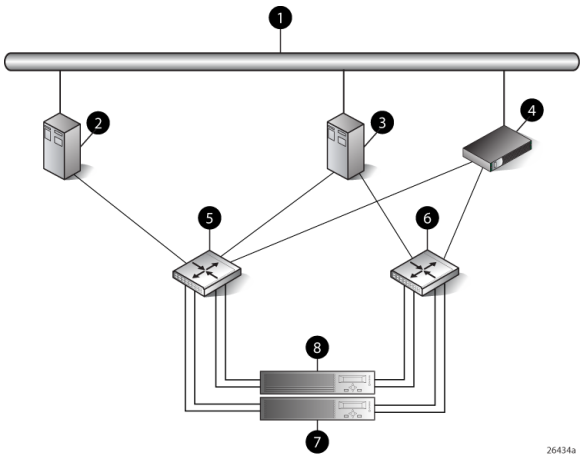

1 Network interconnection	5 SAN switch 1
2 Host 1	6 SAN switch 2
3 Host 2	7 Controller A
4 Management server	8 Controller B

Tru64 UNIX configuration

Requirements

- Switch zoning or controller level SSP must be used to ensure each HBA has exclusive access to its LUNs.
- All nodes with direct connection to a disk must have the same access paths available to them.
- Single HBA server can be in the same fabric as servers with multiple HBAs.
- In the use of snapshots and snapclones, the source virtual disk and all associated snapshots and snapclones must be presented to the single host that are zoned with the same controller. In the case of snapclones, after the cloning process has completed and the clone becomes an ordinary virtual disk, you may present that virtual disk as you would any other ordinary virtual disk.

HBA configuration

- Host 1 is single HBA host with Tru64.
- Host 2 is a dual HBA host.

See Figure 54 (page 153).

Risks

- For nonclustered nodes with a single HBA, a path failure from the HBA to the SAN switch will result in a loss of connection with storage devices.
- If a host crashes or experiences a power failure, or if the path is interrupted, data will be lost. Upon re-establishment of the path, a retransmit can be performed to recover whatever data may have been lost during the outage. The option to retransmit data after interruption is application-dependent.

NOTE: For additional risks, see Table 60 (page 160).

Figure 54 Tru64 UNIX configuration

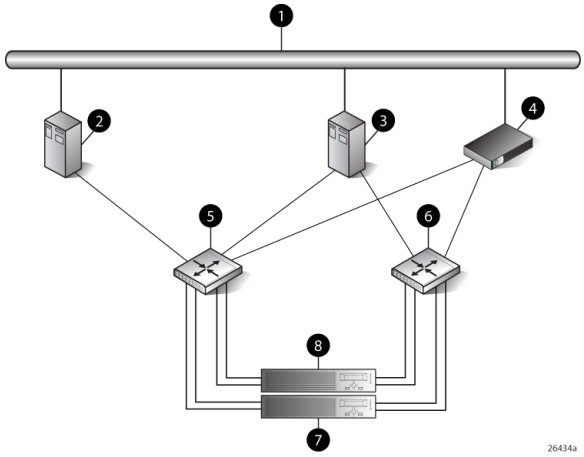

1 Network interconnection
2 Host 1
3 Host 2
4 Management server
5 SAN switch 1
6 SAN switch 2
7 Controller A
8 Controller B

OpenVMS configuration

Requirements

- Switch zoning or controller level SSP must be used to ensure each single path HBA has an exclusive path to its LUNs.
- All nodes with direct connection to a disk must have the same access paths available to them.
- Single path HBA server can be in the same fabric as servers with multiple HBAs.
- In the use of snapshots and snapclones, the source virtual disk and all associated snapshots and snapclones must be presented to the single path hosts that are zoned with the same controller. In the case of snapclones, after the cloning process has completed and the clone becomes an ordinary virtual disk, you may present that virtual disk as you would any other ordinary virtual disk.

HBA configuration

- Host 1 is a single path HBA host.
- Host 2 is a dual HBA host.

See Figure 55 (page 154).

Risks

- For nonclustered nodes with a single path HBA, a path failure from the HBA to the SAN switch will result in a loss of connection with storage devices.

NOTE: For additional risks, see Table 60 (page 160).

Limitations

- HP P6000 Continuous Access is not supported with single path configurations.

Figure 55 OpenVMS configuration

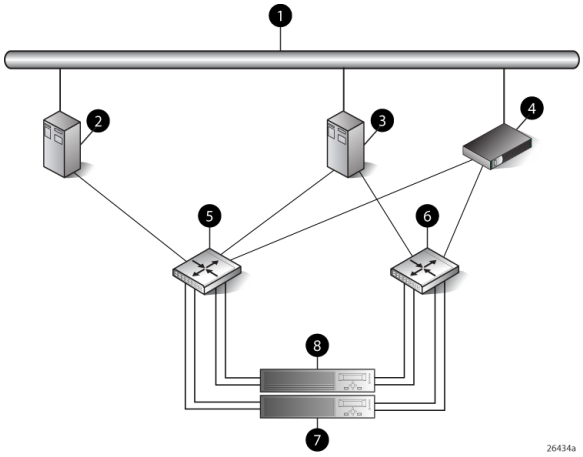

1 Network interconnection	5 SAN switch 1
2 Host 1	6 SAN switch 2
3 Host 2	7 Controller A
4 Management server	8 Controller B

Linux (32-bit) configuration

Requirements

- Switch zoning or controller level SSP must be used to ensure each single path HBA has an exclusive path to its LUNs.
- All nodes with direct connection to a disk must have the same access paths available to them.
- Single path HBA server can be in the same fabric as servers with multiple HBAs.
- In the use of snapshots and snapclones, the source virtual disk and all associated snapshots and snapclones must be presented to the single path hosts that are zoned with the same controller. In the case of snapclones, after the cloning process has completed and the clone becomes an ordinary virtual disk, you may present that virtual disk as you would any other ordinary virtual disk.

HBA configuration

- Host 1 is a single path HBA.
- Host 2 is a dual HBA host with multi-pathing software.

See Figure 56 (page 155).

Risks

- Single path failure may result in data loss or disk corruption.

NOTE: For additional risks, see Table 61 (page 160).

Limitations

- HP P6000 Continuous Access is not supported with single path configurations.
- Single path HBA server is not part of a cluster.
- Booting from the SAN is supported on single path HBA servers.

Figure 56 Linux (32-bit) configuration

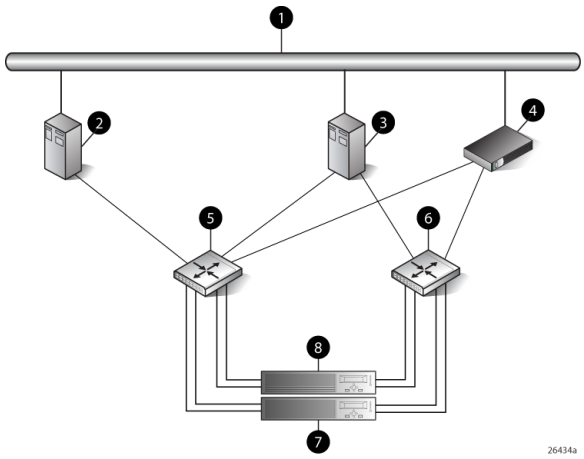

1 Network interconnection
2 Host 1
3 Host 2
4 Management server
5 SAN switch 1
6 SAN switch 2
7 Controller A
8 Controller

Linux (64-bit) configuration

Requirements

- Switch zoning or controller level SSP must be used to ensure each single path HBA has an exclusive path to its LUNs.
- All nodes with direct connection to a disk must have the same access paths available to them.
- Single path HBA server can be in the same fabric as servers with multiple HBAs.
- In the use of snapshots and snapclones, the source virtual disk and all associated snapshots and snapclones must be presented to the single path hosts that are zoned with the same controller. In the case of snapclones, after the cloning process has completed and the clone becomes an ordinary virtual disk, you may present that virtual disk as you would any other ordinary virtual disk.
- Linux 64-bit servers can support up to14 single or dual path HBAs per server. Switch zoning and SSP are required to isolate the LUNs presented to each HBA from each other.

HBA configuration

- Host 1 and 2 are single path HBA hosts.
- Host 3 is a dual HBA host with multi-pathing software.

See Figure 57 (page 156).

Risks

- Single path failure may result in data loss or disk corruption.

NOTE: For additional risks, see Table 61 (page 160).

Limitations

- HP P6000 Continuous Access is not supported with single path configurations.
- Single path HBA server is not part of a cluster.
- Booting from the SAN is supported on single path HBA servers.

Figure 57 Linux (64-bit) configuration

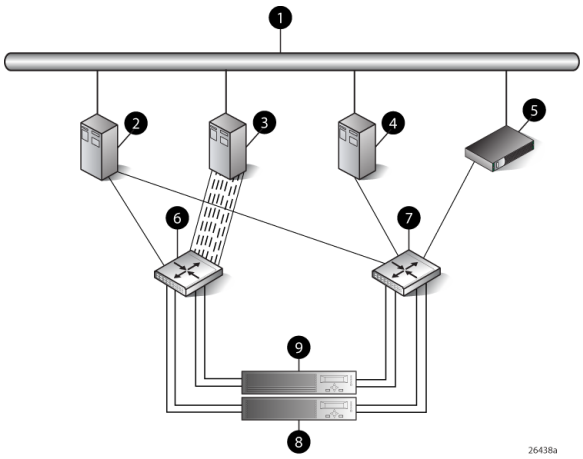

1 Network interconnection
2 Host 3
3 Host 2
4 Host 1
5 Management server
6 SAN switch 1
7 SAN switch 2
8 Controller A
9 Controller B

IBM AIX configuration

Requirements

- Switch zoning or controller level SSP must be used to ensure each single path HBA has an exclusive path to its LUNs.
- Single path HBA server can be in the same fabric as servers with multiple HBAs.
- Single path HBA server cannot share LUNs with any other HBAs.
- In the use of snapshots and snapclones, the source virtual disk and all associated snapshots and snapclones must be presented to the single path hosts that are zoned with the same controller. In the case of snapclones, after the cloning process has completed and the clone becomes an ordinary virtual disk, you may present that virtual disk as you would any other ordinary virtual disk.
- HBA must be properly configured to work in a single HBA server configuration. The single path adapter driver from the AIX 2.0B EVA Kit should be installed: PC1000.image.

HBA configuration

- Host 1 is a single path HBA host.
- Host 2 is a dual HBA host with multi-pathing software.

See Figure 58 (page 157).

Risks

- Single path failure may result in loss of data accessibility and loss of host data that has not been written to storage.
- Controller shutdown results in loss of data accessibility and loss of host data that has not been written to storage.

NOTE: For additional risks, see Table 62 (page 161).

Limitations

- HP P6000 Continuous Access is not supported with single path configurations.
- Single path HBA server is not part of a cluster.
- Booting from the SAN is not supported.

Figure 58 IBM AIX Configuration

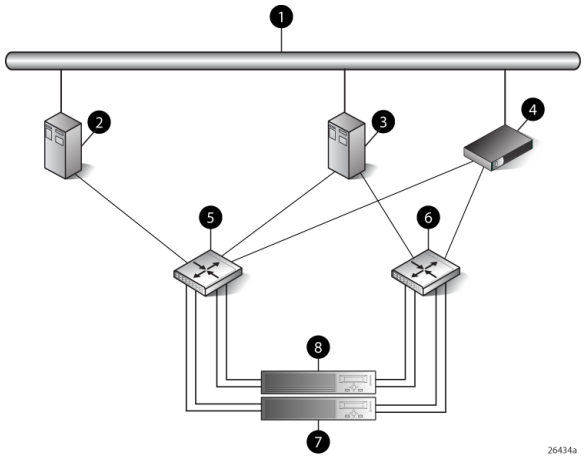

1 Network interconnection	5 SAN switch 1
2 Single HBA server	6 SAN switch 2
3 Dual HBA server	7 Controller A
4 Management server	8 Controller B

VMware configuration

Requirements

- Switch zoning or controller level SSP must be used to ensure each single path HBA has an exclusive path to its LUNs.
- All nodes with direct connection to a disk must have the same access paths available to them.
- Single path HBA server can be in the same fabric as servers with multiple HBAs.
- In the use of snapshots and snapclones, the source virtual disk and all associated snapshots and snapclones must be presented to the single path hosts that are zoned with the same controller. In the case of snapclones, after the cloning process has completed and the clone becomes an ordinary virtual disk, you may present that virtual disk as you would any other ordinary virtual disk.

HBA configuration

- Host 1 is a single path HBA.
- Host 2 is a dual HBA host with multi-pathing software.

See Figure 59 (page 158).

Risks

- Single path failure may result in data loss or disk corruption.

NOTE: For additional risks, see Table 63 (page 161).

Limitations

- HP P6000 Continuous Access is not supported with single path configurations.
- Single path HBA server is not part of a cluster.
- Booting from the SAN is supported on single path HBA servers.

Figure 59 VMware configuration

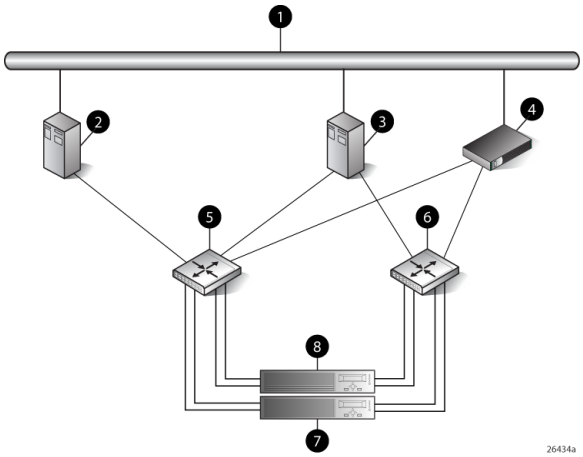

1 Network interconnection
2 Single HBA server
3 Dual HBA server
4 Management server
5 SAN switch 1
6 SAN switch 2
7 Controller A
8 Controller B

Failure scenarios

HP-UX

Table 57 HP-UX failure scenarios

Fault stimulus	Failure effect
Server failure (host power-cycled)	Extremely critical event on UNIX. Can cause loss of system disk.
Switch failure (SAN switch disabled)	Short term: Data transfer stops. Possible I/O errors. Long term: Job hangs, cannot umount disk, fsck failed, disk corrupted, need mkfs disk.
Controller failure	Short term: Data transfer stops. Possible I/O errors. Long term: Job hangs, cannot umount disk, fsck failed, disk corrupted, need mkfs disk.
Controller restart	Short term: Data transfer stops. Possible I/O errors. Long term: Job hangs, cannot umount disk, fsck failed, disk corrupted, need mkfs disk.

Table 57 HP-UX failure scenarios *(continued)*

Fault stimulus	Failure effect
Server path failure	Short term: Data transfer stops. Possible I/O errors. Long term: Job hangs, cannot umount disk, fsck failed, disk corrupted, need mkfs disk.
Storage path failure	Short term: Data transfer stops. Possible I/O errors. Long term: Job hangs, replace cable, I/O continues. Without cable replacement job must be aborted; disk seems error free.

Windows Server

Table 58 Windows Server failure scenarios

Fault stimulus	Failure effect
Server failure (host power-cycled)	OS runs a command called chkdsk when rebooting. Data lost, data that finished copying survived.
Switch failure (SAN switch disabled)	Write delay, server hangs until I/O is cancelled or cold reboot.
Controller failure	Write delay, server hangs or reboots. One controller failed, other controller and shelves critical, shelves offline. Volume not accessible. Server cold reboot, data lost. Check disk when rebooting.
Controller restart	Controller momentarily in failed state, server keeps copying. All data copied, no interruption. Event error warning error detected during paging operation.
Server path failure	Write delay, volume inaccessible. Host hangs and restarts.
Storage path failure	Write delay, volume disappears, server still running. When cables plugged back in, controller recovers, server finds volume, data loss.

Oracle Solaris

Table 59 Oracle Solaris failure scenarios

Fault stimulus	Failure effect
Server failure (host power-cycled)	Check disk when rebooting. Data loss, data that finished copying survived.
Switch failure (SAN switch disabled)	Short term: Data transfer stops. Possible I/O errors. Long term: Repeated error messages on console, no access to CDE. System reboot causes loss of data on disk. Must newfs disk.
Controller failure	Short term: Data transfer stops. Possible I/O errors. Long term: Repeated error messages on console, no access to CDE. System reboot causes loss of data on disk. Must newfs disk.
Controller restart	Short term: Data transfer stops. Possible I/O errors. Long term: Repeated error messages on console, no access to CDE. System reboot causes loss of data on disk. Must newfs disk.
Server path failure	Short term: Data transfer stops. Possible I/O errors. Long term: Repeated error messages on console, no access to CDE. System reboot causes loss of data on disk. Must newfs disk.
Storage path failure	Short term: Job hung, data lost. Long term: Repeated error messages on console, no access to CDE. System reboot causes loss of data on disk. Must newfs disk.

OpenVMS and Tru64 UNIX

Table 60 OpenVMS and Tru64 UNIX failure scenarios

Fault stimulus	Failure effect
Server failure (host power-cycled)	All I/O operations halted. Possible data loss from unfinished or unflushed writes. File system check may be needed upon reboot.
Switch failure (SAN switch disabled)	OpenVMS—OS will report the volume in a Mount Verify state until the MVTIMEOUT limit is exceeded, when it then marks the volume as Mount Verify Timeout. No data is lost or corrupted. Tru64 UNIX—All I/O operations halted. I/O errors are returned back to the applications. An I/O failure to the system disk can cause the system to panic. Possible data loss from unfinished or unflushed writes. File system check may be needed upon reboot.
Controller failure	I/O fails over to the surviving path. No data is lost or corrupted.
Controller restart	OpenVMS—OS will report the volume in a Mount Verify state until the MVTIMEOUT limit is exceeded, when it then marks the volume as Mount Verify Timeout. No data is lost of corrupted. Tru64 UNIX—I/O retried until controller back online. If maximum retries exceeded, I/O fails over to the surviving path. No data is lost or corrupted.
Server path failure	OpenVMS—OS will report the volume in a Mount Verify state until the MVTIMEOUT limit is exceeded, when it then marks the volume as Mount Verify Timeout. No data is lost or corrupted. Tru64 UNIX—All I/O operations halted. I/O errors are returned back to the applications. An I/O failure to the system disk can cause the system to panic. Possible data loss from unfinished or unflushed writes. File system check may be needed upon reboot.
Storage path failure	OpenVMS—OS will report the volume in a Mount Verify state until the MVTIMEOUT limit is exceeded, when it then marks the volume as Mount Verify Timeout. No data is lost or corrupted. Tru64 UNIX—I/O fails over to the surviving path. No data is lost or corrupted.

Linux

Table 61 Linux failure scenarios

Fault stimulus	Failure effect
Server failure (host power-cycled)	OS reboots, automatically checks disks. HSV disks must be manually checked unless auto mounted by the system.
Switch failure (SAN switch disabled)	Short: I/O suspended, possible data loss. Long: I/O halts with I/O errors, data loss. HBA driver must be reloaded before failed drives can be recovered, fsck should be run on any failed drives before remounting.
Controller failure	Short term: I/O suspended, possible data loss. Long term: I/O halts with I/O errors, data loss. Cannot reload driver, need to reboot system, fsck should be run on any failed disks before remounting.
Controller restart	Short term: I/O suspended, possible data loss. Long term: I/O halts with I/O errors, data loss. Cannot reload driver, need to reboot system, fsck should be run on any failed disks before remounting.

Table 61 Linux failure scenarios *(continued)*

Fault stimulus	Failure effect
Server path failure	Short: I/O suspended, possible data loss. Long: I/O halts with I/O errors, data loss. HBA driver must be reloaded before failed drives can be recovered, fsck should be run on any failed drives before remounting.
Storage path failure	Short: I/O suspended, possible data loss. Long: I/O halts with I/O errors, data loss. HBA driver must be reloaded before failed drives can be recovered, fsck should be run on any failed drives before remounting.

IBM AIX

Table 62 IBM AIX failure scenarios

Fault stimulus	Failure effect
Server failure (host power-cycled)	Check disk when rebooting. Data loss, data that finished copying survived
Switch failure (SAN switch disabled)	Short term: Data transfer stops. Possible I/O errors. Long term: Repeated error messages in errpt output. System reboot causes loss of data on disk. Must crfs disk.
Controller failure	Short term: Data transfer stops. Possible I/O errors. Long term: Repeated error messages in errpt output. System reboot causes loss of data on disk. Must crfs disk.
Controller restart	Short term: Data transfer stops. Possible I/O errors. Long term: Repeated error messages in errpt output. System reboot causes loss of data on disk. Must crfs disk.
Server path failure	Short term: Data transfer stops. Possible I/O errors. Long term: Repeated error messages in errpt output. System reboot causes loss of data on disk. Must crfs disk.
Storage path failure	Short term: Data transfer stops. Possible I/O errors. Long term: Repeated error messages in errpt output. System reboot causes loss of data on disk. Must crfs disk.

VMware

Table 63 VMware failure scenarios

Fault stimulus	Failure effect
Server failure (host power-cycled)	OS reboots, automatically checks disks. HSV disks must be manually checked unless auto mounted by the system.
Switch failure (SAN switch disabled)	Short: I/O suspended, possible data loss. Long: I/O halts with I/O errors, data loss. HBA driver must be reloaded before failed drives can be recovered, fsck should be run on any failed drives before remounting.
Controller failure	Short term: I/O suspended, possible data loss. Long term: I/O halts with I/O errors, data loss. Cannot reload driver, need to reboot system, fsck should be run on any failed disks before remounting.
Controller restart	Short term: I/O suspended, possible data loss. Long term: I/O halts with I/O errors, data loss. Cannot reload driver, need to reboot system, fsck should be run on any failed disks before remounting.

Table 63 VMware failure scenarios *(continued)*

Fault stimulus	Failure effect
Server path failure	Short: I/O suspended, possible data loss. Long: I/O halts with I/O errors, data loss. HBA driver must be reloaded before failed drives can be recovered, fsck should be run on any failed drives before remounting.
Storage path failure	Short: I/O suspended, possible data loss. Long: I/O halts with I/O errors, data loss. HBA driver must be reloaded before failed drives can be recovered, fsck should be run on any failed drives before remounting.

Glossary

This glossary defines terms used in this guide or related to this product and is not a comprehensive glossary of computer terms.

A

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

adapter
See controller.

AL_PA
Arbitrated Loop Physical Address. A 1-byte value the arbitrated loop topology uses to identify the loop ports. This value becomes the last byte of the address identifier for each public port on the loop.

allocation policy
Storage system rules that govern how virtual disks are created. Allocate Completely and Allocate on Demand are the two rules used in creating virtual disks.

- Allocate Completely—The space a virtual disk requires on the physical disks is reserved, even if the virtual disk is not currently using the space.
- Allocate on Demand—The space a virtual disk requires on the physical disks is not reserved until needed.

ambient temperature
The air temperature in the area where a system is installed. Also called intake temperature or room temperature.

ANSI
American National Standards Institute. A non-governmental organization that develops standards (such as SCSI I/O interface standards and Fibre Channel interface standards) used voluntarily by many manufacturers within the United States.

arbitrated loop
A Fibre Channel topology that links multiple ports (up to 126) together on a single shared simplex media. Transmissions can only occur between a single pair of nodes at any given time. Arbitration is the scheme that determines which node has control of the loop at any given moment

arbitrated loop physical address
See AL_PA.

arbitrated loop topology
See arbitrated loop.

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.

asynchronous
Events scheduled as the result of a signal requesting the event or that which is without any specified time relation.

audible alarm
The Environmental Monitoring Unit (EMU) alarm that sounds when there is a drive enclosure element condition report. The audible alarm can be muted or disabled.

B

backplane
An electronic printed circuit board that distributes data, control, power, and other signals to element connectors.

bad block
A data block that contains a physical defect.

bad block replacement
A replacement routine that substitutes defect-free disk blocks for those found to have defects. This process takes place in the controller and is transparent to the host.

bail lock
Part of the power supply AC receptacle that engages the AC power cord connector to ensure that the cord cannot be accidentally disconnected.

baud The maximum rate of signal state changes per second on a communication circuit. If each signal state change corresponds to a code bit, then the baud rate and the bit rate are the same. It is also possible for signal state changes to correspond to more than one code bit so the baud rate may be lower than the code bit rate.

bay The physical location of an element, such as a drive, I/O module, EMU or power supply in a drive enclosure. Each bay is numbered to define its location.

bidirectional *Also* called Bi-Di. The movement of optical signals in opposite directions through a common fiber cable such as the data flow path typically on a parallel printer port. A parallel port can provide two-way data flow for disk drives, scanning devices, FAX operations and even parallel modems.

block Also called a sector. The smallest collection of consecutive bytes addressable on a disk drive. In integrated storage elements, a block contains 512 bytes of data, error codes, flags, and the block address header.

blower A variable speed airflow device that pulls air into an enclosure or element. It usually pulls air in from the front and exhausts the heated air out the rear.

C

cabinet An alternate term used for a rack.

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 each end of the cable, the cable assembly is referred to as a jumper.

CAC Corrective Action Code. An HP P6000 Command View graphical user interface (GUI) display component that defines the action required to correct a problem.
See also read cache,write cache, and mirrored cache.

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.

cache battery A rechargeable unit mounted within a controller enclosure that supplies back-up power to the cache module in case of primary power shortage.

cache battery indicator

1. An orange light emitting diode (indicator) that illuminates on the controller operator control panel (OCP) to define the status of the HSV Controller cache batteries.
2. An amber status indicator that illuminates on a cache battery. When illuminated, it indicates that one or more cache battery cells have failed and the battery must be replaced with a new battery.

carrier A drive-enclosure-compatible assembly containing a disk drive or other storage devices.

client A software program that uses the services of another software program. The HP P6000 Command View client is a standard internet browser.

communication logical unit number (LUN) *See* console LUN.

condition report A three-element code generated by the EMU in the form where e.t. is the element type (a hexadecimal number), en. is the element number (a decimal number), and ec is the condition code (a decimal number).

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 enclosure A unit that holds one or more controllers, power supplies, blowers, cache batteries, transceivers, and connectors.

controller event A significant occurrence involving any storage system hardware or software component reported by the controller to HP P6000 Command View.

controller fault indicator An amber fault indicator that illuminates on the controller OCP to indicate when there is an HSV Controller fault.

controller pair Two interconnected controller modules which together control the disk enclosures in the storage system.

corrective action code *See* CAC.

CRITICAL Condition A drive enclosure EMU condition that occurs when one or more drive enclosure elements have failed or are operating outside of their specifications. The failure of the element makes continued normal operation of at least some elements in the enclosure impossible. Some enclosure elements may be able to continue normal operations. Only an UNRECOVERABLE condition has precedence. This condition has precedence over NONCRITICAL errors and INFORMATION condition.

CRU Customer Replaceable Unit. A storage system element that a user can replace without using special tools or techniques, or special training.

customer replaceable unit *See* CRU.

D

data entry mode The state in which controller information can be displayed or controller configuration data can be entered. On the Enterprise Storage System, the controller mode is active when the LCD on the HSV Controller OCP is Flashing.

default disk group The first disk group created at the time the system in initialized. The default disk group can contain the entire set of physical disks in the array or just a few of the disks.
See also disk group.

Detailed Fault View An HSV Controller OCP display that permits a user to view detailed information about a controller fault.

device channel A channel used to connect storage devices to a host I/O bus adapter or intelligent controller.

device ports Controller pair device ports connected to the storage system's physical disk drive array through the Fibre Channel drive enclosure. *Also* called a device-side port.

device-side ports *See* device ports.

DIMM Dual Inline Memory Module. A small circuit board holding memory chips.

dirty data The write-back cached data that has not been written to storage media even though the host operation processing the data has completed.

disk drive A carrier-mounted storage device supporting random access to fixed size blocks of data.

disk drive blank A carrier that replaces a disk drive to control airflow within a drive enclosure whenever there is less than a full complement of storage devices.

disk failure protection A method by which a controller pair reserves drive capacity to take over the functionality of a failed or failing physical disk. For each disk group, the controllers reserve space in the physical disk pool equivalent to the selected number of physical disk drives.

disk group A physical disk drive set or pool in which a virtual disk is created. A disk group may contain all the physical disk drives in a controller pair array or a subset of the array.

disk migration state A physical disk drive operating state. A physical disk drive can be in a stable or migration state:

- Stable—The state in which the physical disk drive has no failure nor is a failure predicted.
- Migration—The state in which the disk drive is failing, or failure is predicted to be imminent. Data is then moved off the disk onto other disk drives in the same disk group.

disk replacement delay The time that elapses between a drive failure and when the controller starts searching for spare disk space. Drive replacement seldom starts immediately in case the "failure" was a glitch or temporary condition.

drive blank *See* disk drive blank.

drive enclosure A unit that holds storage system devices such as disk drives, power supplies, blowers, I/O modules, transceivers, or EMUs.

drive enclosure *See* drive enclosure.

drive enclosure event A significant operational occurrence involving a hardware or software component in the drive enclosure. The drive enclosure EMU reports these events to the controller for processing.

dual power supply configuration *See* redundant power configuration.

dual-loop A configuration where each drive is connected to a pair of controllers through two loops. These two Fibre Channel loops constitute a loop pair.

dynamic capacity expansion A storage system feature that provides the ability to increase the size of an existing virtual disk. Before using this feature, you must ensure that your operating system supports capacity expansion of a virtual disk (or LUN).

E

EIA Electronic Industries Alliance. A standards organization specializing in the electrical and functional characteristics of interface equipment.

EIP Event Information Packet. The event information packet is an HSV element hexadecimal character display that defines how an event was detected. Also called the EIP type.

electromagnetic interference *See* EMI.

electrostatic discharge *See* ESD.

element

1. In a drive enclosure, a device such as an EMU, power supply, disk, blower, or I/O module. The object can be controlled, interrogated, or described by the enclosure services process.
2. In the Open SAN Manager, a controllable object, such as the Enterprise storage system.

EMI Electromagnetic Interference. The impairment of a signal by an electromagnetic disturbance.

EMU Environmental Monitoring Unit. An element which monitors the status of an enclosure, including the power, air temperature, and blower status. The EMU detects problems and displays and reports these conditions to a user and the controller. In some cases, the EMU implements corrective action.

enclosure A unit used to hold various storage system devices such as disk drives, controllers, power supplies, blowers, an EMU, I/O modules, or blowers.

enclosure address bus An Enterprise storage system bus that interconnects and identifies controller enclosures and disk drive enclosures by their physical location. Enclosures within a reporting group can exchange environmental data. This bus uses enclosure ID expansion cables to assign enclosure numbers to each enclosure. Communications over this bus do not involve the Fibre Channel drive enclosure bus and are, therefore, classified as out-of-band communications.

enclosure number (En) One of the vertical rack-mounting positions where the enclosure is located. The positions are numbered sequentially in decimal numbers starting from the bottom of the cabinet. Each disk enclosure has its own enclosure number. A controller pair shares an enclosure number. If the system has an expansion rack, the enclosures in the expansion rack are numbered from 15 to 24, starting at the bottom.

enclosure services Those services that establish the mechanical environmental, electrical environmental, and external indicators and controls for the proper operation and maintenance of devices with an enclosure as described in the *SES SCSI-3 Enclosure Services Command Set (SES), Rev 8b, American National Standard for Information Services*.

Enclosure Services Interface *See* ESI.

Enclosure Services Processor *See* ESP.

Enterprise Virtual Array The Enterprise Virtual Array is a product that consists of one or more storage systems. Each storage system consists of a pair of HSV controllers and the disk drives they manage. A storage system within the Enterprise Virtual Array can be formally referred to as an Enterprise storage system, or generically referred to as the storage system.

Enterprise Virtual Array rack A unit that holds controller enclosures, disk drive enclosures, power distribution supplies, and enclosure address buses that, combined, comprise an Enterprise storage system solution. *Also* called the Enterprise storage system rack.
See also rack.

environmental monitoring unit *See* EMU.

error code The portion of an EMU condition report that defines a problem.

ESD Electrostatic Discharge. The emission of a potentially harmful static electric voltage as a result of improper grounding.

ESI Enclosure Services Interface. The SCSI-3 engineering services interface implementation developed for Storage products. A bus that connects the EMU to the disk drives.

ESP Enclosure Services Processor. An EMU that implements an enclosure's services process.

event Any significant change in the state of the Enterprise storage system hardware or software component reported by the controller to HP P6000 Command View.
See also controller event, drive enclosure event, management agent event, and termination event..

Event Information Packet *See* EIP.

Event Number *See* Evt No.

Evt No. Event Number. A sequential number assigned to each Software Code Identification (SWCID) event. It is a decimal number in the range 0-255.

exabyte A unit of storage capacity that is the equivalent of 2^{60} bytes or 1,152,921,504,606,846,976 bytes. One exabyte is equivalent to 1,024 petabytes.

F

fabric A Fibre Channel fabric or two or more interconnected Fibre Channels 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.

fan The variable speed airflow device that cools an enclosure or element by forcing ambient air into an enclosure or element and forcing heated air out the other side.
See also blower.

Fault Management Code *See* FMC.

FC HBA Fibre Channel Host Bus Adapter. An interchangeable term for Fibre Channel adapter.
See also FCA.

FCA Fibre 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.

FCC Federal Communications Commission. The federal agency responsible for establishing standards and approving electronic devices within the United States.

FCP Fibre Channel Protocol. The mapping of SCSI-3 operations to Fibre Channel.

fiber The optical media used to implement Fibre Channel.

fiber optic cable A transmission medium designed to transmit digital signals in the form of pulses of light. Fiber optic cable is noted for its properties of electrical isolation and resistance to electrostatic contamination.

fiber optics The technology where light is transmitted through glass or plastic (optical) threads (fibers) for data communication or signaling purposes.

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.

Fibre Channel drive enclosure Fibre Channel Arbitrated Loop. The American National Standards Institute's (ANSI) document that specifies arbitrated loop topology operation.

Fibre Channel Loop An enclosure that provides twelve-port central interconnect for Fibre Channel Arbitrated Loops following the ANSI Fibre Channel drive enclosure standard.

field replaceable unit *See* FRU.

flush The act of writing dirty data from cache to a storage media.

FMC Fault Management Code. The HP P6000 Command View display of the Enterprise Storage System error condition information.

form factor A storage industry dimensional standard for 3.5inch (89 mm) and 5.25inch (133 mm) high storage devices. Device heights are specified as low-profile (1inch or 25.4 mm), half-height (1.6inch or 41 mm), and full-height (5.25inch or 133 mm).

FPGA Field Programmable Gate Array. A programmable device with an internal array of logic blocks surrounded by a ring of programmable I/O blocks connected together through a programmable interconnect.

frequency The number of cycles that occur in one second expressed in Hertz (Hz). Thus, 1 Hz is equivalent to one cycle per second.

FRU Field Replaceable Unit. A hardware element that can be replaced in the field. This type of replacement can require special training, tools, or techniques. Therefore, FRU procedures are usually performed only by an Authorized Service Representative.

H

HBA Host Bus Adapter.
See also 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 link indicator The HSV Controller display that indicates the status of the storage system Fibre Channel links.

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.

hot-pluggable A method of element replacement whereby the complete system remains operational during element removal or insertion. Replacement does not interrupt data transfers to other elements.

hub A communications infrastructure device to which nodes on a multi-point bus or loop are physically connected. It is used to improve the manageability of physical cables.

I

I/O module Input/Output module. The enclosure element that is the Fibre Channel drive enclosure interface to the host or controller. I/O modules are bus speed specific, either 1 Gb or 2 Gb.

IDX A 2-digit decimal number portion of the HSV controller termination code display that defines one of 32 locations in the Termination Code array that contains information about a specific event. *See also* param and TC.

in-band communication The method of communication between the EMU and controller that utilizes the Fibre Channel drive enclosure bus.

INFORMATION condition A drive enclosure EMU condition report that may require action. This condition is for information only and does not indicate the failure of an element. All condition reports have precedence over an INFORMATION condition.

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.

input/output module *See* I/O module.

intake temperature *See* ambient temperature.

interface A set of protocols used between components such as cables, connectors, and signal levels.

J

JBOD Just a Bunch of Disks. A number of disks connected to one or more controllers.

L

indicator Light Emitting Diode. A semiconductor diode, used in an electronic display, that emits light when a voltage is applied to it.

LAN Local area network. A group of computers and associated devices that share a common communications line and typically share the resources of a single processor or server within a small geographic area.

laser A device that amplifies light waves and concentrates them in a narrow, very intense beam.

Last Fault View An HSV Controller display defining the last reported fault condition.

Last Termination Error Array *See* LTEA.

License Key A WWN-encoded sequence that is obtained from the license key fulfillment website.

link A connection between ports on Fibre Channel devices. The link is a full duplex connection to a fabric or a simplex connection between loop devices.

logon *Also* called login, it is a procedure whereby a user or network connection is identified as being an authorized network user or participant.

loop *See* arbitrated loop.

loop ID Seven-bit values numbered contiguous from 0 to 126 decimal that represent the 127 valid AL_PA values on a loop (not all 256 hexadecimal values are allowed as AL_PA values per Fibre Channel).

loop pair A Fibre Channel attachment between a controller and physical disk drives. Physical disk drives connect to controllers through paired Fibre Channel arbitrated loops. There are two loop pairs, designated loop pair 1 and loop pair 2. Each loop pair consists of two loops (called loop A and loop B) that operate independently during normal operation, but provide mutual backup in case one loop fails.

LTEA Last Termination Event Array. A two-digit HSV Controller number that identifies a specific event that terminated an operation. Valid numbers range from 00 to 31.

LUN Logical Unit Number. A SCSI convention used to identify elements. The host sees a virtual disk as a LUN. The LUN address a user assigns to a virtual disk for a particular host will be the LUN at which that host will see the virtual disk.

M

management agent The HP P6000 Command View software that controls and monitors the Enterprise storage system. The software can exist on more than one management server in a fabric. Each installation is a management agent.

management agent event Significant occurrence to or within the management agent software, or an initialized storage cell controlled or monitored by the management agent.

mean time between failures *See* MTBF.

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.

micro meter *See* µm.

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.

MTBF Mean Time Between Failures. The average time from start of use to first failure in a large population of identical systems, components, or devices.

multi-mode fiber A fiber optic cable with a diameter large enough (50 microns or more) to allow multiple streams of light to travel different paths from the transmitter to the receiver. This transmission mode enables bidirectional transmissions.

N

Network Storage Controller *See* NSC.

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

non-OFC (Open Fibre Control) A laser transceiver whose lower-intensity output does not require special open Fibre Channel mechanisms for eye protection. The Enterprise storage system transceivers are non-OFC compatible.

NONCRITICAL Condition A drive enclosure EMU condition report that occurs when one or more elements inside the enclosure have failed or are operating outside of their specifications. The failure does not affect continued normal operation of the enclosure. All devices in the enclosure continue to operate according to their specifications. The ability of the devices to operate correctly may be reduced if additional failures occur. UNRECOVERABLE and CRITICAL errors have precedence over this condition. This condition has precedence over INFORMATION condition. *Early correction can prevent the loss of data.*

NSC Network Storage Controller. The HSV Controllers used by the Enterprise storage system.

NVRAM Nonvolatile Random Access Memory. Memory whose contents are not lost when a system is turned Off or if there is a power failure. This is achieved through the use of UPS batteries or implementation technology such as flash memory. NVRAM is commonly used to store important configuration parameters.

O

occupancy alarm level A percentage of the total disk group capacity in blocks. When the number of blocks in the disk group that contain user data reaches this level, an event code is generated. The alarm level is specified by the user.

OCP Operator Control Panel. The element that displays the controller's status using indicators and an LCD. Information selection and data entry is controlled by the OCP push-button.

online/nearonline An online drive is a normal, high-performance drive, while a near-online drive is a lower-performance drive.

operator control panel *See* OCP.

P

param That portion of the HSV controller termination code display that defines:

- The 2-character parameter identifier that is a decimal number in the 0 through 30 range.
- The 8-character parameter code that is a hexadecimal number.

See also IDX and TC.

password A security interlock where the purpose is to allow:

- A management agent to control only certain storage systems
- Only certain management agents to control a storage system

PDM Power Distribution Module. A thermal circuit breaker-equipped power strip that distributes power from a PDU to Enterprise Storage System elements.

PDU Power Distribution Unit. The rack device that distributes conditioned AC or DC power within a rack.

physical disk A disk drive mounted in a 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-colored A convention of applying the color of port or red wine to a CRU tab, lever, or handle to identify the unit as hot-pluggable.

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.

power distribution module *See* PDM.

power distribution unit *See* PDU.

power supply An element that develops DC voltages for operating the storage system elements from either an AC or DC source.

preferred address An AL_PA which a node port attempts to acquire during loop initialization.

preferred path A preference for which controller of the controller pair manages the virtual disk. This preference is set by the user 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.

protocol The conventions or rules for the format and timing of messages sent and received.

Q

quiesce The act of rendering bus activity inactive or dormant. For example, “quiesce the SCSI bus operations during a device warm-swap.”

R

rack A floorstanding structure primarily designed for, and capable of, holding and supporting storage system equipment. All racks provide for the mounting of panels per Electronic Industries Alliance (EIA) *Standard RS310C*.

rack-mounting unit A measurement for rack heights based upon a repeating hole pattern. It is expressed as “U” spacing or panel heights. Repeating hole patterns are spaced every 1.75 inches (44.45 mm) and based on EIA’s *Standard RS310C*. For example, a 3U unit is 5.25inches (133.35 mm) high, and a 4U unit is 7.0inches (177.79 mm) high.

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.

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.

reconstruction The process of regenerating the contents of a failed member data. The reconstruction process writes the data to a spare set disk and incorporates the spare set disk into the mirrorset, striped mirrorset or RAID set from which the failed member came.

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 works 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.

redundant power configuration A capability of the Enterprise storage system racks and enclosures to allow continuous system operation by preventing single points of power failure.

- For a rack, two AC power sources and two power conditioning units distribute primary and redundant AC power to enclosure power supplies.
- For a controller or drive enclosure, two power supplies ensure that the DC power is available even when there is a failure of one supply, one AC source, or one power conditioning unit. Implementing the redundant power configuration provides protection against the loss or corruption of data.

reporting group An Enterprise Storage System controller pair and the associated disk drive enclosures. The Enterprise Storage System controller assigns a unique decimal reporting group number to each EMU on its loops. Each EMU collects disk drive environmental information from its own sub-enclosure and broadcasts the data over the enclosure address bus to all members of the reporting group. Information from enclosures in other reporting groups is ignored.

room temperature *See* ambient temperature.

S

SCSI-3 The ANSI standard that defines the operation and function of Fibre Channel systems.

SCSI-3 Enclosure Services *See* SES.

selective presentation The process whereby a controller presents a virtual disk only to the host computer which is authorized access.

serial transmission A method of transmission in which each bit of information is sent sequentially on a single channel rather than simultaneously as in parallel transmission.

SES SCSI-3 Enclosures Services. Those services that establish the mechanical environment, electrical environment, and external indicators and controls for the proper operation and maintenance of devices within an enclosure.

snapclone A virtual disk that can be manipulated while the data is being copied. Only an Active member of a virtual disk family can be snapcloned.

The Snapclone, like a snapshot, reflects the contents of the source virtual disk at a particular point in time. Unlike the snapshot, the Snapclone is an actual clone of the source virtual disk and immediately becomes an independent Active member of its own virtual disk family.

snapshot A temporary virtual disk (Vdisk) 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. Up to seven snapshots 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, and virtual disk family.

SSN Storage System Name. An HP P6000 Command View-assigned, unique 20-character name that identifies a specific storage system.

storage carrier *See* carrier.

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 Name *See* SSN.

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

symbols and numbers

3U A unit of measurement representing three "U" spaces. "U" spacing is used to designate panel or enclosure heights. Three "U" spaces is equivalent to 5.25 inches (133 mm).
See also rack-mounting unit.

µm A symbol for micrometer; one millionth of a meter. For example, 50 µm is equivalent to 0.000050 m.

T

TC Termination Code. An Enterprise Storage System controller 8-character hexadecimal display that defines a problem causing controller operations to halt.
See also IDX and param.

Termination Code *See* TC.

termination event Occurrences that cause the storage system to cease operation.

terminator Interconnected elements that form the ends of the transmission lines in the enclosure address bus.

topology An interconnection scheme that allows multiple Fibre Channel ports to communicate. Point-to-point, arbitrated loop, and ed fabric are all Fibre Channel topologies.

transceiver The device that converts electrical signals to optical signals at the point where the fiber cables connect to the FC elements such as hubs, controllers, or adapters.

U

uninitialized system A state in which the storage system is not ready for use.
See also initialization.

UNRECOVERABLE Condition A drive enclosure EMU condition report that occurs when one or more elements inside the enclosure have *failed* and have disabled the enclosure. The enclosure may be incapable of recovering or bypassing the failure and will require repairs to correct the condition.

This is the highest level condition and has precedence over all other errors and requires *immediate corrective action*.

unwritten cached data *Also* called unflushed data.
See also dirty data.

UPS Uninterruptible Power Supply. A battery-operated power supply guaranteed to provide power to an electrical device in the event of an unexpected interruption to the primary power supply. Uninterruptible power supplies are usually rated by the amount of voltage supplied and the length of time the voltage is supplied.

V

virtual disk Variable disk capacity that is defined and managed by the array controller and presented to hosts as a disk. May be called Vdisk in the user interface.

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, and virtual disk snapshot.

virtual disk snapshot *See* snapshot.

Vraid0 A virtualization technique that provides no data protection. Data 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 where the 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.

W

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

D

E

W

X

Z