

HP 4400 Enterprise Virtual Array User Guide

Abstract

This document describes the HP 4400 Enterprise Virtual Array (EVA4400) and provides information about operating the EVA4400. It is intended for users who install, operate, and manage EVA4400 storage systems.



© Copyright 2008, 2013 Hewlett-Packard Development Company, L.P.

The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Warranty

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

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

Acknowledgements

Itanium® is a registered trademark of Intel Corporation or its subsidiaries in the United States and other countries.

Microsoft® and Windows® are U.S. registered trademarks of Microsoft Corporation.

Oracle® is a registered U.S. trademark of Oracle Corporation or its affiliates.



Contents

1 EVA4400 hardware.....	9
Physical layout of the storage system.....	9
M6412 disk enclosures.....	10
Enclosure layout.....	10
I/O modules.....	11
I/O module status indicators.....	11
Fiber optic Fibre Channel cables.....	12
Copper Fibre Channel cables.....	12
Fibre Channel disk drives.....	13
Disk drive status indicators.....	13
Disk drive blank.....	13
Controller enclosures.....	14
Management module.....	15
Controller status indicators.....	16
HSV300 controller status LEDs.....	17
Power supplies.....	17
Fan module.....	18
Battery module.....	18
HSV controller cabling.....	19
Storage system racks	19
Rack configurations.....	20
Power distribution—Modular PDUs.....	20
PDUs.....	22
PDU A.....	22
PDU B.....	23
PDMs.....	23
Rack AC power distribution.....	24
Moving and stabilizing a rack.....	25
2 EVA4400 operation.....	27
Best practices.....	27
Operating tips and information.....	27
Reserving adequate free space.....	27
Using FATA disk drives.....	27
Using solid state disk drives.....	27
EVA4400 port does not log in to the embedded switch.....	27
Emulex HBAs require unique zoning.....	27
Creating 16 TB or greater virtual disks in Windows 2008.....	27
Importing Windows dynamic disk volumes.....	28
Losing a path to a dynamic disk.....	28
Microsoft Windows 2003 MSCS cluster installation.....	28
Array I/O can halt after a power cycle with an HP-UX host.....	28
Reboot can be required after EVA power cycle to resume access to LUNs.....	28
Host port connection limit on B-series 3200 and 3800 switches.....	29
Failback preference setting for HSV controllers.....	30
Changing virtual disk failover/failback setting.....	31
Implicit LUN transition.....	32
Storage system shutdown and startup.....	32
Shutting down the storage system from HP P6000 Command View.....	32
Shutting down the storage system from the array controller.....	33
Starting the storage system.....	33
Connecting to the management module.....	34

Connecting through a public network.....	34
Connecting through a private network.....	35
Changing the default operating mode.....	36
Accessing the HP P6000 Control Panel through HP P6000 Command View.....	37
Saving storage system configuration data.....	37
Adding disk drives to the storage system.....	39
Handling fiber optic cables.....	39
3 Configuring application servers.....	41
Overview.....	41
Clustering.....	41
Multipathing.....	41
Installing Fibre Channel adapters.....	41
Testing connections to the EVA.....	42
Adding hosts.....	42
Creating and presenting virtual disks.....	42
Verifying virtual disk access from the host.....	43
Configuring virtual disks from the host.....	43
HP-UX.....	43
Scanning the bus.....	43
Creating volume groups on a virtual disk using vgcreate.....	44
IBM AIX.....	45
Accessing IBM AIX utilities.....	45
Adding hosts.....	45
Creating and presenting virtual disks.....	45
Verifying virtual disks from the host.....	45
Linux.....	46
HBA drivers.....	46
Verifying virtual disks from the host.....	46
OpenVMS.....	46
Updating the AlphaServer console code, Integrity Server console code, and Fibre Channel FCA firmware.....	46
Verifying the Fibre Channel adapter software installation.....	46
Console LUN ID and OS unit ID.....	46
Adding OpenVMS hosts.....	47
Scanning the bus.....	48
Configuring virtual disks from the OpenVMS host.....	49
Setting preferred paths.....	49
Oracle Solaris.....	49
Loading the operating system and software.....	49
Configuring FCAs with the Oracle SAN driver stack.....	50
Configuring Emulex FCAs with the lpfc driver.....	50
Configuring QLogic FCAs with the qla2300 driver.....	52
Fabric setup and zoning.....	53
Oracle StorEdge Traffic Manager (MPxIO)/Oracle Storage Multipathing.....	53
Configuring with Veritas Volume Manager.....	53
Configuring virtual disks from the host.....	55
Verifying virtual disks from the host.....	56
Labeling and partitioning the devices.....	57
VMware.....	58
Configuring the EVA4400 with VMware host servers.....	58
Configuring an ESX server	58
Loading the FCA NVRAM.....	58
Setting the multipathing policy.....	59
Specifying DiskMaxLUN.....	60

Verifying connectivity.....	60
Verifying virtual disks from the host.....	60
Verifying virtual disks from the host.....	60
HP EVA P6000 Software Plug-in for VMware VAAI.....	61
System prerequisites.....	61
Enabling vSphere Storage API for Array Integration (VAAI).....	61
Installing the VAAI Plug-in.....	62
Installation overview.....	62
Installing the HP EVA VAAI Plug-in using ESX host console utilities.....	63
Installing the HP VAAI Plug-in using vCLI/vMA.....	64
Installing the VAAI Plug-in using VUM.....	66
Uninstalling the VAAI Plug-in.....	68
Uninstalling VAAI Plug-in using the automated script (hpeva.pl).....	68
Uninstalling VAAI Plug-in using vCLI/vMA (vihostupdate).....	68
Uninstalling VAAI Plug-in using VMware native tools (esxupdate).....	68
4 Replacing array components.....	70
Customer self repair	70
Parts-only warranty service.....	70
Best practices for replacing hardware components.....	70
Component replacement videos.....	70
Verifying component failure.....	70
Identifying the spare part.....	70
Replaceable parts.....	71
Replacing the failed component.....	73
Replacement instructions.....	73
5 Single path implementation.....	75
Installation requirements.....	75
Recommended mitigations.....	75
Supported configurations.....	76
General configuration components.....	76
Connecting a single path HBA server to a switch in a fabric zone.....	76
HP-UX configuration.....	78
Requirements.....	78
HBA configuration.....	79
Risks.....	79
Limitations.....	79
Windows Server 2003 (32-bit), Windows Server 2008 (32-bit), Windows Server 2012 (32-bit) configurations.....	80
Requirements.....	80
HBA configuration.....	80
Risks.....	80
Limitations.....	80
Windows Server 2008 (64-bit) and Windows Server 2003 (64-bit) configurations.....	81
Requirements.....	81
HBA configuration.....	81
Risks.....	82
Limitations.....	82
Oracle Solaris configuration.....	82
Requirements.....	82
HBA configuration.....	83
Risks.....	83
Limitations.....	83
OpenVMS configuration.....	84
Requirements.....	84

HBA configuration.....	84
Risks.....	85
Limitations.....	85
Xen configuration.....	85
Requirements.....	85
HBA configuration.....	86
Risks.....	86
Limitations.....	86
Linux (32-bit) configuration.....	86
Requirements.....	87
HBA configuration.....	87
Risks.....	87
Limitations.....	87
Linux (Itanium) configuration.....	88
Requirements.....	88
HBA configuration.....	88
Risks.....	89
Limitations.....	89
IBM AIX configuration.....	89
Requirements.....	89
HBA configuration.....	90
Risks.....	90
Limitations.....	90
VMware configuration.....	91
Requirements.....	91
HBA configuration.....	91
Risks.....	91
Limitations.....	91
Mac OS configuration.....	92
Failure scenarios.....	92
HP-UX.....	92
Windows Servers.....	93
Oracle Solaris.....	93
OpenVMS.....	94
Linux.....	94
IBM AIX.....	95
VMware.....	95
Mac OS.....	96
6 Error messages.....	97
7 Support and other resources.....	106
Contacting HP.....	106
HP technical support.....	106
Subscription service.....	106
Documentation feedback.....	106
Related information.....	106
Documents.....	106
Websites.....	106
Typographic conventions	108
Customer self repair.....	108
Rack stability.....	109
A Regulatory compliance notices.....	110
Regulatory compliance identification numbers.....	110
Federal Communications Commission notice.....	110

FCC rating label.....	110
Class A equipment.....	110
Class B equipment.....	110
Declaration of Conformity for products marked with the FCC logo, United States only.....	111
Modification.....	111
Cables.....	111
Canadian notice (Avis Canadien).....	111
Class A equipment.....	111
Class B equipment.....	111
European Union notice.....	111
Japanese notices.....	112
Japanese VCCI-A notice.....	112
Japanese VCCI-B notice.....	112
Japanese VCCI marking.....	112
Japanese power cord statement.....	112
Korean notices.....	112
Class A equipment.....	112
Class B equipment.....	113
Taiwanese notices.....	113
BSMI Class A notice.....	113
Taiwan battery recycle statement.....	113
Turkish recycling notice.....	113
Vietnamese Information Technology and Communications compliance marking.....	113
Laser compliance notices.....	114
English laser notice.....	114
Dutch laser notice.....	114
French laser notice.....	114
German laser notice.....	115
Italian laser notice.....	115
Japanese laser notice.....	115
Spanish laser notice.....	116
Recycling notices.....	116
English recycling notice.....	116
Bulgarian recycling notice.....	116
Czech recycling notice.....	116
Danish recycling notice.....	116
Dutch recycling notice.....	117
Estonian recycling notice.....	117
Finnish recycling notice.....	117
French recycling notice.....	117
German recycling notice.....	117
Greek recycling notice.....	117
Hungarian recycling notice.....	118
Italian recycling notice.....	118
Latvian recycling notice.....	118
Lithuanian recycling notice.....	118
Polish recycling notice.....	118
Portuguese recycling notice.....	119
Romanian recycling notice.....	119
Slovak recycling notice.....	119
Spanish recycling notice.....	119
Swedish recycling notice.....	119
Battery replacement notices.....	120
Dutch battery notice.....	120
French battery notice.....	120

German battery notice.....	121
Italian battery notice.....	121
Japanese battery notice.....	122
Spanish battery notice.....	122
B Non-standard rack specifications.....	123
Internal component envelope.....	123
EIA310-D standards.....	123
EVA cabinet measures and tolerances.....	123
Weights, dimensions and component CG measurements.....	123
Airflow and recirculation.....	124
Component airflow requirements.....	124
Rack airflow requirements.....	124
Configuration standards.....	124
UPS selection.....	124
Shock and vibration specifications.....	125
Glossary.....	126
Index.....	137

1 EVA4400 hardware


The EVA4400 contains the following hardware components:

- EVA controller enclosure—Contains power supplies, cache batteries, fans, and HSV controllers.
- Fibre Channel disk enclosure—Contains disk drives, power supplies, fans, midplane, and I/O modules.
- Fibre Channel Arbitrated Loop cables—Provide connectivity to the EVA controller enclosure and the Fibre Channel disk enclosures.
- Rack—Several free standing racks are available.

Physical layout of the storage system

The basic physical components are shown in [Figure 1 \(page 9\)](#). The disk drives are installed in the disk enclosures.

Figure 1 Storage system hardware components (back view)



1. Controller enclosure (showing HSV300)
2. Disk enclosures
3. Power distribution unit (at rear of enclosure)

M6412 disk enclosures

The M6412 disk enclosure contains the disk drives used for data storage; a storage system contains multiple disk enclosures. The major components of the enclosure are:


- 12-bay enclosure
- Dual-loop, Fibre Channel disk enclosure I/O modules
- Copper Fibre Channel cables
- Fibre Channel disk drives and drive blanks
- Power supplies
- Fan modules

NOTE: An EVA4400 requires a minimum of one disk shelf with eight disk drives.

Enclosure layout


The disk drives mount in **bays** in the front of the enclosure. The bays are numbered sequentially from top to bottom and left to right. A drive is referred to by its bay number (see [Figure 2 \(page 10\)](#)). Enclosure status indicators are located at the right of each disk. [Figure 3 \(page 10\)](#) shows the front and [Figure 4 \(page 11\)](#) shows the rear view of the disk enclosure.

Figure 2 Disk drive bay numbering



15815


Figure 3 Disk enclosure (front view without bezel ears)



16314

- | | |
|-----------------------------|--------------------------|
| 1. Rack-mounting thumbscrew | 4. UID push button |
| 2. Disk drive release | 5. Enclosure status LEDs |
| 3. Drive LEDs | |

Figure 4 Disk enclosure (rear view)




- | | | |
|---|-------------------------|--------------------------|
| 1. Power supply 1 | 5. Fan 1 status LED | 9. Enclosure status LEDs |
| 2. Power supply 1 status LED | 6. I/O module A | 10. Fan 2 |
| 3. Fan 1 | 7. I/O module B | 11. Power push button |
| 4. Enclosure product number and serial number | 8. Rear UID push button | 12. Power supply 2 |

I/O modules

Two I/O modules provide the interface between the disk enclosure and the host controllers, see [Figure 5 \(page 11\)](#). For redundancy, only dual-controller, dual-loop operation is supported. Each controller is connected to both I/O modules in the disk enclosure.

Figure 5 I/O module detail



- | | |
|---|----------------------------------|
| 1. Double 7-segment display: enclosure ID | 4. Manufacturing diagnostic port |
| 2. 4-Gb I/O ports | 5. I/O module status LEDs |
| 3. Port 1 (P1), Port 2 (P2) status LEDs | |

Each I/O module has two ports that can transmit and receive data for bidirectional operation. Activating a port requires connecting a Fibre Channel cable to the port. The port function depends upon the loop.



I/O module status indicators

There are five status indicators on the I/O module. See [Figure 5 \(page 11\)](#). The status indicator states for an operational I/O module are shown in [Table 1 \(page 11\)](#). [Table 2 \(page 12\)](#) shows the status indicator states for a non-operational I/O module.

Table 1 Port status LEDs

Status LED	Description
Green (left)	<ul style="list-style-type: none"> Solid green—Active link. Flashing green—Locate: remotely asserted by application client.
Amber (right)	<ul style="list-style-type: none"> Solid amber—Module fault, no synchronization. Flashing amber—Module fault.

Table 2 I/O module status LEDs

Status LED	Description
UID	<ul style="list-style-type: none"> Locate. Flashing blue—Remotely asserted by application client.
	<ul style="list-style-type: none"> Module health indicator: Flashing green—I/O module powering up. Solid green—Normal operation. Green off—Firmware malfunction.
	<ul style="list-style-type: none"> Fault indicator: Flashing amber—Warning condition (not visible when solid amber showing). Solid amber—Replace FRU. Amber off—Normal operation.

Fiber optic Fibre Channel cables

The Enterprise Virtual Array uses orange, 50- μ m, multi-mode, fiber optic cables for connection to the SAN or the host, where there is a direct connection to the host. The fiber optic cable assembly consists of two 2-m fiber optic strands and small form-factor connectors on each end. See [Figure 6 \(page 12\)](#).

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 inch)
- 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” \(page 39\)](#).

Figure 6 Fiber Optic Fibre Channel cable




Copper Fibre Channel cables

The Enterprise Virtual Array uses copper Fibre Channel cables to interconnect disk shelves. The cables are available in 0.6 meter (1.97 ft.) and 2.0 meter (6.56 ft.) lengths. Copper cables provide

performance comparable to fiber optic cables. Copper cable connectors differ from fiber optic small form-factor connectors (see [Figure 7 \(page 13\)](#)).

Figure 7 Copper Fibre Channel cable



Fibre Channel disk drives

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


- Dual-ported 4-Gb/s Fibre Channel controller interface that allows up to 96 disk drives to be supported per array controller enclosure
- Compact, direct-connect design for maximum storage density and increased reliability and signal integrity
- Both online high-performance disk drives and FATA disk drives supported in a variety of capacities and spindle speeds
- Better vibration damping for improved performance

Up to 12 disk drives can be installed in a disk enclosure.

Disk drive status indicators

Two status indicators display drive operational status. [Figure 8 \(page 13\)](#) identifies the disk drive status indicators and [Table 3 \(page 13\)](#) describes them.

Figure 8 Disk status indicators



1. Bi-color (amber/blue)
2. Green

Table 3 Disk status indicator LED descriptions

Drive LED	Description
Bi-color (top)	<ul style="list-style-type: none">• Slow flashing blue (0.5 Hz)—Used to locate drive.• Solid amber—Drive fault.
Green (bottom)	<ul style="list-style-type: none">• Flashing—Drive is spinning up or down and is not ready.• Solid—Drive is ready to perform I/O operations.• Flickering—Indicates drive activity.

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.

Controller enclosures

The EVA4400 contains either the HSV300 or HSV300-S controller enclosure. 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. A single enclosure contains two controllers. The HSV300 has 4 Gb host port capability; the HSV300-S has 8 Gb host port capability with the Brocade 5410 switch embedded within the controller. Other than the embedded switch, all controller enclosure components are the same for the HSV300 and HSV300-S.


Figure 9 (page 14) and Figure 10 (page 14) show the bezel and front view of the HSV300 controller enclosure respectively. Figure 11 (page 15) shows the back view of the HSV300 controller enclosure. Figure 12 (page 15) shows the back view of the HSV300-S controller enclosure.

Figure 9 Controller enclosure (front bezel)




- | | |
|--------------------------|--------------------------|
| 1. Battery 1 status LEDs | 4. Battery 2 status LEDs |
| 2. Fan 1 status LEDs | 5. Enclosure status LEDs |
| 3. Fan 2 status LEDs | 6. Front UID push button |

Figure 10 Controller enclosure (front view with bezel removed)




- | | |
|---|-------------------------------|
| 1. Rack-mounting thumbscrew | 8. Fan 1 normal operation LED |
| 2. Enclosure product number and serial number | 9. Fan 1 fault LED |
| 3. World Wide Number (WWN) | 10. Fan 2 |
| 4. Battery 1 | 11. Battery 2 |
| 5. Battery normal operation LED | 12. Enclosure status LEDs |
| 6. Battery fault LED | 13. Front UID push button |
| 7. Fan 1 | |

Figure 11 HSV300 controller enclosure (back view)



- | | |
|----------------------------------|---|
| 1. Power supply 1 | 9. Enclosure power push button |
| 2. HSV300 controller 1 | 10. Power supply 2 |
| 3. Management module status LEDs | 11. Host ports, FP1, FP2, connection to front end (host or SAN) |
| 4. Ethernet port | 12. DP1-A port, back-end connection to A loop |
| 5. Management module | 13. DP1-B port, back-end connection to B loop |
| 6. HSV300 controller 2 | 14. Manufacturing diagnostic port |
| 7. Rear UID push button | 15. HSV300 controller status and fault LEDs |
| 8. Enclosure status LEDs | |

Figure 12 HSV300-S controller enclosure (back view)



- | | |
|----------------------------------|--|
| 1. Power supply 1 | 10. Power supply 2 |
| 2. HSV300-S controller 1 | 11. Switch ports, 1, 2, ... up to 10 for connection to front end (switch or SAN) |
| 3. Management module status LEDs | 12. DPI-A port, back-end connection to A loop |
| 4. Ethernet port | 13. DPI-B port, back-end connection to B loop |
| 5. Management module | 14. Console port (switch management), upper connection |
| 6. HSV300-S controller 2 | 15. Ethernet port (switch management) |
| 7. Rear UID push button | 16. Manufacturing diagnostic port |
| 8. Enclosure status LEDs | 17. HSV300-S controller status and fault LEDs |
| 9. Enclosure power push button | |

Management module

The HP P6000 Control Panel provides a direct interface to the management module within each controller. From the HP P6000 Control Panel you can display storage system status and configuration information, shut down the storage system, and manage the password. For tasks to perform with the HP P6000 Control Panel, see the HP P6000 Control Panel online help.

The HP P6000 Control Panel provides two levels of administrator access and an interface for firmware updates to the management module. For additional details about the HP P6000 Control Panel, see the HP P6000 Control Panel online help.

Table 4 (page 16) describes the port LED indicators for the management module Ethernet port (callouts 3 and 4 in Figure 11 (page 15) and Figure 12 (page 15)).

Table 4 Management module Ethernet port LED indicators

LED color	Location	LED function	LED state	Status
Green	Left	Link state indicator	Off	No link detected.
			Solid green	Link detected.
Amber	Right	Link activity indicator	Off	No activity.
			Blinking amber	Normal activity.

Controller status indicators

The status indicators display the operational status of the controller. The function of each indicator is described in Table 8 (page 17). During initial setup, the status indicators might not be fully operational.

Each port on the rear of the controller has an associated status indicator located directly above it. Table 5 (page 16) lists the port and its status description for the HSV300. Table 6 (page 16) lists the port and its status descriptions for the HSV300-S.

Table 5 HSV300 controller port status indicators

Port	Description
Fibre Channel host ports	<ul style="list-style-type: none"> Green—Normal operation Amber—No signal detected Off—No SFP¹ detected or the Direct Connect HP P6000 Control Panel setting is incorrect
Fibre Channel device ports	<ul style="list-style-type: none"> Green—Normal operation Amber—No signal detected or the controller has failed the port Off—No SFP¹ detected Flashing green—Link recovery in progress (the link has been detected as down or failed and attempts to re-initialize the link are in progress)

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

Table 6 HSV300-S controller port status indicators

Port	Description
Fibre Channel switch ports	<ul style="list-style-type: none"> Green on—Normal operation or loopback port Green flashing—Normal online I/O activity Amber on—Faulted port, disabled due to diagnostics or <code>Portdisable</code> command Amber flashing—Port with no synchronization, receiving light but not yet online or segmented port Off—No SFP¹, no cable, no license detected
Fibre Channel device ports	<ul style="list-style-type: none"> Green—Normal operation Amber—No signal detected or the controller has failed the port Off—No SFP¹ detected

The HSV300-S controller includes an Ethernet port for the embedded switch (callout 15 in Figure 12 (page 15)). Table 7 (page 17) describes the LED functions of the Ethernet port.

Table 7 Embedded switch management Ethernet port LED indicators

LED color	Location	LED function	LED state	Status
Green	Right	Port speed indicator	Off	Port speed is 10 Mb/s and 100 Mb/s.
			Solid green	No link detected.
Amber	Left	Link state or activity indicator	Solid amber	No link detected.
			Blinking amber	Link detected.

HSV300 controller status LEDs

Figure 13 (page 17) shows the location of the controller status LEDs; Table 8 (page 17) describes them.

Figure 13 Controller status LEDs








Table 8 Controller status LEDs

Item	LED	Indication
1	UID	Blue LED used to identify a certain controller within the enclosure.
2		Controller health OK. Solid green LED after boot.
3		Blinking amber LED used to identify a controller failure either from a voltage issue or driven by firmware.
4	N/A	Not used.
5	N/A	Cache status. Slow flashing green LED shows cache is being powered by standby power.
6		Cache battery status. Slow flashing amber when cache is being powered by battery (during AC power loss). Faster flashing amber when testing. Solid amber when the controller determines a battery fault.

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 14 Power supply




- 1. Power supply
- 2. AC input connector
- 3. Latch
- 4. Status indicator (green—Normal operation; amber—Failure or no power)
- 5. Handle

Fan module

Fan modules provide the cooling necessary to maintain the proper operating temperature within the controller enclosure. If one fan fails, the remaining fan is capable of cooling the enclosure.

Figure 15 Fan module pulled out



- 1. Green—Fan normal operation LED
- 2. Amber—Fan fault LED


Table 9 Fan status indicators

Status indicator	Fault indicator	Description
Green	Solid green	Normal operation.
	Blinking	Maintenance in progress.
	Off	Amber is on or blinking, or the enclosure is powered down.
Amber	On	Fan failure. Green will be off. (Green and amber are not on simultaneously except for a few seconds after power-up.)

Battery module

Battery modules provide power to the controllers in the enclosure.

Figure 16 Battery module pulled out



1. Green—Normal operation LED

2. Amber—Fault LED

Each battery module provides power to the controller directly across from it in the enclosure.

Table 10 Battery status indicators

Status indicator	Fault indicator	Description
Green	Solid green	Normal operation.
	Blinking	Maintenance in progress.
	Off	Amber is on or blinking, or the enclosure is powered down.
Amber	Solid amber	Battery failure; no cache hold-up. Green will be off.
	Blinking amber	Battery degraded; replace soon. Green will be off. (Green and amber are not on simultaneously except for a few seconds after power-up.)

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 5 \(page 16\)](#) 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 enclosures or loops (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 disk enclosures to the loops in the main cabinet. Copper cables are used the controllers (mirror port) and the controllers and the disk enclosures or loops.

Storage system racks

All storage system components are mounted in a rack. Each configuration includes one controller enclosure holding both controllers (the controller pair), FC cables the controller and the disk enclosures. Each controller pair and all associated disk enclosures form a single storage system.

The rack provides the capability for mounting standard 483 mm (19 inch) wide controller and disk enclosures.

NOTE: Racks and rack-mountable components are typically described using *U* measurements. *U* measurements are used to designate panel or enclosure heights. The *U* measurement is a standard of 41 mm (1.6 inches).

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 contains several disk enclosures included in the storage system. See “Storage system hardware components (back view)” (page 9) for a typical EVA4400 rack configuration. The standard rack for the EVA4400 is the 42U HP 10000 Intelligent Series rack. The EVA4400 is also supported with 22U, 36U, 42U 5642, and 47U racks. The 42U 5642 is a field-installed option and the 47U rack must be assembled onsite because the cabinet height creates shipping difficulties.

For more information on HP rack offerings for the EVA4400, see <http://h18004.www1.hp.com/products/servers/proliantstorage/racks/index.html>.

Power distribution—Modular PDUs

NOTE: This section describes the most common power distribution system for EVA4400s. For information about other options, see the HP power distribution units website:

<http://h18004.www1.hp.com/products/servers/proliantstorage/power-protection/pdu.html>

AC power is distributed to the rack through a dual 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 (100–127 VAC or 220–240 VAC $\pm 10\%$, 50 or 60-Hz, $\pm 5\%$). The following figures illustrate the most common compatible 60-Hz and 50-Hz wall receptacles. For more information about PDU support, see the *HP 4400 Enterprise Virtual Array QuickSpecs*.

NOTE:

This section describes 30-A, single phase power. You can order other voltage, amperage, and phase configurations if you have a different power infrastructure.



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



NEMA L5-30R receptacle, 3-wire, 30-A, 60-Hz



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

- The standard power configuration for any HP 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.
 - One dual PDU assembly. Each PDU connects to a different wall receptacle.
 - Four to eight (depending on the rack) Power Distribution Modules (PDMs) per rack. PDMs are split evenly on both sides of the rack. Each set of PDMs connects to a different PDU.
 - Eight PDMs for 42U, 47U, and 42U 5642 racks
 - Six PDMs for 36U racks
 - Four PDMs for 22U racks
 - The disk enclosure power supplies on the left (disk PS 1) connect to the PDMs on the left with a gray, 66 cm (26 inch) power cord.
 - The disk enclosure power supplies on the right (disk PS 2) connect to the PDMs on the right with a black, 66 cm (26 inch) power cord.
 - Each controller enclosure has two power supplies:
 - Controller PS 1 connects to a PDM on the left with a gray, 152 cm (60 inch) power cord.
 - Controller PS 2 connects to a PDM on the right with a black, 66 cm (26 inch) power cord.

NOTE: Disk 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.

PDU


Each Enterprise Virtual Array rack has either a 50- or 60-Hz, dual PDU mounted at the bottom rear of the rack. The PDU placement is back-to-back, plugs facing toward the front (Figure 17 (page 22)), with circuit breaker switches facing the back (Figure 18 (page 22)).

- 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. One option is the NEMA L5-30R receptacle, 3-wire, 30-A, 60-Hz 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. The voltages are then routed to PDMs, sometimes called AC power strips, mounted on the two vertical rails in the rear of the rack.

Figure 17 Dual PDU—front view



- | | |
|-------------------|-------------------------------|
| 1. PDU B | 4. Power receptacle schematic |
| 2. PDU A | 5. Power cord |
| 3. AC receptacles | |

Figure 18 Dual PDU—rear view



- | | |
|----------|-------------------------|
| 1. PDU B | 3. Main circuit breaker |
| 2. PDU A | 4. Circuit breakers |

PDU A

PDU A connects to AC PDM A1–A4.

A PDU A failure:

- Disables the power distribution circuit
- Removes power from the left side of the rack
- Disables disk enclosure PS 1
- Disables controller PS 1

PDU B

PDU B connects to AC PDM B1–B4.

A PDU B failure:

- Disables the power distribution circuit
- Removes power from the right side of the rack
- Disables disk enclosure PS 2
- Disables controller PS 2

PDMs

Depending on the rack, there can be up to eight PDMs mounted in the rear of the rack:


- The PDMs on the left vertical rail connect to PDU A
- The PDMs on the right vertical rail connect to PDU B

Each PDM has seven AC receptacles. The PDMs distribute the AC power from the PDUs to the enclosures. Two power sources exist for each controller pair and disk 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 19 Rack PDM



1. Power receptacles
2. 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 20 \(page 24\)](#)).

- PDMs A1 through A4 connect to receptacles A through D on PDU A. Power cords connect these PDMs to the left power supplies on the disk enclosures (disk PS 1) and to the left power supply on the controller enclosure (controller PS 1).
- PDMs B1 through B4 connect to receptacles A through D on PDU B. Power cords connect these PDMs to the right power supplies on the disk enclosures (disk PS 2) and to the right power supply on the controller enclosure (controller PS 2).

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

Figure 20 Rack AC power distribution



1. PDM A1
2. PDM A2
3. PDM A3
4. PDM A4
5. PDU B

6. PDM B1
7. PDM B2
8. PDM B3
9. PDM B4
10. PDU A

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 inches) wide for the 60.3 cm (23.7 inch) wide, 42U rack. A vertical clearance of 203.2 cm (80 inches) should ensure sufficient clearance for the 200 cm (78.7 inch) 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 inch) wide and 1000 mm (39.37 inch) deep (see [Figure 21](#) (page 25)).

Figure 21 Single rack configuration floor space requirements




- | | |
|------------------------------|------------------------------------|
| 1. Front door | 5. Rear service area depth 300 mm |
| 2. Rear door | 6. Rack depth 1,000 mm |
| 3. Rack width 600 mm | 7. Front service area depth 406 mm |
| 4. Service area width 813 mm | 8. Total rack depth 1,706 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 22](#) (page 26)).
2. Repeat [Step 1](#) for the other feet.

Figure 22 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 21 \(page 25\)](#)).
- 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.

2 EVA4400 operation

Best practices

For useful information on managing and configuring your storage system, see the *HP 4400/6400/8400 Enterprise Virtual Array configuration Best practices white paper* available at: <http://h18006.www1.hp.com/storage/arraywhitepapers.html>

Operating tips and information

Reserving adequate free space

To ensure efficient storage system operation, reserve some unallocated capacity, or free space, 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 4400/6400/8400 Enterprise Virtual Array configuration Best practices white paper*. See “Best practices” (page 27).

Using FATA disk drives

FATA drives are designed for lower duty cycle applications such as near online data replication for backup. Do not use these drives as a replacement for EVA's high performance, standard duty cycle, Fibre Channel drives. This practice could shorten the life of the drive. For more information on FATA drives, see the *HP 4400/6400/8400 Enterprise Virtual Array configuration Best practices white paper*. See “Best practices” (page 27).

Using solid state disk drives

- Supported in the EVA4400 and EVA6400/8400 only, running a minimum controller software version of 09500000 for the 72 GB drive and 09534000 for the 200 GB and 400 GB drives
- SSD drives must be in a separate disk group
- The SSD disk group supports a minimum of 6 and a maximum of 8 drives per array
- SSD drives can only be configured with Vraid5 or Vraid1 (Vraid1 requires controller software version 09534000)
- Supported with HP P6000 Business Copy
- Not supported with HP P6000 Continuous Access
- Dynamic Capacity Management extend and shrink features are not supported

Use of these devices in unsupported configurations can lead to unpredictable results, including unstable array operation or data loss.

EVA4400 port does not log in to the embedded switch

It has been observed that an EVA4400 host port will occasionally not log in to the embedded switch. If you encounter this issue, restart the controller to bring the port online.

Emulex HBAs require unique zoning

When multiple Emulex HBAs are connected to the same EVA4400, each HBA requires unique zoning.

Creating 16 TB or greater virtual disks in Windows 2008

When creating a virtual disk that is 16 TB or greater in Windows 2008, ensure that the Allocation unit size field is set to something other than Default in the Windows New Simple Volume wizard.

The recommended setting is 16K. If this field is set to Default, you will receive the following error message:

The format operation did not complete because the cluster count is higher than expected.

Importing Windows dynamic disk volumes

If you create a snapshot, snapclone, or mirrorclone with a Windows 2003 RAID-spanned dynamic volume on the source virtual disk, and then try to import the copy to a Windows 2003 x64 (64-bit) system, it will import with Dynamic Foreign status. The following message displays in the DiskPart utility:

The disk management services could not complete the operation.

This error occurs because the 64-bit version of DiskPart fails to import dynamic RAID sets on a new server.

To avoid this issue, use the 32-bit version of DiskPart instead of the 64-bit version. Copy DiskPart from a 32-bit x86 Windows system, located in `C:\WINDOWS\system32`. Place the DiskPart utility in a temporary folder on the 64-bit x64 Windows system.

Losing a path to a dynamic disk

If you are using Windows 2003 with dynamic disks and a path to the EVA virtual disk is temporarily lost, the Logical Disk Manager (LDM) will erroneously show a failed dynamic volume. For more information, see the following issue on the Microsoft knowledge base website:

<http://support.microsoft.com/kb/816307>

To resolve the issue, reboot the Windows 2003 server to restore the dynamic volume.

Microsoft Windows 2003 MSCS cluster installation

The MSCS cluster installation wizard on Windows 2003 can fail to find the shared quorum device and disk resources might 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:

- Follow the workaround recommendation described in the 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)*, available for download on the Microsoft website:

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

- Use the MPIO DSM CLI to set the load balancing policy for each LUN to NLB.

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

Array I/O can halt after a power cycle with an HP-UX host

I/O halts after power cycling an HSV300-S enclosure connected to an HP-UX 11.23 host with HP-UX driver versions earlier than B.11.23.08.02. To resolve this issue, either upgrade to version B.11.23.08.02, or reset the HBAs on any server affected after an HSV300-S is power cycled.

Reboot can be required after EVA power cycle to resume access to LUNs

If an EVA4400 experiences a power cycle while the management GUI for Windows MPIO DSM is open, the LUNs might not appear under Disk Management after they recover. This issue is seen when the MPIO DSM Manager snap-in is open at the time of a power cycle. To reduce the probability of this occurring, HP recommends that you close the GUI when you are not performing management actions.

If this problem occurs, follow these steps:

1. Close the HP MPIO DSM Manager GUI.
2. Close Disk Management.
3. Stop and restart the Virtual Disk services.
4. Open Disk Management, and then rescan or Diskpart rescan to enumerate the LUNs.

If these steps are not successful, reboot the server.

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

A maximum of three EVA4400 host ports are supported on a single B-Series 3200 or 3800 switch running version 3.2.x. HP recommends that you connect only one storage host port on a B-Series switch. However, if you must connect more than one storage host port to a switch, separate the connection into different quadrants. Otherwise the connections can drop following an array controller resynchronization or when an event impacts the fabric, such as rebooting or adding a switch.

To 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, re-establish the connection as follows:
 1. Disconnect the Fibre Channel cable from the failed port.
 2. Wait 10 seconds, and then reconnect the cable.

The port relogs into the fabric and re-establishes a connection to the array.

Failback preference setting for HSV controllers

Table 11 (page 30) describes the failback preference mode for the controllers.

Table 11 Failback preference settings

Setting	Point in time	Behavior
No Preference	At initial presentation	The units are alternately brought online to Controller 1 or to Controller 2.
	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 1 or to Controller 2.
	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 1.
	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 1.
	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 2.
	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 2.
	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 1.
	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 1.
	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 2 and set to Path A are brought online to Controller 1. 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 2.
	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 2.
	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 1 and set to Path B are brought online to Controller 2. This is a one-time occurrence. If the host then moves the LUN using SCSI commands, the LUN will remain where moved.

Table 12 (page 31) describes the failback default behavior and supported settings when ALUA-compliant multipath software is running with each operating system. Recommended settings may vary depending on your configuration or environment.

Table 12 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 13 \(page 31\)](#) identifies the presentation behavior that results when the failover/failback setting for a virtual disk is changed.

NOTE: If the new setting moves the presentation of the virtual disk to a new controller, any snapshots or snapclones associated with the virtual disk are also moved.

Table 13 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 2, it is moved to Controller 1. If the disk is on Controller 1, it remains there.
Path B Failover	If the disk is currently presented on Controller 1, it is moved to Controller 2. If the disk is on Controller 2, it remains there.
Path A Failover/Failback	If the disk is currently presented on Controller 2, it is moved to Controller 1. If the disk is on Controller 1, it remains there.
Path B Failover/Failback	If the disk is currently presented on Controller 1, it is moved to Controller 2. If the disk is on Controller 2, it remains there.

Implicit LUN transition

Implicit LUN transition automatically transfers management of a virtual disk to the array controller that receives the most read requests for that virtual disk. This improves performance by reducing the overhead incurred when servicing read I/Os on the non-managing controller. Implicit LUN transition is enabled in XCS.

When creating a virtual disk, one controller is selected to manage the virtual disk. Only this managing controller can issue I/Os to a virtual disk in response to a host read or write request. If a read I/O request arrives on the non-managing controller, the read request must be transferred to the managing controller for servicing. The managing controller issues the I/O request, caches the read data, and mirrors that data to the cache on the non-managing controller, which then transfers the read data to the host. Because this type of transaction, called a proxy read, requires additional overhead, it provides less than optimal performance. (There is little impact on a write request because all writes are mirrored in both controllers' caches for fault protection.)

With implicit LUN transition, when the array detects that a majority of read requests for a virtual disk are proxy reads, the array transitions management of the virtual disk to the non-managing controller. This improves performance because the controller receiving most of the read requests becomes the managing controller, reducing proxy read overhead for subsequent I/Os.

Implicit LUN transition is disabled for all members of an HP Continuous Access EVA DR group. Because HP Continuous Access EVA requires that all members of a DR group be managed by the same controller, it would be necessary to move all members of the DR group if excessive proxy reads were detected on any virtual disk in the group. This would impact performance and create a proxy read situation for the other virtual disks in the DR group. Not implementing implicit LUN transition on a DR group may cause a virtual disk in the DR group to have excessive proxy reads.

Storage system shutdown and startup

You can shut down the array from HP P6000 Command View or from the array controller.

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 disk enclosures
5. Disconnects the system from HP P6000 Command View

NOTE: The storage system may take several minutes (up to 15) 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 from HP P6000 Command View

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.

Shutting down the storage system from the array controller

1. Push and hold the enclosure power button on the rear of the EVA4400 (see callout 9 in [Figure 11 \(page 15\)](#) or [Figure 12 \(page 15\)](#)).
 2. Wait 4 seconds.
The power button and the green LED start to blink.
-
- NOTE:** You must press and hold the power switch button for a minimum of 4 seconds. Otherwise, nothing will happen.
-
3. Between 4 and 10 seconds, release the button. The button continues to blink and the firmware starts the shutdown process (LED becomes solid amber, flushes the cache, disables the batteries).
 4. After 10 seconds, the power shuts down.

Starting the storage system

1. Verify that each Fibre Channel fabric switch connected to the HSV controllers 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 can take several minutes.
2. Power on the circuit breakers on both EVA rack PDUs. Verify that all disk enclosures are operating properly. The status indicator and the power indicator should be on (green).
3. Power on the controller enclosure, which powers on both controllers.
4. After allowing the EVA4400 to start, connect to the HP P6000 Control Panel to verify the display of the storage system name and WWN.
5. 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 allow HP P6000 Command View to discover the array.

NOTE: If the storage system is still not visible, check zoning and cabling, and verify correct configuration of server connectivity. If necessary, reboot the management server to re-establish the communication link.

6. 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.

With XCS 10000000 or later and HP P6000 Command View 9.4 or later, there is a feature in the HP P6000 Control Panel that enables the controllers to boot automatically when power is applied after a full shutdown. See the HP P6000 Control Panel online help or user guide for details about setting this feature. To further clarify the use of this feature:

- If this feature is disabled and you turn on power to the array from the rack power distribution unit (PDU), only the disk enclosures boot up. With this feature enabled, the controllers will also boot up, making the entire array ready for use.
- If, after setting this feature, you remove the management module from its slot and reinsert it to reset power or you restart the management module from the HP P6000 Control Panel, only the controllers will automatically boot up after a full shutdown. In this scenario, you must ensure

that the disk enclosures are powered up first; otherwise, the controller boot up process may be interrupted.

- After setting this HP P6000 Control Panel feature, if you have to shut down the array, perform the following steps:
 1. Use HP P6000 Command View to shut down the controllers and disk enclosures.
 2. Turn off power from the rack power distribution unit (PDU).
 3. Turn on power from the rack PDU.

After startup of the management module, the controllers will automatically start.

Connecting to the management module


You can connect to the management module through a public or a private network.

- ❗ **IMPORTANT:** To aggregate additional HP Command View servers into a SPoG (Single Pane of Glass) when the servers reside on a different IP network from the SPoG web server, some additional configuration settings may be needed. To configure the system if the HP Command View servers reside on a different network and you would like them visible in the SPoG, see Discovery URI under Discovery configuration settings in the *HP Management Integration Framework Administrator Guide*.

If you use a laptop to connect to the management module, configure the laptop to have an address in the same IP range as the management module (for example, 192.168.0.2 with a subnet mask of 255.255.255.0).

The management module has an MDI-X port that supports straight-through or crossover Ethernet cables. Use a Cat5e or greater cable to connect to the management module from its Ethernet jack (2, [Figure 23 \(page 34\)](#)) to the management server.

Figure 23 Management module



1. Status LEDs
2. Ethernet jack
3. Reset button

Connecting through a public network

1. Initialize the EVA4400 storage system, using HP P6000 SmartStart or HP P6000 Command View.
2. If currently connected, disconnect the public network LAN cable from the back of the management module in the controller enclosure.
3. Press and hold the recessed Reset button (3, [Figure 23 \(page 34\)](#)) for 4–5 seconds.

If you are running HP Command View EVA 9.3 or later, the green LED on the management module (1, [Figure 23 \(page 34\)](#)) blinks to indicate the configuration reset has started. The reset may take up to 2 minutes to complete. When the reset is completed, the green LED turns solid. This sets IP addresses of 192.168.0.1/24 (IPv4) and fd50:f2eb:a8a::7/48 (IPv6). If you are running a version earlier than HP Command View EVA 9.3 on the management module, the amber LED will flash momentarily when the reset is completed.

-
- ① **IMPORTANT:** At initial setup, you cannot browse to the HP P6000 Control Panel until you perform this step.
-
4. Do one of the following:
- Temporarily connect a LAN cable from a private network to the management module.
 - Temporarily connect a laptop computer to the management module using a LAN patch cable.
5. Browse to `https://192.168.0.1:2373` or `https://[fd50:f2eb:a8a::7]:2373/` and log in as an HP EVA administrator. HP recommends that you either change or delete the default IPv4 or IPv6 addresses to avoid duplicate address detection issues on your network. The default user name is `admin`. No password is required during the initial setup.
- The HP P6000 control panel GUI appears.
-

NOTE:

- If you are running a version of management module software earlier than 0001.1000 (which corresponds to XCS 09004000 or earlier), use port 2372.
 - If you change the password for the administrator or user account for the HP P6000 Control Panel, be sure to record the new passwords since they cannot be cleared without resetting the management module. HP recommends that you change the default passwords.
-

6. Select **Administrator Options > Configure Network Options**.
7. Enter an IP address and other network settings that apply.
-

NOTE: The management module firmware reserves IP addresses for internal communications. If you are using HP Command View EVA 9.2 or earlier, the reserved IP addresses are 10.0.0.1 through 10.0.0.3. Later versions of the management module firmware will also use these addresses if the array has older controller firmware loaded. The latest versions of controller and management module firmware, however, use the reserved IP address range from 10.253.251.230 through 10.253.251.249. The management module cannot be configured to use these addresses, nor can it communicate with external systems that fall within the reserved address range (for example, web client, router, DNS server, and so on).

8. Click **Save Changes**.

The IP address changes immediately, causing you to lose connectivity to the HP P6000 Control Panel.

The new IP address is stored and remains in effect, even when the storage system is shut down or restarted.

NOTE: If you uninitialize the storage system, wait one minute before reinitializing or you may lose the IP address.

9. Remove the LAN cable to the private network and reconnect the cable to the public network.
10. From a computer on the public network, browse to `https://newly configured ip address:2373` and log in.
- The HP P6000 control panel GUI appears.
-

Connecting through a private network

1. Press and hold the recessed Reset button (3, [Figure 23 \(page 34\)](#)) for 4–5 seconds.
- If you are running HP Command View EVA 9.3 or later, the green LED on the management module (1, [Figure 23 \(page 34\)](#)) blinks to indicate the configuration reset has started. The reset may take up to 2 minutes to complete. When the reset is completed, the green LED turns solid. This sets IP addresses of 192.168.0.1/24 (IPv4) and fd50:f2eb:a8a::7/48 (IPv6). If

you are running a version earlier than HP Command View EVA 9.3 on the management module, the amber LED will flash momentarily when the reset is completed.

2. Browse to `https://192.168.0.1:2373` and log in as an HP EVA administrator. HP recommends that you either change or delete the default IPv4 or IPv6 addresses to avoid duplicate address detection issues on your network. The default user name is `admin`. No password is required.

The HP P6000 control panel GUI appears.

NOTE: If you are running a version of management module software earlier than 0001.1000 (which corresponds to XCS 09004000 or earlier), use port 2372. At initial setup, you cannot browse to the HP P6000 Control Panel until you perform this step.

3. Select **Administrator Options > Configure Network Options**.
4. Enter an IP address and other network settings that apply.

NOTE: The management module firmware reserves IP addresses for internal communications. If you are using HP Command View EVA 9.2 or earlier, the reserved IP addresses are 10.0.0.1 through 10.0.0.3. Later versions of the management module firmware will also use these addresses if the array has older controller firmware loaded. The latest versions of controller and management module firmware, however, use the reserved IP address range from 10.253.251.230 through 10.253.251.249. The management module cannot be configured to use these addresses, nor can it communicate with external systems that fall within the reserved address range (for example, web client, router, DNS server, and so on).

5. Click **Save Changes**.

The IP address changes immediately, causing you to lose connectivity to the HP P6000 Control Panel. The new IP address is stored and remains in effect, even when the storage system is shut down or restarted.

NOTE: If you uninitialize the storage system, wait one minute before reinitializing or you may lose the IP address.

- ① **IMPORTANT:** The IP address will be lost if the storage system is uninitialized or the management module is reset.

6. From a computer on the private network, browse to `https://newly configured ip address:2373` and log in.

The HP P6000 control panel GUI appears.

Changing the default operating mode

By default, the EVA4400 is shipped to operate in a Fibre Channel switch environment and is configured in *fabric* mode. If you choose to connect the EVA4400 directly to a server, you must change the default operating mode of the EVA4400 to *direct* mode. If you do not change this mode, the EVA4400 will be unable to communicate with your server. Use the HP P6000 Control Panel to change the default operating mode.

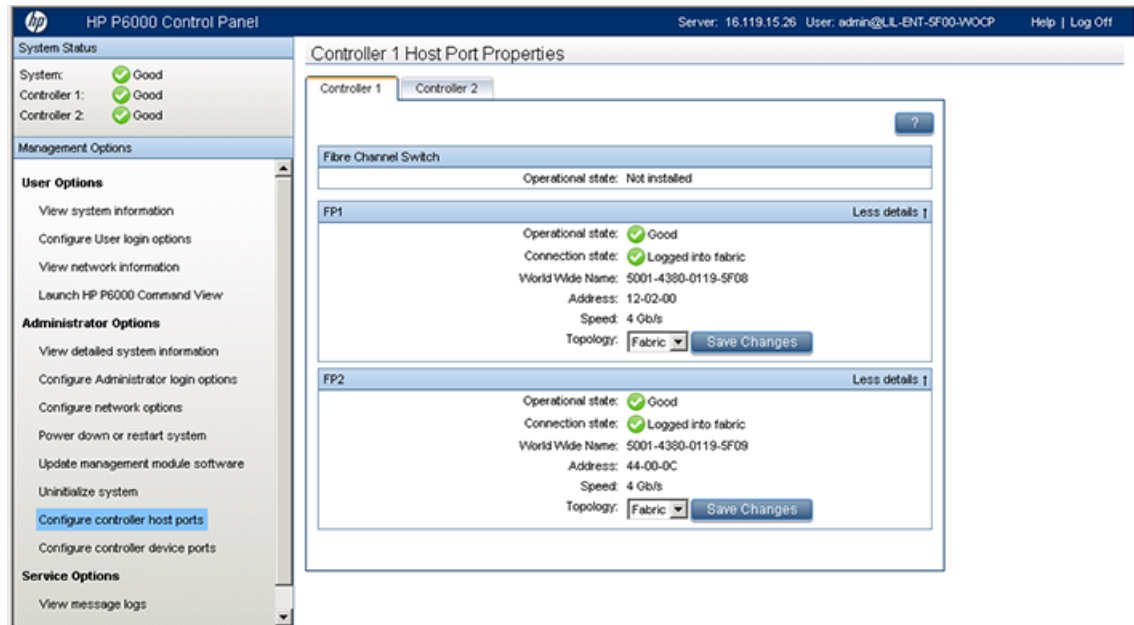
NOTE: You must have administrator privilege to change the settings in the HP P6000 Control Panel. Change your browser settings for the HP P6000 Control Panel window as described in the *HP P6000 Command View Installation Guide*.

To change the default operating mode:

1. Connect to the management module using one of the methods described in “Connecting through a public network” (page 34) or “Connecting through a private network” (page 35).
2. Log into the HP P6000 Control Panel as an administrator. The default username is admin and the password field is blank. For security reasons, change the password after you log in.
3. Select **Administrator Options > Configure controller host ports**.

The HP P6000 Control Panel screen appears.

Figure 24 HP P6000 Control Panel Configure controller host ports page



- In the **Topology** box, select **Direct** from the drop-down menu.
- Click **Save Changes**.
- Change the default operating mode for the other controller.
- Close the HP P6000 Control Panel and remove the Ethernet cable from the server.

Accessing the HP P6000 Control Panel through HP P6000 Command View

1. Select the EVA4400 storage system to access.
2. Select **Hardware→Controller Enclosure→Management Module**.
3. Click **Launch HP P6000 Control Panel**.

Rarely, the Launch HP P6000 Control Panel button may not appear due to invalid array state or if the management module IP address is unknown to HP P6000 Command View.

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 if re-initializing the storage system becomes necessary. The configuration data is saved to a series of files, which should be stored in a location other than on the storage system.

You can perform this procedure from the management server where HP P6000 Command View is installed from any host running HP Storage System Scripting Utility and connected to the management server.

NOTE: For more information on using the HP Storage System Scripting Utility, see the *HP Storage System Scripting Utility Reference*. See [“Documents” \(page 106\)](#).

1. Double-click 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 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 HP Storage System Scripting 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*.

The following examples illustrate how to save and restore the storage system configuration data using SSSU on a Windows host.

Example 1 Saving configuration data on a Windows host

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 SSSU command window.
-

Example 2 Restoring configuration data on a Windows host

If it is necessary to restore the storage system configuration, it can be done using the following procedure.

1. Double-click 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. Use files in sequential order. For example, use Step1A before Step1B, and so on. Files that are not needed for configuration data are not created, so there is no need to restore them.
-

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.

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

- Set the add disk option to manual.
- Add disks one at a time, waiting a minimum of 60 seconds between disks.
- Distribute disks vertically and as evenly as possible to all disk enclosures.
- Unless otherwise indicated, use the `SET DISK_GROUP` command in the HP Storage System Scripting Utility to add new disks to existing disk groups.
- Add disks in groups of eight.
- 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.

See the disk drive replacement instructions for the steps to add a disk drive. See [“Replacement instructions” \(page 73\)](#) for a link to this document.

Handling fiber optic cables

This section provides protection 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 should be installed whenever the transceivers or cables are disconnected.

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.

One of the many sources for cleaning equipment specifically designed for fiber optic connectors is:

Alcoa Fujikura Ltd. 1-888-385-4587 (North America)

011-1-770-956-7200 (International)

3 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.

NOTE: You can use HP P6000 SmartStart to configure Windows application servers. See the *HP 4400 Enterprise Virtual Array Installation Guide* or the HP P6000 SmartStart documentation for 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 EVA4400 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 14 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
	Microsoft Windows 2012

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 P6000 Command View 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 EVA4400.
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 [“Documents” \(page 106\)](#) 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 EVA4400 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 EVA4400 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/rdisk/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/rdisk/c10t0d2
disk      47 0/6/0/0.39.28.0.0.0.3  sdisk       CLAIMED   DEVICE      HP HSV300 /dev/dsk/c10t0d3
                                /dev/rdisk/c10t0d3
disk      48 0/6/0/0.39.28.0.0.0.4  sdisk       CLAIMED   DEVICE      HP HSV300 /dev/dsk/c10t0d4
                                /dev/rdisk/c10t0d4
disk      49 0/6/0/0.39.28.0.0.0.5  sdisk       CLAIMED   DEVICE      HP HSV300 /dev/dsk/c10t0d5
                                /dev/rdisk/c10t0d5
disk      50 0/6/0/0.39.28.0.0.0.6  sdisk       CLAIMED   DEVICE      HP HSV300 /dev/dsk/c10t0d6
                                /dev/rdisk/c10t0d6
disk      51 0/6/0/0.39.28.0.0.0.7  sdisk       CLAIMED   DEVICE      HP HSV300 /dev/dsk/c10t0d7
                                /dev/rdisk/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

HBA drivers

For most configurations and the latest version of Linux distributions, native HBA drivers are the supported drivers. *Native driver* means the driver that is included with the OS distribution.

NOTE: The term *inbox driver* is also sometimes used and means the same as *native driver*.

However, in some configurations, it may require the use of an out-of-box driver, which typically requires a driver package be downloaded and installed on the host. In those cases, follow the documentation of the driver package for instruction. Driver support information can be found on the Single Point of Connectivity Knowledge (SPOCK) website:

<http://www.hp.com/storage/spock>

NOTE: Registration is required to access SPOCK.

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.
 - Execute the following command (where *X* is the SCSI host enumerator of the HBA):
`echo "- - -" > /sys/class/scsi_host/host[X]/scan`

- 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 EVA4400 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 15 \(page 47\)](#) 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 15 Comparing console LUN to OS unit ID

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

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> drvdfg 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 EVA4400 console LUN can be seen without any virtual disks presented. The LUN appears as \$1\$GGAx (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 PGAO.5000-1FE1-0027-0A38 (BRCK18), primary path.			
Error count	0	Operations completed	145
Path PGAO.5000-1FE1-0027-0A3A (BRCK18).			
Error count	0	Operations completed	338
Path PGAO.5000-1FE1-0027-0A3E (BRCK18).			
Error count	0	Operations completed	276
Path PGAO.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 Oracle 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, go log in to 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 EVA4400 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:

```
fcg-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” \(page 50\)](#).

Configuring QLogic FCAs with the qla2300 driver

See the latest Enterprise Virtual Array release notes 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 (see <http://www.hp.com/storage/spock>).
2. 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).
3. 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 an Oracle-branded QLogic FCA, the configuration file is `\kernel\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.

4. 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.

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 `SANsurfer` utility.

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

The following example for an EVA4400 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.

6. 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.

7. 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/QLLogic_Corporation/drvutil/qla2300/qlreconfig -d qla2300 -s
```

8. 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 50).

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)/Oracle Storage Multipathing

Oracle StorEdge Traffic Manager (MPxIO)/Oracle 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 50) 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   HSV300, HSV400, HSV450

# vxddladm listsupport libname=libvxhpevale.so
ATTR_NAME            ATTR_VALUE
=====
LIBNAME               libvxhpevale.so
VID                   HP
PID                   HSV300, HSV400, HSV450
ARRAY_TYPE            A/A-A-HP
ARRAY_NAME            EVA4400, EVA6400, EVA8400

# vxddmpadm listapm all | grep HP
dmphpalua            dmphpalua            1            A/A-A-HP            Active
# vxddmpadm 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

# vxddmpadm listenclosure all
ENCLR_NAME            ENCLR_TYPE            ENCLR_SNO            STATUS            ARRAY_TYPE
=====
Disk                  Disk                  DISKS                CONNECTED         Disk
EVA44000              EVA4400              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 I/O policy”](#) shows the commands you can use to check and change the I/O policy.

Example 4 Setting the I/O policy

```
# vxddmpadm getattr arrayname EVA4400 iopolicy
ENCLR_NAME      DEFAULT      CURRENT
=====
EVA44000        Round-Robin   Round-Robin

# vxddmpadm setattr arrayname EVA4400 iopolicy=adaptive

# vxddmpadm getattr arrayname EVA4400 iopolicy
ENCLR_NAME      DEFAULT      CURRENT
=====
EVA44000        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](#).

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/c5t600508B4001030E40000500000B20000d0s2`
 - 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/c3t50001FE1002709E8d5s2`
`/dev/rdsk/c3t50001FE1002709ECd5s2`
`/dev/rdsk/c4t50001FE1002709E9d5s2`
`/dev/rdsk/c4t50001FE1002709EDd5s2`
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 4650001F
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/c3t50001FE1002709E8d5s2`
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
c3t50001FE1002709FCd2: 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 <HP-HSV400-0952 cyl 126 alt 2 hd 128 sec 128>
   /pci@1f,4000/QLGC,qla@4/fp@0,0/ssd@w50001fe1002709f8,1
2. c2t50001FE1002709F8d2 <HP-HSV400-0952 cyl 126 alt 2 hd 128 sec 128>
   /pci@1f,4000/QLGC,qla@4/fp@0,0/ssd@w50001fe1002709f8,2
3. c2t50001FE1002709FCd1 <HP-HSV400-0952 cyl 126 alt 2 hd 128 sec 128>
   /pci@1f,4000/QLGC,qla@4/fp@0,0/ssd@w50001fe1002709fc,1
4. c2t50001FE1002709FCd2 <HP-HSV400-0952 cyl 126 alt 2 hd 128 sec 128>
   /pci@1f,4000/QLGC,qla@4/fp@0,0/ssd@w50001fe1002709fc,2
5. c3t50001FE1002709F9d1 <HP-HSV400-0952 cyl 126 alt 2 hd 128 sec 128>
   /pci@1f,4000/lpfc@5/fp@0,0/ssd@w50001fe1002709f9,1
6. c3t50001FE1002709F9d2 <HP-HSV400-0952 cyl 126 alt 2 hd 128 sec 128>
   /pci@1f,4000/lpfc@5/fp@0,0/ssd@w50001fe1002709f9,2
7. c3t50001FE1002709FDd1 <HP-HSV400-0952 cyl 126 alt 2 hd 128 sec 128>
   /pci@1f,4000/lpfc@5/fp@0,0/ssd@w50001fe1002709fd,1
8. c3t50001FE1002709FDd2 <HP-HSV400-0952 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 Oracle 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 55)).
 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 1 through 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

Configuring the EVA4400 with VMware host servers

To configure an EVA4400 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.

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.

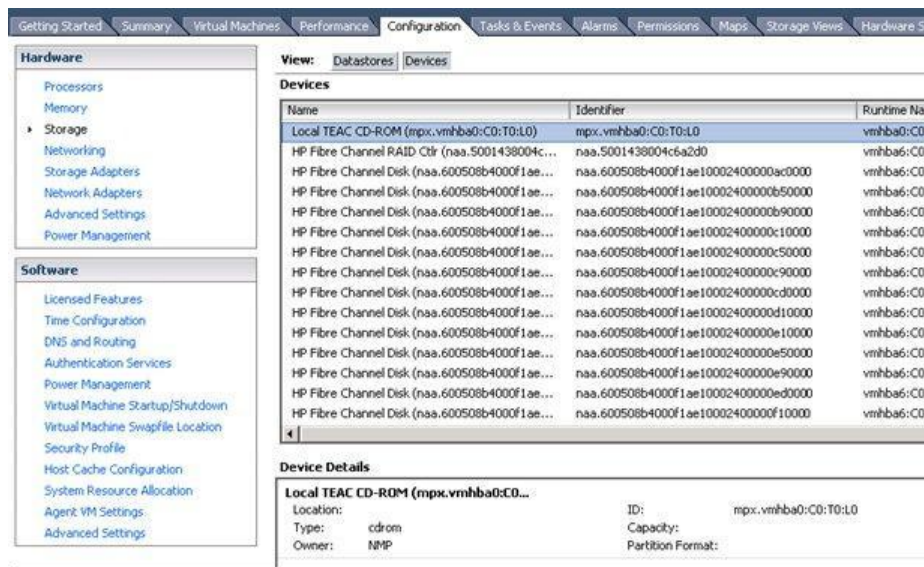
The output lists all SCSI devices detected by the server. An EVA4400 LUN entry looks similar to the following:

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

Verifying virtual disks from the host

Use the VMware vCenter management GUI to check all devices (see [Figure 25 \(page 61\)](#)).

Figure 25 Verifying virtual disks from the host



HP EVA P6000 Software Plug-in for VMware VAAI

The vSphere Storage API for Array Integration (VAAI) is included in VMware vSphere solutions. VAAI can be used to offload certain functions from the target VMware host to the storage array. With the tasks being performed more efficiently by the array instead of the target VMware host, performance can be greatly enhanced.

The HP EVA P6000 Software Plug-in for VMware VAAI (VAAI Plug-in) enables the offloading of the following functions (primitives) to the EVA:

- Full copy—Enables the array to make full copies of data within the array, without the ESX server having to read and write the data.
- Block zeroing—Enables the array to zero out a large number of blocks to speed up provisioning of virtual machines.
- Hardware assisted locking—Provides an alternative means to protect the metadata for VMFS cluster file systems, thereby improving the scalability of large ESX server farms sharing a datastore.

System prerequisites

VMware operating system:	ESX/ESXi 4.1 or later
VMware management station:	VMware vCenter 4.1,
VMware administration tools:	ESX/ESXi 4.1 environments: vCLI 4.1 (Windows or Linux)
HP P6000 controller software:	XCS 10100000 or later

Enabling vSphere Storage API for Array Integration (VAAI)

To enable the VAAI primitives, do the following:

NOTE: By default, the three VAAI primitives are enabled.

NOTE: The EVA VAAI Plug-In is required with vSphere 4.1 in order to permit discovery of the EVA VAAI capability. This is not required for vSphere 5.

1. Install the XCS controller software.

2. Enable the primitives from the ESX server.
Enable and disable these primitives through the following advanced settings:
 - DataMover.HardwareAcceleratedMove (full copy)
 - DataMover.HardwareAcceleratedInit (block zeroing)
 - VMFS3.HardwareAccelerated Locking (hardware assisted locking)
 For more information about the vSphere Storage API for Array Integration (VAAI), see the VMware documentation.
3. Install the HP EVA VAAI Plug-in.
For information about installing the VAAI Plug-in, see [“Installing the VAAI Plug-in” \(page 62\)](#).

Installing the VAAI Plug-in

Depending on user preference and environment, choose one of the following three methods to install the HP EVA VAAI Plug-in:

- Using ESX host console utilities
- vCLI/vMA
- Using VUM

The following table compares the three VAAI Plug-in installation methods:

Table 16 Comparison of installation methods

Installation method	Required deployment tools	Host Operating System	Client operating system	VMware commands used	Scriptable
ESX host console utilities—Local console	N/A	ESX 4.1	N/A	esxupdate esxcli	Yes (eva-vaaip.sh)
ESX host console utilities—Remote console	SSH tool, such as PuTTY		Any computer running SSH		
VMware CLI (vCLI)	VMware vSphere CLI	ESX 4.1, ESXi 4.1	Windows XP Windows Vista Windows Server 2003 Windows Server 2008 Linux x86 Linux x64	vicfg-hostops.pl vihostupdate.pl	Yes (eva-vaaip.pl)
VM Appliance (vMA)	N/A		N/A		
VMware Update Manager (VUM)	VMware vSphere Server VMware Update Manager	ESX 4.1, ESXi 4.1	Windows Server 2003, Windows Server 2008	VUM graphical user interface	No

Installation overview

Regardless of installation method, key installation tasks include:

1. Obtaining the HP VAAI Plug-in software bundle from the HP website.
2. Extracting files from HP VAAI Plug-in software bundle to a temporary location on the server.
3. Placing the target VMware host in maintenance mode.
4. Invoking the software tool to install the HP VAAI Plug-in.

Automated installation steps include:

- a. Installing the HP VAAI plug-in driver (hp_vaaip_p6000) on the target VMware host.
- b. Adding VIB details to the target VMware host.

- c. Creating VAAI claim rules.
 - d. Loading and executing VAAI claim rules.
5. Restarting the target VMware host.
6. Taking the target VMware host out of maintenance mode.

After installing the HP VAAI Plug-in, the operating system will execute all VAAI claim rules and scan every five minutes to check for any array volumes that may have been added to the target VMware host. If new volumes are detected, they will become VAAI enabled.

Installing the HP EVA VAAI Plug-in using ESX host console utilities

NOTE: This installation method is supported for use only with VAAI Plug-in Version 1.00, in ESX/ESXi 4.1 environments. This is required for ESX 4.1, but not for ESX 5i.

1. Obtain the VAAI Plug-in software package and save to a local folder on the target VMware host:
 - a. Go to the HP Support Download drivers and software website at <http://www.hp.com/support/downloads>.
 - b. Navigate through the display to locate and then download the HP EVA P6000 Software Plug-in for VMware VAAI to a temporary folder on the server. (Example folder location: /root/vaaip)
2. Install the VAAI Plug-in.

From the ESX service console, enter a command using the following syntax:

```
esxupdate --bundle hp_vaaip_p6000-xxx.zip --maintenancemode update
```

(where `hp_vaaip_p6000-xxx.zip` represents the filename of the VAAI Plug-in.)
3. Restart the target VMware host.

4. Verify the installation:

a. Check for new HP P6000 claim rules.

Using the service console, enter:

```
esxcli corestorage claimrule list -c VAAI
```

The return display will be similar to the following:

Rule	Class	Rule	Class	Type	Plugin	Matches
VAAI		5001	runtime	vendor	hp_vaaip_p6000	vendor=HP model=HSV
VAAI		5001	file	vendor	hp_vaaip_p6000	vendor=HP model=HSV

b. Check for claimed storage devices.

Using the service console, enter:

```
esxcli vaaip device list
```

The return display will be similar to the following:

```
aa.600c0ff00010e1cbc7523f4d01000000
Device Display Name: HP iSCSI Disk (naa.600c0ff00010e1cbc7523f4d01000000)
VAAI Plugin Name: hp_vaaip_P6000

naa.600c0ff000da030b521bb64b01000000
Device Display Name: HP Fibre Channel Disk (naa.600c0ff000da030b521bb64b01000000)
VAAI Plugin Name: hp_vaaip_P6000
```

c. Check the VAAI status on the storage devices.

Using the service console, enter:

```
esxconfig-scsidevs -l | egrep "Display Name:|VAAI Status:"
```

The return display will be similar to the following:

```
Display Name: Local TEAC CD-ROM (mpx.vmhba5:C0:T0:L0)
VAAI Status: unknown
Display Name: HP Serial Attached SCSI Disk (naa.600508b1001052395659314e39440200)
VAAI Status: unknown
Display Name: HP Serial Attached SCSI Disk (naa.600c0ff0001087439023704d01000000)
VAAI Status: supported
Display Name: HP Serial Attached SCSI Disk (naa.600c0ff0001087d28323704d01000000)
VAAI Status: supported
Display Name: HP Fibre Channel Disk (naa.600c0ff000f00186a622b24b01000000)
VAAI Status: unknown
```

Table 17 VAAI device status values

Value	Description
Unknown	The array volume is hosted by a non-supported VAAI Array (such as the EVA P6300 or P6500)
Supported	The array volume is hosted by a supported VAAI array, and all three VAAI commands completed successfully.
Not supported	The array volume is hosted by a supported VAAI array, but all three VAAI commands did not complete successfully.

NOTE: VAAI device status is "Unknown" until all VAAI primitives are attempted by ESX on the device and completed successfully. Upon completion, VAAI device status becomes "Supported."

Installing the HP VAAI Plug-in using vCLI/vMA

NOTE: This installation method is supported for use only with VAAI Plug-in version 1.00, in ESX/ESXi 4.1 environments.

1. Obtain the VAAI Plug-in software package and save to a local folder on the target VMware host:
 - a. Go to the HP Support Downloads website at <http://www.hp.com/support/downloads>.
 - b. Navigate through the display to locate and then download the HP EVA P6000 Software Plug-in for VMware VAAI to a temporary folder on the server. (Example folder location: /root/vaaip)
2. Enter maintenance mode.
Enter a command using the following syntax:

```
vicfg-hostops.pl --server Host_IP_Address --username User_Name --password Account_Password -o enter
```
3. Install the VAAI Plug-in using vihostupdate.
Enter a command using the following syntax:

```
vihostupdate.pl --server Host_IP_Address --username User_Name --password Account_Password --bundle hp_vaaip_p6000_offline-bundle-xyz --install
```
4. Restart the target VMware host.
Enter a command using the following syntax:

```
vicfg-hostops.pl --server Host_IP_Address --username User_Name --password Account_Password -o reboot -f
```
5. Exit maintenance mode.
Enter a command using the following syntax:

```
vicfg-hostops.pl --server Host_IP_Address --username User_Name --password Account_Password -o exit
```
6. Verify the claimed VAAI device.
 - a. Check for new HP P6000 claim rules.
Enter a command using the following syntax:

```
esxcli --server Host_IP_Address --username User_Name --password Account_Password corestorage claimrule list -c VAAI
```

The return display will be similar to the following:

Rule	Class	Rule	Class	Type	Plugin	Matches
VAAI		5001	runtime	vendor	hp_vaaip_p6000	vendor=HP model=HSV
VAAI		5001	file	vendor	hp_vaaip_p6000	vendor=HP model=HSV
 - b. Check for claimed storage devices.
List all devices claimed by the VAAI Plug-in.
Enter a command using the following syntax:

```
esxcli --server Host_IP_Address --username User_Name --password Account_Password vaaip device list
```

The return display will be similar to the following:

```
naa.600c0ff00010e1cbc7523f4d01000000
  Device Display Name: HP iSCSI Disk (naa.600c0ff00010e1cbc7523f4d01000000)
  VAAI Plugin Name: hp_vaaip_p6000

naa.600c0ff000da030b521bb64b01000000
  Device Display Name: HP Fibre Channel Disk (naa.600c0ff000da030b521bb64b01000000)
  VAAI Plugin Name: hp_vaaip_p6000
```
 - c. Check the VAAI status on the storage devices. For more information about VAAI device status values, see [Table 17 \(page 64\)](#) .
Use the vCenter Management Station as listed in the following section.

NOTE:

- This installation method is supported for use with VAAI Plug-in versions 1.00 and 2.00, in ESX/ESXi 4.1 environments.
 - HP recommends installing the plug-in using VMware Update Manager.
-

Installing the VAAI Plug-in using VUM consists of two steps:

1. “Importing the VAAI Plug-in to the vCenter Server” (page 66)
2. “Installing the VAAI Plug-in on each ESX/ESXi host” (page 67)

Importing the VAAI Plug-in to the vCenter Server

1. Obtain the VAAI Plug-in software package and save it on the system that has VMware vSphere client installed:
 - a. Go to the HP Download drivers and support website at <http://www.hp.com/support/downloads>.
 - b. Locate the HP EVA P6000 Software Plug-in for VMware VAAI and then download it to a temporary folder on the server.
 - c. Expand the contents of the downloaded .zip file into the temporary folder and locate the HP EVA VAAI offline bundle file. The filename will be in one of the following formats:
`hp_vaaip_p6000_offline-bundle_xyz.zip`
(where xyz represents the VAAI Plug-in version.)
2. Open VUM:
 - a. Double-click the **VMware vSphere Client** icon on your desktop, and then log in to the vCenter Server using administrator privileges.
 - b. Click the **Home** icon in the navigation bar.
 - c. In the Solutions and Applications pane, click the **Update Manager** icon to start VUM.

NOTE: If the Solutions and Applications pane is missing, the VUM Plug-in is not installed on your vCenter Client system. Use the vCenter Plug-ins menu to install VUM.

3. Import the Plug-in:
 - a. Select the **Patch Repository** tab.
 - b. Click **Import Patches** in the upper right corner. The Import Patches dialog window opens.
 - c. Browse to the extracted HP EVA VAAI offline bundle file. The filename will be in the following format:
`hp_vaaip_p6000_offline-bundle_xyz.zip`
(where xyz represents the VAAI Plug-in version).
 - d. Wait for the import process to complete.
 - e. Click **Finish**.

4. Create a new Baseline set for this offline plug-in:
 - a. Select the **Baselines and Groups** tab.
 - b. Above the left pane, click **Create**.
 - c. In the New Baseline window:
 - Enter a name and a description. (Example: HP P6000 Baseline and VAAI Plug-in for HP EVA)
 - Select **Host Extension**.
 - Click **Next** to proceed to the Extensions window.
 - d. In the Extensions window:
 - Select **HP EVA VAAI Plug-in for VMware vSphere x.x**, where x.x represents the plug-in version.
 - Click the down arrow to add the plug-in in the Extensions to Add panel at the bottom of the display.
 - Click **Next** to proceed.
 - Click **Finish** to complete the task and return to the **Baselines and Groups** tab.

The HP P6000 Baseline should now be listed in the left pane.

Importing the VAAI Plug-in is complete. To install the plug-in, see [“Installing the VAAI Plug-in on each ESX/ESXi host” \(page 67\)](#).

Installing the VAAI Plug-in on each ESX/ESXi host

1. From the vCenter Server, click the **Home** icon in the navigation bar.
2. Click the **Hosts and Clusters** icon in the Inventory pane.
3. Click the DataCenter that has the ESX/ESXi hosts that you want to stage.
4. Click the **Update Manager** tab. VUM automatically evaluates the software recipe compliance for all ESX/ESXi Hosts.
5. Above the right pane, click **Attach** to open the **Attach Baseline or Group** dialog window. Select the HP P6000 Baseline entry, and then click **Attach**.
6. To ensure that the patch and extensions compliance content is synchronized, again click the DataCenter that has the ESX/ESXi hosts that you want to stage. Then, in the left panel, right-click the **DataCenter** icon and select **Scan for Updates**. When prompted, ensure that Patches and Extensions is selected, and then click **Scan**.
7. Stage the installation:
 - a. Click **Stage** to open the Stage Wizard.
 - b. Select the target VMware hosts for the extension that you want to install, and then click **Next**.
 - c. Click **Finish**.
8. Complete the installation:
 - a. Click **Remediate** to open the Remediation Wizard.
 - b. Select the target VMware host that you want to remediate, and then click **Next**.
 - c. Make sure that the HP EVA VAAI extension is selected, and then click **Next**.
 - d. Fill in the related information, and then click **Next**.
 - e. Click **Finish**.

Installing the VAAI Plug in is complete. View the display for a summary of which ESX/ESXi hosts are compliant with the vCenter patch repository.

NOTE:

- In the Tasks & Events section, the following tasks should have a Completed status: Remediate entry, Install, and Check.
 - If any of the above tasks has an error, click the task to view the detail events information.
-

Verifying VAAI status

1. From the vCenter Server, click the **Home Navigation** bar and then click **Hosts and Clusters**.
2. Select the target VMware host from the list and then click the **Configuration** tab.
3. Click the **Storage Link** under Hardware.

For more information about VAAI device status values, see [Table 17 \(page 64\)](#).

Uninstalling the VAAI Plug-in

Procedures vary, depending on user preference and environment:

Uninstalling VAAI Plug-in using the automated script (hpeva.pl)

1. Enter maintenance mode.
2. Query the installed VAAI Plug-in to determine the name of the bulletin to uninstall.
Enter a command using the following syntax:

```
c:\>hpeva.pl --server Host_IP_Address --username User_Name --password Account_Password --query
```
3. Uninstall the VAAI Plug-in.
Enter a command using the following syntax:

```
c:\>hpeva.pl --server Host_IP_Address --username User_Name --password Account_Password --bulletin Bulletin_Name --remove
```
4. Restart the host.
5. Exit maintenance mode.

Uninstalling VAAI Plug-in using vCLI/vMA (vihostupdate)

1. Enter maintenance mode.
2. Query the installed VAAI Plug-in to determine the name of the VAAI Plug-in bulletin to uninstall.
Enter a command using the following syntax:

```
c:\>vihostupdate.pl --server Host_IP_Address --username User_Name --password Account_Password --query
```
3. Uninstall the VAAI Plug-in.
Enter a command using the following syntax:

```
c:\>vihostupdate.pl --server Host_IP_Address --username User_Name --password Account_Password --bulletin 0-HPQ-ESX-4.1.0-hp-vaaip-p6000-1.0.10 --remove
```
4. Restart the host.
5. Exit maintenance mode.

Uninstalling VAAI Plug-in using VMware native tools (esxupdate)

1. Enter maintenance mode.
2. Query the installed VAAI Plug-in to determine the name of the VAAI Plug-in bulletin to uninstall.
Enter a command using the following syntax:

```
$host# esxupdate --vib-view query | grep hp-vaaip-p6000
```

3. Uninstall the VAAI Plug-in.

Enter a command using the following syntax:

```
$host# esxupdate remove -b VAAI_Plug_In_Bulletin_Name  
--maintenancemode
```

4. Restart the host.
5. Exit maintenance mode.

4 Replacing array components

Customer self repair

Table 18 (page 71) and Table 19 (page 72) identify hardware components that 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 components, videos of the procedures have been produced. To view the videos, go to the following website and navigate to your product:

<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 the HP fault monitoring 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 System Event Analyzer online help for additional information.

Identifying the spare part

Parts have a nine-character spare part number on their label. See Figure 26 (page 71). For some spare parts, the part number is available in HP P6000 Command View. Alternatively, the HP call center will assist in identifying the correct spare part number.

Figure 26 Example of typical product label



8053A-ST

1. Spare part number

Replaceable parts

This product contains the replaceable parts listed in [“Controller enclosure replacement parts” \(page 71\)](#) and [“M6412-A disk enclosure replaceable parts” \(page 72\)](#). Parts that are available for customer self repair (CSR) are indicated as follows:

- ✓ Mandatory CSR 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 18 Controller enclosure replacement parts

Description	Spare part number (non RoHS/RoHS)	CSR status
4Gb array controller (HSV300)	461488-005	•
4Gb array controller (HSV300-S)	460586-005	•
HSV300-S embedded switch console port cable	316131-001	✓
1GB cache DIMM	466263-001 (uses 512 Mb DRAM chips)	•
	635205-001 (uses 1 Gb DRAM chips) ¹	
Array battery	460581-001	✓
Array power supply	519842-001	✓
Array fan module	460583-001	✓
Array management module	460584-005	✓
Array LED membrane display	461489-001	•
Array midplane	461490-005	•
Array riser assembly	461491-001	•
Array power UID	466264-001	•
Array front bezel	460585-001	✓

¹ Requires XCS 09522000 or later.

Table 19 M6412–A disk enclosure replaceable parts

Description	Spare part number (non RoHS/RoHS)	CSR status
4 Gb FC disk shelf midplane	461492-001, 635829-001	•
4 Gb FC disk shelf backplane	461493-001, 461493-005	•
SPS-BD Front UID	399053-001	•
SPS-BD Power UID with cable	399054-001	•
SPS-BD Front UID Interconnect PCA with cable	399055-001	•
4 Gb FC disk shelf I/O module	461494-001, 461494-005	•
FC disk shelf fan module	468715-001	✓
FC disk shelf power supply	405914-001	✓
Filler panel for unused drive bay	389015-001	✓
Disk drive 300 GB, 10K, EVA M6412–A Enclosure, Fibre Channel	537582-001	✓
Disk drive 450 GB, 10K, EVA M6412–A Enclosure, Fibre Channel	518734-001	✓
Disk drive 600 GB, 10K, EVA M6412–A Enclosure, Fibre Channel	518735-001	✓
Disk drive 146 GB, 15K, EVA M6412–A Enclosure, Fibre Channel	454410-001	✓
Disk drive 300 GB, 15K, EVA M6412–A Enclosure, Fibre Channel	454411-001	✓
Disk drive 400 GB, 15K, EVA M6412–A Enclosure, Fibre Channel	466277-001	✓
Disk drive 450 GB, 15K, EVA M6412–A Enclosure, Fibre Channel	454412-001	✓
Disk drive 600 GB, 15K, EVA M6412–A Enclosure, Fibre Channel	495808-001	✓
Disk drive 1 TB, 7.2K, EVA M6412–A Enclosure, FATA	671148-001	✓
Disk drive 72 GB, EVA M6412–A Enclosure, SSD	515189-001	✓
Disk drive 200 GB, EVA M6412–A Enclosure, SSD	595336-001	✓
Disk drive 400 GB, EVA M6412–A Enclosure, SSD	595337-001	✓
Disk drive 600 GB, 15K, Fibre Channel	621832-001	✓
Rail kit, 2U form factor	383663-001	✓
SPS-CABLE ASSY, 4Gb COPPER, FC, 2.0m	432374-001	✓
SPS-CABLE ASSY, 4Gb COPPER, FC, 0.6m	432375-001	✓
SPS-CABLE ASSY, 4Gb COPPER, FC, 0.41 m	496917-001	✓

For more information about CSR, contact your local service provider or 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 (ESD). 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 instructions

Printed instructions are shipped with the replacement part. Instructions for all replaceable components are also included on the documentation CD that ships with the EVA4400 and posted on the HP website. For the latest information, HP recommends that you obtain the instructions from the HP Business Support Center Manuals website at <http://www.hp.com/support/manuals>. Under Storage, select **Disk Storage Systems**, and then under EVA Storage, select **HP 4400 Enterprise Virtual Array**. The manuals page for the EVA4400 appears. Scroll to the Service and maintenance information section where the following replacement instructions are posted:

- *HP Controller Enclosure 4Gb Array Controller Replacement Instructions*
- *HP Controller Enclosure 4Gb Array Controller with Embedded Switch Replacement Instructions*
- *HP Controller Enclosure Battery Replacement Instructions*
- *HP Controller Enclosure Cache DIMM Replacement Instructions*
- *HP Controller Enclosure Fan Module Replacement Instructions*
- *HP Controller Enclosure LED Display Replacement Instructions*
- *HP Controller Enclosure Management Module Replacement Instructions*
- *HP Controller Enclosure Midplane Replacement Instructions*
- *HP Controller Enclosure Power Supply Replacement Instructions*
- *HP Controller Enclosure Riser Assembly Replacement Instructions*
- *HP Disk Enclosure Backplane Replacement Instructions*
- *HP Disk Enclosure Fan Module Replacement Instructions*
- *HP Disk Enclosure Front UID Interconnect Board (with cable) Replacement Instructions*
- *HP Disk Enclosure Front UID Replacement Instructions*

- *HP Disk Enclosure I/O Module Replacement Instructions*
- *HP Disk Enclosure Midplane Replacement Instructions*
- *HP Disk Enclosure Power Supply Replacement Instructions*
- *HP Fibre Channel Disk Drive Replacement Instructions*
- *HP Power UID Replacement Instructions*

5 Single path implementation

This chapter provides guidance for connecting servers with a single path HBA to the EVA storage system with no multipath software installed. A single path HBA is defined as:

- A single HBA port to a switch with no multipathing software installed
- A single HBA port to a switch with multipathing software installed

HBA LUNs are not shared by any other HBA in the server or in the SAN. Failover action is different depending on which single path method is employed.

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 multipath EVA environment, configurations should follow these recommendations.

The purpose of single HBA configurations for non-mission critical storage access is to control costs. This chapter describes the configurations, limitations, and failover characteristics of single HBA servers under different operating systems. Several of the descriptions are based on a single HBA configuration resulting in a single path to the device, but OpenVMS has native multipath features by default.

NOTE: Tru64 UNIX is not supported with the EVA4400.

With OpenVMS, a single HBA configuration will result in two paths to the device by having connections to both EVA controllers. Single HBA configurations are not single path configurations with these operating systems.

In addition, cluster configurations for OpenVMS provide enhanced availability and security. To achieve availability within cluster configurations, configure each member with its own HBAs and connectivity to shared LUNs. For further information on cluster configurations and attributes, see the appropriate operating system guide and the *HP SAN Design Reference 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 following HP website:

<http://www.hp.com/country/us/eng/prodserv/storage.html>

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 2012, Microsoft Windows Server 2008, Microsoft Windows Server, Linux and IBM AIX operating systems, the zones consist of the single path HBA systems and one HSV controller port.
- On OpenVMS, the zones consist of the single HBA systems and two HSV controller ports. This results 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 multipath software, high data availability and fault tolerance are achieved. In single path HBA server configurations, neither multipath 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 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 *HP SAN Design Reference Guide* for heterogeneous SANs, located at:

<http://www.hp.com/go/sdgmanuals>

General configuration components

All configurations require the following components:

- XCS controller 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 the OpenVMS 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://www.hp.com/go/sdgmanuals>

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 27 (page 77) and Figure 28 (page 78) 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 multipath software that manages the two HBAs and their connections to the switch, 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 28 (page 78), servers with OpenVMS can be zoned with two controllers.


Figure 27 Single path HBA server without OpenVMS



26463a

- | | |
|-------------------------------|------------------|
| 1. Network interconnection | 6. SAN switch 1 |
| 2. Single HBA server (Host 1) | 7. SAN switch 2 |
| 3. Dual HBA server (Host 2) | 8. Fabric zone |
| 4. Management server | 9. Controller A |
| 5. Multiple single HBA paths | 10. Controller B |

Figure 28 Single path HBA server with OpenVMS



- | | |
|-------------------------------|------------------|
| 1. Network interconnection | 6. SAN switch 1 |
| 2. Single HBA server (Host 1) | 7. SAN switch 2 |
| 3. Dual HBA server (Host 2) | 8. Fabric zone |
| 4. Management server | 9. Controller A |
| 5. Multiple single HBA paths | 10. Controller B |

26465a

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 multipathing software.

See [Figure 29](#) (page 79).

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 “[HP-UX](#)” (page 92).

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 29 HP-UX configuration



1. Network interconnection
2. Single HBA server (Host 1)
3. Dual HBA server (Host 2)
4. Management server

5. SAN switch 1
6. SAN switch 2
7. Controller A
8. Controller B

26464a

Windows Server 2003 (32-bit), Windows Server 2008 (32-bit), Windows Server 2012 (32-bit) configurations

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 multipathing software.

See [Figure 30 \(page 81\)](#).

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 [“Windows Servers” \(page 93\)](#).

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 30 Windows Server 2008 (32-bit), Windows Server 2003 (32-bit), and Windows 2000 configuration



- | | |
|-------------------------------|-----------------|
| 1. Network interconnection | 5. SAN switch 1 |
| 2. Single HBA server (Host 1) | 6. SAN switch 2 |
| 3. Dual HBA server (Host 2) | 7. Controller A |
| 4. Management server | 8. Controller B |

Windows Server 2008 (64-bit) and Windows Server 2003 (64-bit) configurations

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 multipathing software.

See [Figure 31 \(page 82\)](#).

NOTE: Single path HBA servers running the Windows Server 2003 (x64) or Microsoft Windows Server 2008 (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 “Windows Servers” (page 93).

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 31 Windows Server 2008 (64-bit) or Windows Server 2003 (64-bit) configuration



1. Network interconnection
2. Management server
3. Host 1
4. Host 2
5. Host 3

6. SAN switch 1
7. Multiple single HBA paths
8. SAN switch 2
9. Controller A
10. 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 multipathing software.

See [Figure 32 \(page 84\)](#).

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 [“Oracle Solaris” \(page 93\)](#).

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 32 Oracle Solaris configuration



- | | |
|-------------------------------|-----------------|
| 1. Network interconnection | 5. SAN switch 1 |
| 2. Single HBA server (Host 1) | 6. SAN switch 2 |
| 3. Dual HBA server (Host 2) | 7. Controller A |
| 4. Management server | 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 33 \(page 85\)](#).

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 “OpenVMS” (page 94).

Limitations

- HP P6000 Continuous Access is not supported with single path configurations.

Figure 33 OpenVMS configuration



26464a

- | | |
|-------------------------------|-----------------|
| 1. Network interconnection | 5. SAN switch 1 |
| 2. Single HBA server (Host 1) | 6. SAN switch 2 |
| 3. Dual HBA server (Host 2) | 7. Controller A |
| 4. Management server | 8. Controller B |

Xen 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 multipathing software.

See [Figure 34 \(page 86\)](#).


Risks

- Single path failure may result in data loss or disk corruption.

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 34 Xen configuration



1. Network interconnection
2. Single HBA server (Host 1)
3. Dual HBA server (Host 2)
4. Management server

5. SAN switch 1
6. SAN switch 2
7. Controller A
8. Controller B

Linux (32-bit) configuration

NOTE: The information in this section also applies to RHEL Xen configurations.

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 multipathing software.

See [Figure 35 \(page 88\)](#).

Risks


- Single path failure may result in data loss or disk corruption.

NOTE: For additional risks, see [“Linux” \(page 94\)](#).

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 35 Linux (32-bit) configuration



- | | |
|-------------------------------|-----------------|
| 1. Network interconnection | 5. SAN switch 1 |
| 2. Single HBA server (Host 1) | 6. SAN switch 2 |
| 3. Dual HBA server (Host 2) | 7. Controller A |
| 4. Management server | 8. Controller B |

26464a

Linux (Itanium) 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 to 14 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 is a single path HBA.
- Host 2 is a dual HBA host with multipathing software.

See [Figure 36 \(page 89\)](#).

Risks


- Single path failure may result in data loss or disk corruption.

NOTE: For additional risks, see “Linux” (page 94).

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 36 Linux (Itanium) configuration



1. Network interconnection
2. Single HBA server (Host 1)
3. Dual HBA server (Host 2)
4. Management server

5. SAN switch 1
6. SAN switch 2
7. Controller A
8. Controller B

26464a

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 configuration

- Host 1 is a single path HBA host.
- Host 2 is a dual HBA host with multipathing software.

See [Figure 37 \(page 90\)](#).

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 [“IBM AIX” \(page 95\)](#).

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 37 IBM AIX configuration



1. Network interconnection
2. Single HBA server (Host 1)
3. Dual HBA server (Host 2)
4. Management server

5. SAN switch 1
6. SAN switch 2
7. Controller A
8. Controller B

26464a

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 multipathing software.

See [Figure 38 \(page 92\)](#).

Risks


- Single path failure may result in data loss or disk corruption.

NOTE: For additional risks, see [“VMware” \(page 95\)](#).

Limitations

- HP P6000 Continuous Access is not supported with single path configurations.
- Single HBA path at the host server is not part of a cluster, unless in a VMware High Availability Cluster.
- Booting from the SAN is supported on single path HBA servers.

Figure 38 VMware configuration



- | | |
|-------------------------------|-----------------|
| 1. Network interconnection | 5. SAN switch 1 |
| 2. Single HBA server (Host 1) | 6. SAN switch 2 |
| 3. Dual HBA server (Host 2) | 7. Controller A |
| 4. Management server | 8. Controller B |

Mac OS configuration

For information about Mac OS connectivity, see *Mac OS X Fibre Channel connectivity to the HP Enterprise Virtual Array Storage System Configuration Guide* (to download, see [“Documents”](#) (page 106)).

Failure scenarios

HP-UX

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.

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 Servers

Windows Server 2003, Windows Server 2008, and Windows Server 2012

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

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.

Fault stimulus	Failure effect
Server failure (host power-cycled)	Nonclustered-Processes fail. Clustered—Other nodes running processes that used devices served from the single-path HBA failed over access to a different served path. When the single-path node crashes, only the processes executing on that node fail. In either case, no data is lost or corrupted.
Switch failure (SAN switch disabled)	I/O is suspended or process is terminated across this HBA until switch is back online. No data is lost or corrupted. The operating system 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.
Controller failure	I/O fails over to the surviving controller. No data is lost or corrupted.
Controller restart	I/O is suspended or process is terminated across this HBA until EVA is back online. No data is lost or corrupted. The operating system 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.
Server path failure	If the LUN is not shared, I/O is suspended or process is terminated across this HBA until path is restored. If running OpenVMS 7.3-1 and the LUN is shared, another cluster node having direct access will take over serving the device, resulting in no loss of service. In either case, no data is lost or corrupted. The operating system 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.
Storage path failure	I/O is suspended or process is terminated across this HBA until path is restored. No data is lost or corrupted. The operating system 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.

Linux

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.

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

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

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 term: I/O suspended, possible data loss. Long term: 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.

Fault stimulus	Failure effect
Server path failure	Short term: I/O suspended, possible data loss. Long term: 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 term: I/O suspended, possible data loss. Long term: 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.

Mac OS

Fault stimulus	Failure effect
Server failure (host power-cycled)	OS reboots. Both HFS and StorNext replay journal on filesystem. Disk auto mounted by OS.
Switch failure	Short term: I/O suspended, possible data loss. Long term: I/O halts with I/O errors.
Controller failure	Short term: I/O suspended, possible data loss. Long term: I/O fails over to alternate storage controller if visible (by zoning). Otherwise, I/O halts with I/O errors, data loss. Can require a server reboot for full recovery.
Controller restart	Short term: I/O suspended, possible data loss. Long term: I/O fails over to alternate storage controller if visible (by zoning). Otherwise, I/O halts with I/O errors, data loss. Can require a server reboot for full recovery.
Server path failure	Short term: I/O suspended, possible data loss. Long term: I/O halts with I/O errors, data loss. Can require a server reboot for full recovery.
Storage path failure	Short term: I/O suspended, possible data loss. Long term: I/O fails over to alternate storage controller if available. Otherwise, I/O halts with I/O errors. Can require a server reboot for full recovery.

6 Error messages

This list of error messages is in order by status code value, 0 to 100.

Table 20 Error messages

Status code value	Meaning	How to correct
0 Successful Status	The SCMI command completed successfully.	No corrective action required.
1 Object Already Exists	The object or relationship already exists.	Delete the associated object and try the operation again. Several situations can cause this message: Presenting a LUN to a host: <ul style="list-style-type: none">• Delete the current association or specify a different LUN number. Storage cell initialize: <ul style="list-style-type: none">• Remove or erase disk volumes before the storage cell can be successfully created. Adding a port WWN to a host: <ul style="list-style-type: none">• Specify a different port WWN. Adding a disk to a disk group: <ul style="list-style-type: none">• Delete the specified disk volume before creating a new disk volume.
2 Supplied Buffer Too Small	The command or response buffer is not large enough to hold the specified number of items. This can be caused by a user or program error.	Report the error to product support.
3 Object Already Assigned	The handle is already assigned to an existing object. This can be caused by a user or program error.	Report the error to product support.
4 Insufficient Available Data Storage	There is insufficient storage available to perform the request.	Reclaim some logical space or add physical hardware.
5 Internal Error	An unexpected condition was encountered while processing a request.	Report the error to product support.
6 Invalid status for logical disk	This error is no longer supported.	Report the error to product support.
7 Invalid Class	The supplied class code is of an unknown type. This can be caused by a user or program error.	Report the error to product support.
8 Invalid Function	The function code specified with the class code is of an unknown type.	Report the error to product support.
9 Invalid Logical Disk Block State	The specified command supplied unrecognized values. This can indicate a user or program error.	Report the error to product support.
10 Invalid Loop Configuration	The specified request supplied an invalid loop configuration.	Verify the hardware configuration and retry the request.
11 Invalid parameter	There are insufficient resources to fulfill the request, the requested value is not supported, or the parameters supplied are invalid. This can indicate a user or program error.	Report the error to product support.

Table 20 Error messages *(continued)*

Status code value	Meaning	How to correct
12 Invalid Parameter handle	The supplied handle is invalid. This can indicate a user error, program error, or a storage cell in an uninitialized state. In the following cases, the storage cell is in an uninitialized state, but no action is required: Storage cell discard (informational message): Storage cell look up object count (informational message): Storage cell look up object (informational message):	In the following cases, the message can occur because the operation is not allowed when the storage cell is in an uninitialized state. If you see these messages, initialize the storage cell and retry the operation. Storage cell set device addition policy Storage cell set name Storage cell set time Storage cell set volume replacement delay Storage cell free command lock Storage cell set console lun id
13 Invalid Parameter Id	The supplied identifier is invalid. This can indicate a user or program error.	Report the error to product support.
14 Invalid Quorum Configuration	Quorum disks from multiple storage systems are present.	Report the error to product support.
15 Invalid Target Handle	The supplied target handle is invalid. This can indicate a user or program error (Case 1), or Volume set requested usage (Case 2): The operation could not be completed because the disk has never belonged to a disk group and therefore cannot be added to a disk group.	Case 1: Report the error to product support. Case 2: To add additional capacity to the disk group, use the management software to add disks by count or capacity.
16 Invalid Target Id	The supplied target identifier is invalid. This can indicate a user or program error.	Report the error to product support.
17 Invalid Time	The time value specified is invalid. This can indicate a user or program error.	Report the error to product support.
18 Media is Inaccessible	The operation could not be completed because one or more of the disk media was inaccessible.	Report the error to product support.
19 No Fibre Channel Port	The Fibre Channel port specified is not valid. This can indicate a user or program error.	Report the error to product support.
20 No Image	There is no firmware image stored for the specified image number.	Report the error to product support.
21 No Permission	The disk device is not in a state to allow the specified operation.	The disk device must be in either maintenance mode or in a reserved state for the specified operation to proceed.
22 Storage system not initialized	The operation requires a storage cell to exist.	Create a storage cell and retry the operation.
23 Not a Loop Port	The Fibre Channel port specified is either not a loop port or is invalid. This can indicate a user or program error.	Report the error to product support.
24 Not a Participating Controller	The controller must be participating in the storage cell to perform the operation.	Verify that the controller is a participating member of the storage cell.
25 Objects in your system are in use, and their state prevents the operation you wish to perform.	Several states can cause this message: Case 1: The operation cannot be performed because an association exists a related object, or the object is in a progress state.	Case 1: Either delete the associated object or resolve the in progress state. Case 2: . Report the error to product support. Case 3: Unpresent the LUNs before deleting this virtual disk.

Table 20 Error messages (continued)

Status code value	Meaning	How to correct
	<p>Derived unit create: Case 2: The supplied virtual disk handle is already an attribute of another derived unit. This may indicate a programming error</p> <p>Derived unit discard: Case 3: One or more LUNs are presented to EVA hosts that are based on this virtual disk.</p> <p>Case 4: Logical disk clear data lost: The virtual disk is in the non-mirrored delay window.</p> <p>Case 5: LDAD discard: The operation cannot be performed because one or more virtual disks still exist, the disk group still may be recovering its capacity, or this is the last disk group that exists.</p> <p>Case 6: LDAD resolve condition: The disk group contains a disk volume that is in a data-lost state. This condition cannot be resolved.</p> <p>Case 7: Physical Store erase volume: The disk is a part of a disk group and cannot be erased.</p> <p>Case 8: Storage cell discard: The storage cell contains one or more virtual disks or LUN presentations.</p> <p>Case 9: Storage cell client discard: = The EVA host contains one or more LUN presentations.</p> <p>Case 10: SCVD discard: The virtual disk contains one or more derived units and cannot be discarded. This may indicate a programming error.</p> <p>Case 11: SCVD set capacity: The capacity cannot be modified because the virtual disk has a dependency on either a snapshot or snapclone.</p> <p>Case 12: SCVD set disk cache policy: The virtual disk cache policy cannot be modified while the virtual disk is presented and enabled.</p> <p>Case 13: SCVD set logical disk: The logical disk attribute is already set, or the supplied logical disk is already a member of another virtual disk.</p> <p>Case 14: VOLUME set requested usage: The disk volume is already a member of a disk group or is in the state of being removed from a disk group.</p> <p>Case 15: GROUP discard: The Continuous Access group cannot be discarded as one or more virtual disk members exist.</p>	<p>Case 4: Resolve the delay before performing the operation.</p> <p>Case 5: Delete any remaining virtual disks or wait for the used capacity to reach zero before the disk group can be deleted. If this is the last remaining disk group, uninitialize the storage cell to remove it.</p> <p>Case 6: Report the error to product support.</p> <p>Case 7: The disk must be in a reserved state before it can be erased.</p> <p>Case 8: Delete the virtual disks or LUN presentations before uninitializing the storage cell.</p> <p>Case 9: Delete the LUN presentations before deleting the EVA host.</p> <p>Case 10: Report the error to product support.</p> <p>Case 11: Resolve the situation before attempting the operation again.</p> <p>Case 12: Resolve the situation before attempting the operation again.</p> <p>Case 13: This may indicate a programming error. Report the error to product support.</p> <p>Case 14: Select another disk or remove the disk from the disk group before making it a member of a different disk group.</p> <p>Case 15: Remove the virtual disks from the group and retry the operation.</p>
26 Parameter Object Does Not Exist	<p>The operation cannot be performed because the object does not exist. This can indicate a user or program error.</p> <p>VOLUME set requested usage: The disk volume set requested usage cannot be performed because the disk group does not exist. This can indicate a user or program error.</p>	Report the error to product support.
27 Target Object Does Not Exist	<p>Case 1: The operation cannot be performed because the object does not exist. This can indicate a user or program error.</p> <p>Case 2: DERIVED UNIT discard: The operation cannot be performed because the virtual disk, snapshot, or snapclone does not exist or is still being created.</p> <p>Case 3: VOLUME set requested usage: The operation cannot be performed because the target disk volume does not exist. This can indicate a user or program error.</p>	<p>Case 1: Report the error to product support.</p> <p>Case 2: Retry the request at a later time.</p> <p>Case 3: Report the error to product support.</p> <p>Case 4: Report the error to product support.</p>

Table 20 Error messages *(continued)*

Status code value	Meaning	How to correct
	Case 4: GROUP get name: The operation cannot be performed because the Continuous Access group does not exist. This can indicate a user or program error.	
28 Timeout	A timeout has occurred in processing the request.	Verify the hardware connections and that communication to the device is successful.
29 Unknown ID	The supplied storage cell identifier is invalid. This can indicate a user or program error.	Report the error to product support.
30 Unknown Parameter Handle	The supplied parameter handle is unknown. This can indicate a user or program error.	Report the error to product support.
31 Unrecoverable Media Error	The operation could not be completed because one or more of the disk media had an unrecoverable error.	Report the error to product support.
32 Invalid State	This error is no longer supported.	Report the error to product support.
33 Transport Error	A SCMI transport error has occurred.	Verify the hardware connections, communication to the device, and that the management software is operating successfully.
34 Volume is Missing	The operation could not be completed because the drive volume is in a missing state.	Resolve the condition and retry the request. Report the error to product support.
35 Invalid Cursor	The supplied cursor or sequence number is invalid. This may indicate a user or program error.	Report the error to product support.
36 Invalid Target for the Operation	The specified target logical disk already has an existing data sharing relationship. This can indicate a user or program error.	Report the error to product support.
37 No More Events	There are no more events to retrieve. (This message is informational only.)	No action required.
38 Lock Busy	The command lock is busy and being held by another process.	Retry the request at a later time.
39 Time Not Set	The storage system time is not set. The storage system time is set automatically by the management software.	Report the error to product support.
40 Not a Supported Version	The requested operation is not supported by this firmware version. This can indicate a user or program error.	Report the error to product support.
41 No Logical Disk for Vdisk	The specified SCVD does not have a logical disk associated with it. This can indicate a user or program error.	Report the error to product support.
42 Logical disk Presented	The virtual disk specified is already presented to the client and the requested operation is not allowed.	Delete the associated presentation(s) and retry the request.
43 Operation Denied On Slave	The request is not allowed on the slave controller. This can indicate a user or program error.	Report the error to product support.
44 Not licensed for data replication	This error is no longer supported.	Report the error to product support.
45 Not DR group member	The operation cannot be performed because the virtual disk is not a member of a Continuous Access group.	Configure the virtual disk to be a member of a Continuous Access group and retry the request.

Table 20 Error messages *(continued)*

Status code value	Meaning	How to correct
46 Invalid DR mode	The operation cannot be performed because the Continuous Access group is not in the required mode.	Configure the Continuous Access group correctly and retry the request.
47 The target DR member is in full copy, operation rejected	The operation cannot be performed because at least one of the virtual disk members is in a copying state.	Wait for the copying state to complete and retry the request.
48 Security credentials needed. Please update your system's ID and password in the Storage System Access menu.	The management software is unable to log in to the storage system. The storage system password has been configured.	Use the management software to save the password specified so communication can proceed.
49 Security credentials supplied were invalid. Please update your system's ID and password in the Storage System Access menu.	The management software is unable to login to the device. The storage system password may have been re-configured or removed.	Use the management software to set the password to match the device so communication can proceed.
50 Security credentials supplied were invalid. Please update your system's ID and password in the Storage System Access menu.	The management software is already logged in to the device. (This message is informational only.)	No action required.
51 Storage system connection down	The Continuous Access group is not functioning.	Verify that devices are powered on and that device hardware connections are functioning correctly.
52 DR group empty	No virtual disks are members of the Continuous Access group.	Add one or more virtual disks as members and retry the request.
53 Incompatible attribute	The request cannot be performed because one or more of the attributes specified is incompatible.	Retry the request with valid attributes for the operation.
54 Vdisk is a DR group member	The requested operation cannot be performed on a virtual disk that is already a member of a data replication group.	Remove the virtual disk as a member of a data replication group and retry the request.
55 Vdisk is a DR log unit	The requested operation cannot be performed on a virtual disk that is a log unit.	No action required.
56 Cache batteries failed or missing.	The battery system is missing or discharged.	Report the error to product support.
57 Vdisk is not presented	The virtual disk member is not presented to a client.	The virtual disk member must be presented to a client before this operation can be performed.
58 Other controller failed	Invalid status for logical disk. This error is no longer supported.	Report the error to product support.
59 Maximum Number of Objects Exceeded.	Case 1: The maximum number of items allowed has been reached. Case 2: The maximum number of EVA hosts has been reached. Case 3: The maximum number of port WWNs has been reached.	Case 1: If this operation is still desired, delete one or more of the items and retry the operation. Case 2: If this operation is still desired, delete one or more of the EVA hosts and retry the operation.

Table 20 Error messages *(continued)*

Status code value	Meaning	How to correct
		Case 3: If this operation is still desired, delete one or more of the port WWNs and retry the operation.
60 Max size exceeded	<p>Case 1: The maximum number of items already exist on the destination storage cell.</p> <p>Case 2: The size specified exceeds the maximum size allowed.</p> <p>Case 3: The presented user space exceeds the maximum size allowed.</p> <p>Case 4: The presented user space exceeds the maximum size allowed.</p> <p>Case 5: The size specified exceeds the maximum size allowed.</p> <p>Case 6: The maximum number of EVA hosts already exist on the destination storage cell.</p> <p>Case 7: The maximum number of EVA hosts already exist on the destination storage cell.</p> <p>Case 8: The maximum number of Continuous Access groups already exist.</p>	<p>Case 1: If this operation is still desired, delete one or more of the items on the destination storage cell and retry the operation.</p> <p>Case 2: Use a smaller size and retry the operation.</p> <p>Case 3: No action required.</p> <p>Case 4: No action required.</p> <p>Case 5: Use a smaller size and retry the operation.</p> <p>Case 6: If this operation is still desired, delete one or more of the EVA hosts and retry the operation.</p> <p>Case 7: If this operation is still desired, delete one or more of the virtual disks on the destination storage cell and retry the operation.</p> <p>Case 8: If this operation is still desired, delete one or more of the groups and retry the operation.</p>
61 Password mismatch. Please update your system's password in the Storage System Access menu. Continued attempts to access this storage system with an incorrect password will disable management of this storage system.	The login password entered on the controllers does not match.	Reconfigure one of the storage system controller passwords, then use the management software to set the password to match the device so communication can proceed.
62 DR group is merging	The operation cannot be performed because the Continuous Access connection is currently merging.	Wait for the merge operation to complete and retry the request.
63 DR group is logging	The operation cannot be performed because the Continuous Access connection is currently logging.	Wait for the logging operation to complete and retry the request.
64 Connection is suspended	The operation cannot be performed because the Continuous Access connection is currently suspended	Resolve the suspended mode and retry the request.
65 Bad image header	The firmware image file has a header checksum error.	Retrieve a valid firmware image file and retry the request.
66 Bad image	The firmware image file has a checksum error.	Retrieve a valid firmware image file and retry the request.
67 The firmware image file is too large. Image too large	Invalid status for logical disk. This error is no longer supported.	Retrieve a valid firmware image file and retry the request.
70 Image incompatible with system configuration. Version conflict in upgrade or downgrade not allowed.	The firmware image file is incompatible with the current firmware.	Retrieve a valid firmware image file and retry the request

Table 20 Error messages (continued)

Status code value	Meaning	How to correct
71 Bad image segment	The firmware image download process has failed because of a corrupted image segment.	Verify that the firmware image is not corrupted and retry the firmware download process.
72 Image already loaded	The firmware version already exists on the device.	No action required.
73 Image Write Error	The firmware image download process has failed because of a failed write operation.	Verify that the firmware image is not corrupted and retry the firmware download process.
74 Logical Disk Sharing	<p>Case 1: The operation cannot be performed because the virtual disk or snapshot is part of a snapshot group.</p> <p>Case 2: The operation may be prevented because a snapclone or snapshot operation is in progress. If a snapclone operation is in progress, the parent virtual disk should be discarded automatically after the operation completes. If the parent virtual disk has snapshots, then you must delete the snapshots before the parent virtual disk can be deleted.</p> <p>Case 3: The operation cannot be performed because either the previous snapclone operation is still in progress, or the virtual disk is already part of a snapshot group.</p> <p>Case 4: A capacity change is not allowed on a virtual disk or snapshot that is a part of a snapshot group.</p> <p>Case 5: The operation cannot be performed because the virtual disk or snapshot is a part of a snapshot group.</p>	<p>Case 1: No action required.</p> <p>Case 2: No action required.</p> <p>Case 3: If a snapclone operation is in progress, wait until the snapclone operation has completed and retry the operation. Otherwise, the operation cannot be performed on this virtual disk.</p> <p>Case 4: No action required.</p> <p>Case 5: No action required.</p>
75 Bad Image Size	The firmware image file is not the correct size.	Retrieve a valid firmware image file and retry the request.
76 The controller is temporarily busy and it cannot process the request. Retry the request later.	The controller is currently processing a firmware download. Retry the request once the firmware download process is complete.	Retry the request once the firmware download process is complete.
77 Volume Failure Predicted	The disk volume specified is in a predictive failed state.	Report the error to product support.
78 Invalid object condition for this command.	The current condition or state is preventing the request from completing successfully.	Resolve the condition and retry the request.
79 Snapshot (or snapclone) deletion in progress. The requested operation is currently not allowed. Please try again later.	The current condition of the snapshot, snapclone or parent virtual disk is preventing the request from completing successfully.	Wait for the operation to complete and retry the request.
80 Invalid Volume Usage	Case 1: The disk volume is already a part of a disk group.	Resolve the condition by setting the usage to a reserved state and 80 retry the request. Invalid Volume Usage
	Case 2: The disk volume usage cannot be modified, as the minimum number of disks exist in the disk group.	Report the error to product support.
81 Minimum Volumes In Disk Group	The disk volume usage cannot be modified, as the minimum number of disks exist in the disk group.	Resolve the condition by adding additional disks and retry the request.

Table 20 Error messages *(continued)*

Status code value	Meaning	How to correct
82 Shutdown In Progress	The controller is currently shutting down.	No action required.
83 Controller API Not Ready, Try Again Later	The device is not ready to process the request.	Retry the request at a later time.
84 Is Snapshot	This is a snapshot virtual disk and cannot be a member of a Continuous Access group.	No action required.
85 Cannot add or remove DR group member. Mirror cache must be active for this Vdisk. Check controller cache condition.	An incompatible mirror policy of the virtual disk is preventing it from becoming a member of a Continuous Access group.	Modify the mirror policy and retry the request.
86 Command View EVA has detected this array as inoperative. Contact HP Service for assistance.	Case 1: A virtual disk is in an inoperative state and the request cannot be processed.	Report the error to product support.
	Case 2: The snapclone cannot be associated with a virtual disk that is in an inoperative state.	
	Case 3: The snapshot cannot be associated with a virtual disk that is in an inoperative state.	
87 Disk group inoperative or disks in group less than minimum.	The disk group is in an inoperative state and cannot process the request.	Report the error to product support.
88 Storage system inoperative	The storage system is inoperative and cannot process the request.	Report the error to product support.
89 Failsafe Locked	The request cannot be performed because the Continuous Access group is in a failsafe locked state.	Resolve the condition and retry the request.
90 Data Flush Incomplete	The disk cache data need to be flushed before the condition can be resolved.	Retry the request later.
91 Redundancy Mirrored Inoperative	The disk group is in a redundancy mirrored inoperative state and the request cannot be completed.	Report the error to product support.
92 Duplicate LUN	The LUN number is already in use by another client of the storage system.	Select another LUN number and retry the request.
93 Other remote controller failed	While the request was being performed, the remote storage system controller failed.	Resolve the condition and retry the request. Report the error to product support.
94 Unknown remote Vdisk	The remote storage system specified does not exist.	Correctly select the remote storage system and retry the request.
95 Unknown remote DR group	The remote Continuous Access group specified does not exist.	Correctly select the remote Continuous Access group and retry the request.
96 PLDM failed	The disk metadata was unable to be updated.	Resolve the condition and retry the request. Report the error to product support.
97 Storage system could not be locked. System busy. Try command again.	Another process has already taken the SCMI lock on the storage system.	Retry the request later.

Table 20 Error messages *(continued)*

Status code value	Meaning	How to correct
98 Error on remote storage system.	While the request was being performed, an error occurred on the remote storage system.	Resolve the condition and retry the request
99 The DR operation can only be completed when the source-destination connection is down. If you are doing a destination DR deletion, make sure the connection link to the source DR system is down or do a failover operation to make this system the source.	The request failed because the operation cannot be performed on a Continuous Access connection that is up.	Resolve the condition and retry the request.
100 Login required - password changed.	The management software is unable to log into the device as the password has changed.	The storage system password may have been re-configured or removed. The management software must be used to set the password up to match the device so communication can proceed.

7 Support and other resources

Contacting HP

HP technical support

For world wide 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/wwalerts>

After registering, you will receive email 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

For related documents, see the Manuals page of the HP Business Support Center website:

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

Click **Disk Storage Systems** under Storage, and then click **HP 4400 Enterprise Virtual Array** under P6000/EVA Disk Arrays. For software documentation, click **Storage Software** under Storage, and then click **HP P6000 Command View Software** under Storage Device Management Software.

You can find *HP 4400 Enterprise Virtual Array QuickSpecs* from the main product page website.

For example, to find the QuickSpecs if you are in the U.S., use the following website:

<http://welcome.hp.com/country/us/en/support.html>

Websites

- HP:
<http://www.hp.com>
- HP Storage:
<http://www.hp.com/go/storage>

- HP Partner Locator:
http://www.hp.com/service_locator
- HP Software Downloads:
<http://www.hp.com/support/downloads>
- HP Software Depot:
<http://www.software.hp.com>
- HP Single Point of Connectivity Knowledge (SPOCK):
<http://www.hp.com/storage/spock>
- HP SAN manuals:
<http://www.hp.com/go/sdgmanuals>
- HP Support Center
<http://h20566.www2.hp.com/portal/site/hpsc/public/>

Typographic conventions

Table 21 Document conventions

Convention	Element
Blue text: Table 21 (page 108)	Cross-reference links
Blue, underlined text: http://www.hp.com	Website addresses
Bold text	<ul style="list-style-type: none">• 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
<i>Italic</i> text	Text emphasis
Monospace text	<ul style="list-style-type: none">• File and directory names• System output• Code• Commands, their arguments, and argument values
<i>Monospace, italic</i> text	<ul style="list-style-type: none">• Code variables• Command variables
Monospace, bold text	Emphasized monospace text
WARNING!	An alert that calls attention to important information that if not understood or followed can result in personal injury.
CAUTION:	An alert that calls attention to important information that if not understood or followed can result in data loss, data corruption, or damage to hardware or software.
IMPORTANT:	An alert that calls attention to essential information.
NOTE:	An alert that calls attention to additional or supplementary information.
TIP:	An alert that calls attention to helpful hints and shortcuts.

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 or see the CSR website:

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

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.
-

A Regulatory compliance notices

Regulatory compliance identification numbers

For the purpose of regulatory compliance certifications and identification, this product has been assigned a unique regulatory model number. The regulatory model number can be found on the product nameplate label, along with all required approval markings and information. When requesting compliance information for this product, always refer to this regulatory model number. The regulatory model number is not the marketing name or model number of the product.

Product specific information:

HP _____

Regulatory model number: _____

FCC and CISPR classification: _____

These products contain laser components. See Class 1 laser statement in the [“Laser compliance notices”](#) (page 114) section.

Federal Communications Commission 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.

FCC rating label

The FCC rating label on the device shows the classification (A or B) of the equipment. Class B devices have an FCC logo or ID on the label. Class A devices do not have an FCC logo or ID on the label. After you determine the class of the device, refer to the corresponding statement.

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 this FCC declaration, contact us by mail or telephone:

- Hewlett-Packard Company P.O. Box 692000, Mail Stop 510101 Houston, Texas 77269-2000
- Or call 1-281-514-3333

Modification

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

When provided, 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.

Canadian 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 class 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 class B respecte toutes les exigences du Règlement sur le matériel brouilleur du Canada.


European Union notice

This product complies with the following EU directives:

- Low Voltage Directive 2006/95/EC
- EMC Directive 2004/108/EC

Compliance with these directives implies conformity to applicable harmonized European standards (European Norms) which are listed on the EU Declaration of Conformity issued by Hewlett-Packard for this product or product family.

This compliance is indicated by the following conformity marking placed on the product:



Certificates can be obtained from <http://www.hp.com/go/certificates>.

Hewlett-Packard GmbH, HQ-TRE, Herrenberger Strasse 140, 71034 Boeblingen, Germany

Japanese notices

Japanese VCCI-A notice

この装置は、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。 VCCI-A

Japanese VCCI-B notice

この装置は、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を目的としています。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-B

Japanese VCCI marking



Japanese power cord statement

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。

Please use the attached power cord.
The attached power cord is not allowed to use with other product.

Korean notices

Class A equipment

A급 기기 (업무용 정보통신기기)

이 기기는 업무용으로 전자파적합등록을 한 기기이오니 판매자 또는 사용자는 이 점을 주의하시기 바라며, 만약 잘못판매 또는 구입하였을 때에는 가정용으로 교환하시기 바랍니다.

Class B equipment

B급 기기 (가정용 정보통신기기)

이 기기는 가정용으로 전자파적합등록을 한 기기로서
주거지역에서는 물론 모든지역에서 사용할 수 있습니다.


Taiwanese notices

BSMI Class A 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。

Taiwan battery recycle statement




Turkish recycling notice



Türkiye Cumhuriyeti: EEE Yönetmeliğine Uygundur

Vietnamese Information Technology and Communications compliance marking



Laser compliance notices

English laser notice

This device may contain a laser that is classified as a Class 1 Laser Product in accordance with U.S. FDA regulations and the IEC 60825-1. The product does not emit hazardous laser radiation.



WARNING! Use of controls or adjustments or performance of procedures other than those specified herein or in the laser product's installation guide may result in hazardous radiation exposure. To reduce the risk of exposure to hazardous radiation:

- Do not try to open the module 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 unit.

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.

Dutch laser notice



WAARSCHUWING: dit apparaat bevat mogelijk een laser die is geclassificeerd als een laserproduct van Klasse 1 overeenkomstig de bepalingen van de Amerikaanse FDA en de richtlijn IEC 60825-1. Dit product geeft geen gevaarlijke laserstraling af.

Als u bedieningselementen gebruikt, instellingen aanpast of procedures uitvoert op een andere manier dan in deze publicatie of in de installatiehandleiding van het laserproduct wordt aangegeven, loopt u het risico te worden blootgesteld aan gevaarlijke straling. Het risico van blootstelling aan gevaarlijke straling beperkt u als volgt:

- Probeer de behuizing van de module niet te openen. U mag zelf geen onderdelen repareren.
- Gebruik voor de laserapparatuur geen andere knoppen of instellingen en voer geen andere aanpassingen of procedures uit dan die in deze handleiding worden beschreven.
- Alleen door HP geautoriseerde technici mogen het apparaat repareren.

French laser notice



AVERTISSEMENT : cet appareil peut être équipé d'un laser classé en tant que Produit laser de classe 1 et conforme à la réglementation de la FDA américaine et à la norme 60825-1 de l'IEC. Ce produit n'émet pas de rayonnement dangereux.

L'utilisation de commandes, de réglages ou de procédures autres que ceux qui sont indiqués ici ou dans le manuel d'installation du produit laser peut exposer l'utilisateur à des rayonnements dangereux. Pour réduire le risque d'exposition à des rayonnements dangereux :

- Ne tentez pas d'ouvrir le boîtier renfermant l'appareil laser. Il ne contient aucune pièce dont la maintenance puisse être effectuée par l'utilisateur.
- Tout contrôle, réglage ou procédure autre que ceux décrits dans ce chapitre ne doivent pas être effectués par l'utilisateur.
- Seuls les Mainteneurs Agréés HP sont habilités à réparer l'appareil laser.

German laser notice



VORSICHT: Dieses Gerät enthält möglicherweise einen Laser, der nach den US-amerikanischen FDA-Bestimmungen und nach IEC 60825-1 als Laserprodukt der Klasse 1 zertifiziert ist. Gesundheitsschädliche Laserstrahlen werden nicht emittiert.

Die Anleitungen in diesem Dokument müssen befolgt werden. Bei Einstellungen oder Durchführung sonstiger Verfahren, die über die Anleitungen in diesem Dokument bzw. im Installationshandbuch des Lasergeräts hinausgehen, kann es zum Austritt gefährlicher Strahlung kommen. Zur Vermeidung der Freisetzung gefährlicher Strahlungen sind die folgenden Punkte zu beachten:

- Versuchen Sie nicht, die Abdeckung des Lasermoduls zu öffnen. Im Inneren befinden sich keine Komponenten, die vom Benutzer gewartet werden können.
 - Benutzen Sie das Lasergerät ausschließlich gemäß den Anleitungen und Hinweisen in diesem Dokument.
 - Lassen Sie das Gerät nur von einem HP Servicepartner reparieren.
-

Italian laser notice



AVVERTENZA: AVVERTENZA. Questo dispositivo può contenere un laser classificato come prodotto laser di Classe 1 in conformità alle normative US FDA e IEC 60825-1. Questo prodotto non emette radiazioni laser pericolose.

L'eventuale esecuzione di comandi, regolazioni o procedure difformi a quanto specificato nella presente documentazione o nella guida di installazione del prodotto può causare l'esposizione a radiazioni nocive. Per ridurre i rischi di esposizione a radiazioni pericolose, attenersi alle seguenti precauzioni:

- Non cercare di aprire il contenitore del modulo. All'interno non vi sono componenti soggetti a manutenzione da parte dell'utente.
 - Non eseguire operazioni di controllo, regolazione o di altro genere su un dispositivo laser ad eccezione di quelle specificate da queste istruzioni.
 - Affidare gli interventi di riparazione dell'unità esclusivamente ai tecnici dell'Assistenza autorizzata HP.
-

Japanese laser notice



警告: 本製品には、US FDA規則およびIEC 60825-1に基づくClass 1レーザー製品が含まれている場合があります。本製品は人体に危険なレーザー光は発しません。

本書およびレーザー製品のインストールガイドに示されている以外の方法で制御、調整、使用した場合、人体に危険な光線にさらされる場合があります。人体に危険な光線にさらされないため、以下の項目を守ってください。

- モジュール エンクロージャを開けないでください。ユーザーが取り扱えるコンポーネントは含まれていません。
- 本書に示されている以外の方法で、レーザー デバイスを制御、調整、使用しないでください。
- HPの正規サービス技術者のみが本ユニットの修理を許可されています。

Spanish laser notice



ADVERTENCIA: Este dispositivo podría contener un láser clasificado como producto de láser de Clase 1 de acuerdo con la normativa de la FDA de EE.UU. e IEC 60825-1. El producto no emite radiaciones láser peligrosas.

El uso de controles, ajustes o manipulaciones distintos de los especificados aquí o en la guía de instalación del producto de láser puede producir una exposición peligrosa a las radiaciones. Para evitar el riesgo de exposición a radiaciones peligrosas:

- No intente abrir la cubierta del módulo. Dentro no hay componentes que el usuario pueda reparar.
- No realice más operaciones de control, ajustes o manipulaciones en el dispositivo láser que los aquí especificados.
- Sólo permita reparar la unidad a los agentes del servicio técnico autorizado HP.

Recycling notices

English recycling notice

Disposal of waste equipment by users in private household in the European Union



This symbol means do not dispose of your product with your other household waste. Instead, you should protect human health and the environment by handing over your waste equipment to a designated collection point for the recycling of waste electrical and electronic equipment. For more information, please contact your household waste disposal service

Bulgarian recycling notice

Изхвърляне на отпадъчно оборудване от потребители в частни домакинства в Европейския съюз



Този символ върху продукта или опаковката му показва, че продуктът не трябва да се изхвърля заедно с другите битови отпадъци. Вместо това, трябва да предпазите човешкото здраве и околната среда, като предадете отпадъчното оборудване в предназначен за събирането му пункт за рециклиране на неизползваемо електрическо и електронно борудване. За допълнителна информация се свържете с фирмата по чистота, чиито услуги използвате.

Czech recycling notice

Likvidace zařízení v domácnostech v Evropské unii



Tento symbol znamená, že nesmíte tento produkt likvidovat spolu s jiným domovním odpadem. Místo toho byste měli chránit lidské zdraví a životní prostředí tím, že jej předáte na k tomu určené sběrné pracoviště, kde se zabývají recyklací elektrického a elektronického vybavení. Pro více informací kontaktujte společnost zabývající se sběrem a svozem domovního odpadu.

Danish recycling notice

Bortskaffelse af brugt udstyr hos brugere i private hjem i EU



Dette symbol betyder, at produktet ikke må bortskaffes sammen med andet husholdningsaffald. Du skal i stedet den menneskelige sundhed og miljøet ved at afl evere dit brugte udstyr på et dertil beregnet indsamlingssted for af brugt, elektrisk og elektronisk udstyr. Kontakt nærmeste renovationsafdeling for yderligere oplysninger.

Dutch recycling notice

Inzameling van afgedankte apparatuur van particuliere huishoudens in de Europese Unie



Dit symbool betekent dat het product niet mag worden gedeponeerd bij het overige huishoudelijke afval. Bescherm de gezondheid en het milieu door afgedankte apparatuur in te leveren bij een hiervoor bestemd inzamelpunt voor recycling van afgedankte elektrische en elektronische apparatuur. Neem voor meer informatie contact op met uw gemeentereinigingsdienst.

Estonian recycling notice

Äravisatavate seadmete likvideerimine Euroopa Liidu eramajapidamistes



See märk näitab, et seadet ei tohi visata olmeprügi hulka. Inimeste tervise ja keskkonna säästmise nimel tuleb äravisatav toode tuua elektriliste ja elektrooniliste seadmete käitlemisega egelevasse kogumispunkti. Küsimuste korral pöörduge kohaliku prügikäitlusettevõtte poole.

Finnish recycling notice

Kotitalousjätteiden hävittäminen Euroopan unionin alueella



Tämä symboli merkitsee, että laitetta ei saa hävittää muiden kotitalousjätteiden mukana. Sen sijaan sinun on suojattava ihmisten terveyttä ja ympäristöä toimittamalla käytöstä poistettu laite sähkö- tai elektroniikkajätteen kierrätyspisteeseen. Lisätietoja saat jätehuoltoyhtiöltä.

French recycling notice

Mise au rebut d'équipement par les utilisateurs privés dans l'Union Européenne



Ce symbole indique que vous ne devez pas jeter votre produit avec les ordures ménagères. Il est de votre responsabilité de protéger la santé et l'environnement et de vous débarrasser de votre équipement en le remettant à une déchetterie effectuant le recyclage des équipements électriques et électroniques. Pour de plus amples informations, prenez contact avec votre service d'élimination des ordures ménagères.

German recycling notice

Entsorgung von Altgeräten von Benutzern in privaten Haushalten in der EU



Dieses Symbol besagt, dass dieses Produkt nicht mit dem Haushaltsmüll entsorgt werden darf. Zum Schutze der Gesundheit und der Umwelt sollten Sie stattdessen Ihre Altgeräte zur Entsorgung einer dafür vorgesehenen Recyclingstelle für elektrische und elektronische Geräte übergeben. Weitere Informationen erhalten Sie von Ihrem Entsorgungsunternehmen für Hausmüll.

Greek recycling notice

Απορριψη άχρηστου εξοπλισμού από ιδιώτες χρήστες στην Ευρωπαϊκή Ένωση



Αυτό το σύμβολο σημαίνει ότι δεν πρέπει να απορρίψετε το προϊόν με τα λοιπά οικιακά απορρίμματα. Αντίθετα, πρέπει να προστατέψετε την ανθρώπινη υγεία και το περιβάλλον παραδίδοντας τον άχρηστο εξοπλισμό σας σε εξουσιοδοτημένο σημείο συλλογής για την ανακύκλωση άχρηστου ηλεκτρικού και ηλεκτρονικού εξοπλισμού. Για περισσότερες πληροφορίες, επικοινωνήστε με την υπηρεσία απόρριψης απορριμμάτων της περιοχής σας.

Hungarian recycling notice

A hulladék anyagok megsemmisítése az Európai Unió háztartásaiban



Ez a szimbólum azt jelzi, hogy a készüléket nem szabad a háztartási hulladékkal együtt kidobni. Ehelyett a leselejtezett berendezéseknek az elektromos vagy elektronikus hulladék átvételére kijelölt helyen történő beszállításával megóvja az emberi egészséget és a környezetet. További információt a helyi köztisztasági vállalatától kaphat.



Italian recycling notice

Smaltimento di apparecchiature usate da parte di utenti privati nell'Unione Europea



Questo simbolo avvisa di non smaltire il prodotto con i normali rifiuti domestici. Rispettare la salute umana e l'ambiente conferendo l'apparecchiatura dismessa a un centro di raccolta designato per il riciclo di apparecchiature elettroniche ed elettriche. Per ulteriori informazioni, rivolgersi al servizio per lo smaltimento dei rifiuti domestici.



Latvian recycling notice

Europos Sąjungos namų ūkio vartotojų įrangos atliekų šalinimas



Šis simbolis nurodo, kad gaminio negalima išmesti kartu su kitomis buitinėmis atliekomis. Kad apsaugotumėte žmonių sveikatą ir aplinką, pasenusią nenaudojamą įrangą turite nuvežti į elektrinių ir elektroninių atliekų surinkimo punktą. Daugiau informacijos teiraukitės buitinių atliekų surinkimo tarnybos.



Lithuanian recycling notice

Nolietotu iekārtu iznīcināšanas noteikumi lietotājiem Eiropas Savienības privātajās mājāsaimniecībās



Šis simbols norāda, ka ierīci nedrīkst utilizēt kopā ar citiem mājāsaimniecības atkritumiem. Jums jā rūpējas par cilvēku veselības un vides aizsardzību, nododot lietoto aprīkojumu otrreizējai pārstrādei īpašā lietotu elektrisko un elektronisko ierīču savākšanas punktā. Lai iegūtu plašāku informāciju, lūdzu, sazinieties ar savu mājāsaimniecības atkritumu likvidēšanas dienestu.



Polish recycling notice

Utylizacja zużytego sprzętu przez użytkowników w prywatnych gospodarstwach domowych w krajach Unii Europejskiej



Ten symbol oznacza, że nie wolno wyrzucać produktu wraz z innymi domowymi odpadkami. Obowiązkiem użytkownika jest ochrona zdrowia ludzkiego i środowiska przez przekazanie zużytego sprzętu do wyznaczonego punktu zajmującego się recyklingiem odpadów powstających ze sprzętu elektrycznego i elektronicznego. Więcej informacji można uzyskać od lokalnej firmy zajmującej wywozem nieczystości.



Portuguese recycling notice

Descarte de equipamentos usados por utilizadores domésticos na União Europeia



Este símbolo indica que não deve descartar o seu produto juntamente com os outros lixos domiciliários. Ao invés disso, deve proteger a saúde humana e o meio ambiente levando o seu equipamento para descarte em um ponto de recolha destinado à reciclagem de resíduos de equipamentos eléctricos e electrónicos. Para obter mais informações, contacte o seu serviço de tratamento de resíduos domésticos.

Romanian recycling notice

Casarea echipamentului uzat de către utilizatorii casnici din Uniunea Europeană



Acest simbol înseamnă să nu se arunce produsul cu alte deșeuri menajere. În schimb, trebuie să protejați sănătatea umană și mediul predând echipamentul uzat la un punct de colectare desemnat pentru reciclarea echipamentelor electrice și electronice uzate. Pentru informații suplimentare, vă rugăm să contactați serviciul de eliminare a deșeurilor menajere local.

Slovak recycling notice

Likvidácia vyradených zariadení používateľmi v domácnostiach v Európskej únii



Tento symbol znamená, že tento produkt sa nemá likvidovať s ostatným domovým odpadom. Namiesto toho by ste mali chrániť ľudské zdravie a životné prostredie odovzdaním odpadového zariadenia na zbernom mieste, ktoré je určené na recykláciu odpadových elektrických a elektronických zariadení. Ďalšie informácie získate od spoločnosti zaoberajúcej sa likvidáciou domového odpadu.

Spanish recycling notice

Eliminación de los equipos que ya no se utilizan en entornos domésticos de la Unión Europea



Este símbolo indica que este producto no debe eliminarse con los residuos domésticos. En lugar de ello, debe evitar causar daños a la salud de las personas y al medio ambiente llevando los equipos que no utilice a un punto de recogida designado para el reciclaje de equipos eléctricos y electrónicos que ya no se utilizan. Para obtener más información, póngase en contacto con el servicio de recogida de residuos domésticos.

Swedish recycling notice

Hantering av elektroniskt avfall för hemanvändare inom EU



Den här symbolen innebär att du inte ska kasta din produkt i hushållsavfallet. Värna i stället om natur och miljö genom att lämna in uttjänt utrustning på anvisad samlingsplats. Allt elektriskt och elektroniskt avfall går sedan vidare till återvinning. Kontakta ditt återvinningsföretag för mer information.

Battery replacement notices

Dutch battery notice

Verklaring betreffende de batterij



WAARSCHUWING: dit apparaat bevat mogelijk een batterij.

- Probeer de batterijen na het verwijderen niet op te laden.
- Stel de batterijen niet bloot aan water of temperaturen boven 60° C.
- De batterijen mogen niet worden beschadigd, gedemonteerd, geplet of doorboord.
- Zorg dat u geen kortsluiting veroorzaakt tussen de externe contactpunten en laat de batterijen niet in aanraking komen met water of vuur.
- Gebruik ter vervanging alleen door HP goedgekeurde batterijen.

Batterijen, accu's en accumulators mogen niet worden gedeponeerd bij het normale huishoudelijke afval. Als u de batterijen/accu's wilt inleveren voor hergebruik of op de juiste manier wilt vernietigen, kunt u gebruik maken van het openbare inzamelingssysteem voor klein chemisch afval of ze terugsturen naar HP of een geautoriseerde HP Business of Service Partner.

Neem contact op met een geautoriseerde leverancier of een Business of Service Partner voor meer informatie over het vervangen of op de juiste manier vernietigen van accu's.

French battery notice

Avis relatif aux piles



AVERTISSEMENT : cet appareil peut contenir des piles.

- N'essayez pas de recharger les piles après les avoir retirées.
- Évitez de les mettre en contact avec de l'eau ou de les soumettre à des températures supérieures à 60°C.
- N'essayez pas de démonter, d'écraser ou de percer les piles.
- N'essayez pas de court-circuiter les bornes de la pile ou de jeter cette dernière dans le feu ou l'eau.
- Remplacez les piles exclusivement par des pièces de rechange HP prévues pour ce produit.

Les piles, modules de batteries et accumulateurs ne doivent pas être jetés avec les déchets ménagers. Pour permettre leur recyclage ou leur élimination, veuillez utiliser les systèmes de collecte publique ou renvoyez-les à HP, à votre Partenaire Agréé HP ou aux agents agréés.

Contactez un Revendeur Agréé ou Mainteneur Agréé pour savoir comment remplacer et jeter vos piles.

Hinweise zu Batterien und Akkus



VORSICHT: Dieses Produkt enthält unter Umständen eine Batterie oder einen Akku.

- Versuchen Sie nicht, Batterien und Akkus außerhalb des Gerätes wieder aufzuladen.
- Schützen Sie Batterien und Akkus vor Feuchtigkeit und Temperaturen über 60°.
- Verwenden Sie Batterien und Akkus nicht missbräuchlich, nehmen Sie sie nicht auseinander und vermeiden Sie mechanische Beschädigungen jeglicher Art.
- Vermeiden Sie Kurzschlüsse, und setzen Sie Batterien und Akkus weder Wasser noch Feuer aus.
- Ersetzen Sie Batterien und Akkus nur durch die von HP vorgesehenen Ersatzteile.

Batterien und Akkus dürfen nicht über den normalen Hausmüll entsorgt werden. Um sie der Wiederverwertung oder dem Sondermüll zuzuführen, nutzen Sie die öffentlichen Sammelstellen, oder setzen Sie sich bezüglich der Entsorgung mit einem HP Partner in Verbindung.

Weitere Informationen zum Austausch von Batterien und Akkus oder zur sachgemäßen Entsorgung erhalten Sie bei Ihrem HP Partner oder Servicepartner.

Istruzioni per la batteria



AVVERTENZA: Questo dispositivo può contenere una batteria.

- Non tentare di ricaricare le batterie se rimosse.
- Evitare che le batterie entrino in contatto con l'acqua o siano esposte a temperature superiori a 60° C.
- Non smontare, schiacciare, forare o utilizzare in modo improprio la batteria.
- Non accorciare i contatti esterni o gettare in acqua o sul fuoco la batteria.
- Sostituire la batteria solo con i ricambi HP previsti a questo scopo.

Le batterie e gli accumulatori non devono essere smaltiti insieme ai rifiuti domestici. Per procedere al riciclaggio o al corretto smaltimento, utilizzare il sistema di raccolta pubblico dei rifiuti o restituirli a HP, ai Partner Ufficiali HP o ai relativi rappresentanti.

Per ulteriori informazioni sulla sostituzione e sullo smaltimento delle batterie, contattare un Partner Ufficiale o un Centro di assistenza autorizzato.

バッテリーに関する注意



警告: 本製品はバッテリーを内蔵している場合があります。

- バッテリーを取り外している場合は、充電しないでください。
- バッテリーを水にさらしたり、60°C (140°F) 以上の温度にさらさないでください。
- バッテリーを誤用、分解、破壊したり、穴をあけたりしないでください。
- 外部極を短絡させたり、火や水に投棄しないでください。
- バッテリーを交換する際は、HP指定の製品と交換してください。

バッテリー、バッテリー パック、蓄電池は一般の家庭廃棄物と一緒に廃棄しないでください。リサイクルまたは適切に廃棄するため、公共の収集システム、HP、HPパートナー、またはHPパートナーの代理店にお送りください。

バッテリー交換および適切な廃棄方法についての情報は、HPのサポート窓口にお問い合わせください。

Declaración sobre las baterías



ADVERTENCIA: Este dispositivo podría contener una batería.

- No intente recargar las baterías si las extrae.
- Evite el contacto de las baterías con agua y no las exponga a temperaturas superiores a los 60 °C (140 °F).
- No utilice incorrectamente, ni desmonte, aplaste o pinche las baterías.
- No cortocircuite los contactos externos ni la arroje al fuego o al agua.
- Sustituya las baterías sólo por el repuesto designado por HP.

Las baterías, los paquetes de baterías y los acumuladores no se deben eliminar junto con los desperdicios generales de la casa. Con el fin de tirarlos al contenedor de reciclaje adecuado, utilice los sistemas públicos de recogida o devuélvalas a HP, un distribuidor autorizado de HP o sus agentes.

Para obtener más información sobre la sustitución de la batería o su eliminación correcta, consulte con su distribuidor o servicio técnico autorizado.


B Non-standard rack specifications

This appendix provides information on the requirements when installing the EVA4400 in a non-standard rack. All the requirements must be met to ensure proper operation of the storage system.

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 and clip nuts. [Figure 39 \(page 123\)](#) 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 39 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 28.25 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).

Determining the CG of a configuration may be necessary for safety considerations. CG considerations for CG calculations do not include cables, PDUs 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:

$$\sum d_{\text{component}} W = d_{\text{system CG}} W$$

where $d_{\text{component}}$ = 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 $(10 \times 1.75) + 2.7$ 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, leave 830 inches² (5,350 cm²) of open space evenly distributed from top to bottom to permit adequate airflow (equivalent to the required 64% 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 blank panels to fill empty front panel U-spaces in the rack. This ensures proper airflow. Using a rack without blank 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.

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 22 \(page 124\)](#) are available for use with the EVA and are included in this comparison. [Table 23 \(page 124\)](#) identifies the amount of time each UPS can sustain power under varying loads and with various UPS ERM (Extended Runtime Module) options.

NOTE: The specified power requirements reflect fully loaded enclosures (14 disks) .

Table 22 HP UPS models and capacities

UPS Model	Capacity (in watts)
R1500	1,340
R3000	2,700
R5500	4,500
R12000	12,000

Table 23 UPS operating time limits

Load (percent)	Minutes of operation		
	With standby battery	With 1 ERM	With 2 ERMs
R1500			
100	5	23	49

Table 23 UPS operating time limits *(continued)*

Load (percent)	Minutes of operation		
	With standby battery	With 1 ERM	With 2 ERMs
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	
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

Shock and vibration specifications

Table 24 (page 125) 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 24 (page 125). Transmission of site vibrations through non-HP racks exceeding these limits could cause operational failures of the system components.

Table 24 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.

Glossary

This glossary defines terms used in this guide or related to this product and is not a comprehensive glossary of computer terms.

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 133 mm (5.25 inches).
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.

A

- active member of a virtual disk family** 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 copy, 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** The storage system rules that govern how virtual disks are created. There are two rules:
- 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.
- ALUA** Asymmetric logical unit access. Operating systems that support asymmetric logical unit access work with the array’s active/active functionality to enable any virtual disk to be accessed through either of the array’s two controllers.
- 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 medium. 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** synonym of storage array, storage system, and virtual array. A group of disks in one or more disk enclosures combined with controller software that presents disk storage capacity as one or more virtual disks.
- array controller** *See* controller.
- array controller failover** The process that takes place when one controller assumes the workload of a failed companion controller.
- asynchronous** Events scheduled as the result of a signal requesting the event or that which is without any specified time relation.

B

backplane	An electronic printed circuit board that distributes data, control, power, and other signals among components in an enclosure.
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	The part of the power supply AC receptacle that engages the AC power cord connector to ensure that the cord cannot be accidentally disconnected.
battery	A rechargeable unit mounted within a controller enclosure that supplies backup power to the cache module in case of primary power shortage.
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	See fan.

C

cabinet	An alternate term used for a rack.
cable assembly	<p>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.</p> <ul style="list-style-type: none">• 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.
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	See battery.
carrier	A drive enclosure-compatible assembly containing a disk drive or other storage devices.
client	An intelligent device that requests services from other intelligent devices. In the context of HP P6000 Command View, a client is a computer that is used to access the software remotely using a supported browser.
clone	A full copy of a volume usable by an application.
communication 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.
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 or fans, 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 pair	Two connected controller modules that control a disk array.
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 disk group that is created when the array is initialized. The minimum number of disks the group can contain is eight. The maximum is the number of installed disks.
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	The 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 drive enclosure	A unit that holds storage system devices such as disk drives, power supplies, fans, I/O modules, and transceivers.
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 named group of disks selected from all the available disks in a disk array. One or more virtual disks can be created from a disk group. Also refers to the physical disk locations associated with a parity group.

disk migration state	<p>A physical disk drive operating state. A physical disk drive can be in a stable or migration state:</p> <ul style="list-style-type: none"> • 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.
DR group failover	An operation that reverses data replication direction so that the destination becomes the source and the source becomes the destination. Failovers can be planned or unplanned and can occur between DR groups or managed sets (which are sets of DR groups).
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 fabric	Two independent fabrics providing multipath connections between Fibre Channel end devices.
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	In a disk enclosure, a device such as a, power supply, disk, fan/blower, or I/O module. The object can be controlled, interrogated, or described by the enclosure services process.
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, I/O modules, or fans/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.
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	A sequential number assigned to each Software Code Identification (SWCID) event. It is a decimal number in the range 0-255.
Evt No.	See Event Number.
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 network of Fibre Channel switches or hubs and other devices.
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	See array controller failover or DR group failover.
failsafe	A safe state that devices automatically enter after a malfunction. Failsafe DR groups stop accepting host input and stop logging write history if a group member becomes unavailable.
fan	The variable speed airflow device that cools an enclosure or component by forcing ambient air into an enclosure or component and forcing heated air out the other side.
FATA	Fibre Attached Technology Adapted disk drive.
Fault Management Code	See FMC.
FC HBA	Fibre Channel Host Bus Adapter. See also FCA.
FCA	Fibre Channel Adapter. 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.
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 Channel	A data transfer architecture designed for mass storage devices and other peripheral devices that require high bandwidth.
Fibre Channel adapter	See FCA.
Fibre Channel drive enclosure	An enclosure that provides 12-port central interconnect for Fibre Channel arbitrated loops following the ANSI Fibre Channel disk enclosure standard.
Fibre Channel Loop	Fibre Channel Arbitrated Loop. The American National Standards Institute's (ANSI) document that specifies arbitrated loop topology operation.
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 3.5 inch (89 mm) and 5.25 inch (133 mm) high storage devices. Device heights are specified as low-profile (1 inch or 25.4 mm), half-height (1.6 inch or 41 mm), and full-height (5.25 inch 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. An assembly component that is designed to be replaced on site, without the system having to be returned to the manufacturer for repair.

G

Giga (G)	The notation to represent 10^9 or 1 billion (1,000,000,000).
gigabaud	An encoded bit transmission rate of one billion (10^9) bits per second.

H

HBA	Host Bus Adapter.
host	A computer that runs user applications and uses the information stored on an array.
Host Bus Adapter	Host bus adapter.
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.
host-side ports	See host ports.
hot-pluggable	The ability to add and remove elements or devices to a system or appliance while the appliance is running and have the operating system automatically recognize the change.
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.
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.
in-band communication	The communication that uses the same communications channel as the operational data.

INFORMATION condition	A drive enclosure EMU condition that may require action. This condition is for information purposes only and does not indicate the failure of an element.
initialization	A configuration step that binds the controllers together and establishes preliminary data structures on the array. Initialization also sets up the first disk group, called the default disk group, and makes the array ready for use.
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.
L	
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	<ol style="list-style-type: none"> 1. A connection of ports on fibre channel devices. 2. A full duplex connection to a fabric or a simplex connection of loop devices.
logon	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. (With Fibre Channel, not all 256 hexadecimal values are allowed as AL-PA values.)
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 47.
LUN	Logical unit number. A LUN results from mapping a SCSI logical unit number, port ID, and LDEV ID to a RAID group. The size of the LUN is determined by the emulation mode of the LDEV and the number of LDEVs associated with the LUN. For example, a LUN associated with two OPEN-3 LDEVs has a size of 4,693 MB.
M	
management agent	The HP P6000 Command View software that controls and monitors the HP 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	A 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.
Mega	A notation denoting a multiplier of 1 million (1,000,000).
metadata	The data in the first sectors of a disk drive that the system uses to identify virtual disk members.
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

near-online storage	On-site storage of data on media that takes slightly longer to access than online storage kept on high-speed disk drives.
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 HP Enterprise Storage System transceivers are non-OFC compatible.
NONCRITICAL Condition	An EMU condition that occurs when one or more elements in the drive enclosure fail or are operating outside specifications. The failure does not affect operation of the enclosure; all devices in the enclosure continue to operate according to specifications. If there are additional failures, however, the devices may not operate properly. UNRECOVERABLE and CRITICAL errors take precedence over this condition. This condition takes precedence over the INFORMATION condition. Early correction can prevent the loss of data.
NSC	Network storage controller. The HSV controllers used by the HP 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 pushbutton.
online storage	An allotment of storage space that is available for immediate use, such as a peripheral device that is turned on and connected to a server.
operator control panel	See OCP.

P

param	That portion of the HP HSV controller termination code display that defines: <ul style="list-style-type: none"> • The two-character parameter identifier that is a decimal number in the 0 through 31 range. • The eight-character parameter code that is a hexadecimal number.
password	A security interlock where the purpose is to allow: <ul style="list-style-type: none"> • 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 HP Enterprise Storage System elements.
PDU	Power distribution unit. The rack device that distributes conditioned AC or DC power within a rack.
petabyte	A unit of storage capacity that is the equivalent of 2^{50} , 1,125,899,906,842,624 bytes or 1,024 terabytes.
physical disk	<p>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.</p> <p>The physical disks, in aggregate, are called the array and constitute the storage pool from which the controllers create virtual disks.</p>
physical disk array	See array.
port	A physical connection that allows data to pass between a host and a disk array.
port-colored	Pertaining to the application of 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) <i>Standard RS310C</i> .
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 44.45 mm (1.75 inches) and based on EIA's <i>Standard RS310C</i> . For example, a 3U unit is 133.35 mm (5.25 inches) high, and a 4U unit is 177.79 mm (7.0 inches) 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	<ol style="list-style-type: none"> 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 HP Enterprise Storage System racks and enclosures to allow continuous system operation by preventing single points of power failure.
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.
RoHS	Reduction of Hazardous Substances.
room temperature	See ambient temperature.
RPO	Recovery point objective. The maximum age of the data you want the ability to restore in the event of a disaster. For example, if your RPO is six hours, you want to be able to restore systems back to the state they were in as of no longer than six hours ago. To achieve this objective, you need to make backups or other data copies at least every six hours.

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 where each bit of information is sent sequentially on a single channel, not simultaneously on all channels as occurs 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.
SFP	Small form-factor pluggable transceiver.
solid state disk (SSD)	A high-performance storage device that contains no moving parts. SSD components include either DRAM or EEPROM memory boards, a memory bus board, a CPU, and a battery card.
SSN	Storage system name. A unique 20-character name, assigned by HP P6000 Command View, that identifies a storage system.
storage carrier	See carrier.
storage pool	The aggregated blocks of available storage in the total physical disk array.
storage system	See array.
Storage System Name	See SSN.
switch	An electro-mechanical device that initiates an action or completes a circuit.

T

TC	Termination Code. An eight-character hexadecimal display that identifies why controller operations have halted.
Termination Code	See TC.
termination event	The 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 fibre channel elements such as hubs, controllers, or adapters.

U

UID	Unit identification.
uninitialized system	A state in which the storage system is not ready for use.
UNRECOVERABLE Condition	An EMU condition that occurs when one or more elements in the drive enclosure have failed and have disabled the enclosure. The enclosure may not be able to recover or bypass the failure; this will require repairs to correct the condition. This is the highest-level condition. It takes precedence over all other errors and requires immediate corrective action.
unwritten cached data	Also known as unflushed data. <i>See also</i> 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.
UUID	Unique universal identifier. A unique 128-bit identifier for each component of an array. UUIDs are internal system values that users cannot modify.

V

virtual disk	Variable disk capacity that is defined and managed by the array controller and presentable to hosts as a disk.
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.
Vraid0	Optimized for I/O speed and efficient use of physical disk space, but provides no data redundancy.
Vraid1	Optimized for data redundancy and I/O speed, but uses the most physical disk space.
Vraid5	Provides a balance of data redundancy, I/O speed, and efficient use of physical disk space.
Vraid6	Offers the features of Vraid5 while providing more protection for an additional drive failure, but uses additional physical disk space.

W

World Wide Name	<i>See</i> 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 identifier assigned to a Fibre Channel device.

Index

A

- AC power
 - distributing, 20
- accessing
 - multipathing, 41
 - Secure Path, 41
- adding
 - hosts, 47
- adding hosts, 42

B

- bad image header, 102
- bad image segment, 103
- bad image size, 103
- battery replacement notices, 120
- bays
 - locating, 10
 - numbering, 10
- bidirectional operation, 11

C

- cabling controller, 19
- Cache batteries failed or missing, 101
- Canadian notice, 111
- configuration
 - physical layout, 9
- configuring EVA, 58
- configuring the ESX server, 58
- connection suspended, 102
- connectivity
 - verifying, 60
- connectors
 - power IEC 309 receptacle, 21
 - power NEMA L5-30R, 21
 - power NEMA L6-30R, 21
 - protecting, 39
- controller
 - cabling, 19
 - connectors, 19
 - defined, 9
 - HSV300, 9
 - HSV300-S, 15
- conventions
 - document, 108
 - text symbols, 108
- creating virtual disks, 42
- creating volume groups, 44
- customer self repair, 108
 - parts list, 71

D

- Declaration of Conformity, 111
- disk drives
 - defined, 13
 - reporting status, 13

- disk enclosures
 - defined, 9
 - front view, 10
- DiskMaxLUN, 60
- disks
 - labeling, 57
 - partitioning, 57
- Disposal of waste equipment, European Union, 116
- DMP, 53
- document
 - conventions, 108
 - prerequisites, 106
- documentation
 - providing feedback, 106
- DR group empty, 101
- DR group logging, 102
- DR group merging, 102
- dust covers, using, 40

E

- enclosures
 - bays, 10
- Enterprise rack
 - physical layout, 9
- European Union notice, 111

F

- fabric setup, 53
- FATA drives, using, 27
- FC loops, 11
- FCA
 - configuring, 50
 - configuring QLogic, 52
 - configuring, Emulex, 50
- Federal Communications Commission notice, 110
- feedback
 - documentation, 106
- fiber optics
 - protecting cable connectors, 39
- Fibre Channel Disk Enclosure
 - defined, 9

H

- hardware components, 9
- help
 - obtaining, 106
- high availability
 - HSV Controllers, 14
- hosts
 - adding, 47
- HP
 - technical support, 106
- HP P6000 Command View
 - adding hosts with, 42
 - creating virtual disk with, 42
 - using, 42

HSV Controllers
defined, 9

I

I/O modules
bidirectional, 11
image already loaded, 103
image incompatible with configuration, 102
image too large, 102
image write error, 103
implicit LUN transition, 32
incompatible attribute, 101
invalid
parameter id, 98
quorum configuration, 98
target handle, 98
target id, 98
time, 98
invalid cursor, 100
invalid state, 100
invalid status, 102
invalid target, 100
iopolicy
setting, 54

J

Japanese notices, 112

K

Korean notices, 112

L

laser compliance notices, 114
lock busy, 100
logical disk presented, 100
logical disk sharing, 103
lpfc driver, 50

M

maximum number of objects exceeded, 101
maximum size exceeded, 102
media inaccessible, 98
multipathing
accessing, 41
policy, 59

N

no FC port, 98
no image, 98
no logical disk for Vdisk, 100
no more events, 100
no permission, 98
non-standard rack, specifications, 123
not a loop port, 98
not participating controller, 98

O

object does not exist, 99
objects in use, 98

operation rejected, 101
Oracle SAN driver stack, 50
Oracle StorEdge, 50
Traffic Manager, 53
other controller failed, 101

P

parts
replaceable, 71
password mismatch, 102
PDUs, 20
physical configuration, 9
power connectors
IEC 309 receptacle, 21
NEMA L5-30R, 21
NEMA L6-30R, 21
prerequisites, 106
presenting virtual disks, 42
protecting fiber optic connectors
cleaning supplies, 40
dust covers, 40
proxy reads, 32

Q

qla2300 driver, 52

R

rack
defined, 20
non-standard specifications, 123
physical layout, 9
rack configurations, 20
recycling notices, 116
regulatory compliance
Canadian notice, 111
European Union notice, 111
identification numbers, 110
Japanese notices, 112
Korean notices, 112
laser, 114
recycling notices, 116
Taiwanese notices, 113

S

Secure Path
accessing, 41
security credentials invalid, 101
Security credentials needed, 101
slots see enclosures, bays
spare parts, 72
status, disk drives, 13
storage connection down, 101
storage not initialized, 98
storage system racks, defined, 20
Subscriber's choice, HP, 106
symbols in text, 108
system rack configurations, 20

T

- Taiwanese notices, [113](#)
- technical support
 - HP, [106](#)
 - service locator website, [106](#)
- text symbols, [108](#)
- time not set, [100](#)
- timeout, [100](#)
- transport error, [100](#)

U

- universal disk drives, [13](#)
- unknown id, [100](#)
- unknown parameter handle, [100](#)
- unrecoverable media error, [100](#)
- UPS, selecting, [124](#)

V

- Vdisk DR group member, [101](#)
- Vdisk DR log unit, [101](#)
- Vdisk not presented, [101](#)
- verifying virtual disks, [55](#)
- Veritas Volume Manager, [53](#)
- version not supported, [100](#)
- vgcreate, [44](#)
- virtual disks
 - configuring, [43](#), [49](#), [55](#)
 - presenting, [42](#)
 - verifying, [55](#), [56](#), [60](#)
- VMware
 - VAAI Plug-in, [61](#)
- volume groups, [44](#)
- volume is missing, [100](#)

W

- websites
 - customer self repair, [108](#)
 - HP, [106](#)
 - HP Subscriber's choice for business, [106](#)
 - Oracle documentation, [58](#)
 - Symantec/Veritas, [54](#)
- WWLUN ID
 - identifying, [55](#)

Z

- zoning, [53](#)