Tru64 UNIX

HP StorageWorks Disk Array XP operating system configuration guide

XP48
XP128
XP512
XP1024
XP12000

fourth edition (August 2004)

part number: A5951-96062

This guide describes the requirements and procedures for connecting and configuring the XP family of disk arrays to a Tru64 UNIX system.

Printed in the U.S.A.

HP StorageWorks Disk Array XP Operating System Configuration Guide: Tru64

fourth edition (August 2004)

part number: A5951-96062

Contents

About this guide

This guide describes the requirements and procedures for connecting the XP family of disk arrays to an HP system, and configuring the new disk array for operation with Tru64 UNIX.

Intended audience

This guide is intended for system administrators who have knowledge of the following topics:

- Data processing concepts
- Direct access storage device subsystems and their basic functions
- Disk arrays and RAID technology
- Operating system commands and utilities

Disk arrays

Unless otherwise noted, the term *disk array* refers to these disk arrays:

HP Surestore Disk Array XP512
HP Surestore Disk Array XP48
HP StorageWorks Disk Array XP128
HP StorageWorks Disk Array XP1024
HP StorageWorks XP12000 Disk Array

Related documentation

HP provides the following related documentation:

- *HP StorageWorks Disk Array XP128: Owner's Guide*
- *HP StorageWorks Disk Array XP1024: Owner's Guide*
- *HP StorageWorks XP12000 Disk Array: Owner's Guide*

For information about operating system commands and third-party products, refer to the manufacturer's documentation.

Conventions

This guide uses the following text conventions.

Figure 1	Blue text represents a cross-reference. For the online version of this guide, the reference is linked to the target.
www.hp.com	Underlined, blue text represents a website on the Internet. For the online version of this guide, the reference is linked to the target.
literal	Bold text represents literal values that you type exactly as shown, as well as key and field names, menu items, buttons, file names, application names, and dialog box titles.
variable	Italic type indicates that you must supply a value. Italic type is also used for manual titles.
`input/output`	Monospace font denotes user input and system responses, such as output and messages.
Example	Denotes an example of input or output. The display shown in this guide may not match your configuration exactly.
[ ]	Indicates an optional parameter.
{ }	Indicates that you must specify at least one of the listed options.
|	Separates alternatives in a list of options.

HP technical support

In North America, call technical support at 1-800-652-6672, available 24 hours a day, 7 days a week.

Outside North America, call technical support at the nearest location. Telephone numbers for worldwide technical support are listed on the HP website under support:

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

Be sure to have the following information available before calling:

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

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

HP storage website

For the most current information about HP StorageWorks XP products, visit the support website. Select the appropriate product or solution from this website:

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

For information about product availability, configuration, and connectivity, consult your HP account representative.

HP authorized reseller

For the name of your nearest HP authorized reseller, you can obtain information by telephone:

United States	1-800-345-1518
Canada	1-800-263-5868
Or contact:	www.hp.com

Revision history

May 1999	First release.
September 1999	Added OPEN-8 emulation.
September 2000	Added support for XP512.
August 2004	Added support for XP12000.

Warranty statement

HP warrants that for a period of ninety calendar days from the date of purchase, as evidenced by a copy of the invoice, the media on which the Software is furnished (if any) will be free of defects in materials and workmanship under normal use.

DISCLAIMER. EXCEPT FOR THE FOREGOING AND TO THE EXTENT ALLOWED BY LOCAL LAW, THIS SOFTWARE IS PROVIDED TO YOU "AS IS" WITHOUT WARRANTIES OF ANY KIND, WHETHER ORAL OR WRITTEN, EXPRESS OR IMPLIED. HP SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OR CONDITIONS OF MERCHANTABILITY, SATISFACTORY QUALITY, NON-INFRINGEMENT, TITLE, ACCURACY OF INFORMATIONAL CONTENT, AND FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow exclusions of implied warranties or conditions, so the above exclusion may not apply to you to the extent prohibited by such local laws. You may have other rights that vary from country to country, state to state, or province to province.

WARNING! **YOU EXPRESSLY ACKNOWLEDGE AND AGREE THAT USE OF THE SOFTWARE IS AT YOUR SOLE RISK.** HP DOES NOT WARRANT THAT THE FUNCTIONS CONTAINED IN THE SOFTWARE WILL MEET YOUR REQUIREMENTS, OR THAT THE OPERATION OF THE SOFTWARE WILL BE UNINTERRUPTED, VIRUS-FREE OR ERROR-FREE, OR THAT DEFECTS IN THE SOFTWARE WILL BE CORRECTED. THE ENTIRE RISK AS TO THE RESULTS AND PERFORMANCE OF THE SOFTWARE IS ASSUMED BY YOU. HP DOES NOT WARRANT OR MAKE ANY REPRESENTATIONS REGARDING THE USE OR THE RESULTS OF THE USE OF THE SOFTWARE OR RELATED DOCUMENTATION IN TERMS OF THEIR CORRECTNESS, ACCURACY, RELIABILITY, CURRENTNESS, OR OTHERWISE. NO ORAL OR WRITTEN INFORMATION OR ADVICE GIVEN BY HP OR HP'S AUTHORIZED REPRESENTATIVES SHALL CREATE A WARRANTY.

LIMITATION OF LIABILITY. EXCEPT TO THE EXTENT PROHIBITED BY LOCAL LAW, IN NO EVENT INCLUDING NEGLIGENCE WILL HP OR ITS SUBSIDIARIES, AFFILIATES, DIRECTORS, OFFICERS, EMPLOYEES, AGENTS OR SUPPLIERS BE LIABLE FOR DIRECT, INDIRECT, SPECIAL, INCIDENTAL, CONSEQUENTIAL, PUNITIVE OR OTHER DAMAGES (INCLUDING LOST PROFIT, LOST DATA, OR DOWNTIME COSTS), ARISING OUT OF THE USE, INABILITY TO USE, OR THE RESULTS OF USE OF THE SOFTWARE, WHETHER BASED IN WARRANTY, CONTRACT, TORT OR OTHER LEGAL THEORY, AND WHETHER OR NOT ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Your use of the Software is entirely at your own risk. Should the Software prove defective, you assume the entire cost of all service, repair or correction. Some jurisdictions do not allow the exclusion or limitation of liability for incidental or consequential damages, so the above limitation may not apply to you to the extent prohibited by such local laws.

NOTE. EXCEPT TO THE EXTENT ALLOWED BY LOCAL LAW, THESE WARRANTY TERMS DO NOT EXCLUDE, RESTRICT OR MODIFY, AND ARE IN ADDITION TO, THE MANDATORY STATUTORY RIGHTS APPLICABLE TO THE LICENSE OF THE SOFTWARE TO YOU; PROVIDED, HOWEVER, THAT THE CONVENTION ON CONTRACTS FOR THE INTERNATIONAL SALE OF GOODS IS SPECIFICALLY DISCLAIMED AND SHALL NOT GOVERN OR APPLY TO THE SOFTWARE PROVIDED IN CONNECTION WITH THIS WARRANTY STATEMENT.

1

Installation

Installation of the HP StorageWorks Disk Array XP is performed by your HP service representative and you. The HP service representative installs the disk array and formats the disk devices. You configure the host server for the new devices with assistance from the HP service representative.

Features and requirements

The disk array and host have the following features and requirements.

- HP StorageWorks disk arrays:

  XP48: Up to 48 drives from 72 GB to 8.7 TB, 24 FC ports

  XP128: From 8 to 128 drives for up to 18 TB, 48 FC ports

  XP512: Up to 512 drives from 72 GB to 93 TB, 48 FC ports

  XP1024: From 8 to 1024 drives for up to 149 TB, 64 FC ports

  XP12000: Up to 1152 drives for up to 165 TB, 128 FC ports
- PCI-based AlphaStation or AlphaServer
- Tru64 UNIX, versions 4.0F/G or 5.1A/B with latest supported patches.
- Host Bus Adapters (HBAs): Install adapters and all utilities and drivers. Refer to the adapter documentation for installation details.
- *(Recommended)* HP StorageWorks Command View XP with LUN management feature or Remote Control with the LUN Configuration Manager XP option for configuring disk array ports and paths.
- *(Recommended)* HP StorageWorks Secure Manager XP: Allows the host to access only array devices for which it is authorized.
- Other available XP Software (some may not apply to your system):
  HP StorageWorks Business Copy XP
  HP StorageWorks Continuous Access XP
  HP StorageWorks Continuous Access Extension XP
  HP StorageWorks Auto LUN XP
  HP StorageWorks Data Exchange XP
  HP StorageWorks Resource Manager XP
  HP StorageWorks RAID Manager XP
  HP StorageWorks Cache LUN XP
  HP StorageWorks Auto Path XP
  HP StorageWorks Cluster Extension XP
  HP StorageWorks Performance Advisor XP software

Fibre Channel interface

The XP48, XP128, XP512, XP1024, and XP12000 disk arrays support these 1 Gbps and 2 Gbps Fibre Channel interfaces:

- Short-wave non-OFC (open fiber control) optical interface
- Multimode optical cables with SC or LC connectors
- Public or private arbitrated loop (FC-AL) or fabric direct attach
- Fibre Channel switches

Even though the interface is Fibre Channel, this guide uses the term "SCSI disk" because disk array devices are defined to the host as SCSI disks.

Device types

The disk arrays support the following device types:

- **OPEN-*x* devices:** OPEN-x logical units represent disk devices. Except for OPEN-V, these devices are based on fixed sizes. OPEN-V is a user-defined size. Supported emulations include OPEN-3, OPEN-8, OPEN-9, OPEN-E, OPEN-L, and OPEN-V devices.
- **LUSE devices (OPEN-*x***n*):** Logical Unit Size Expansion (LUSE) allows you to combine 2 to 36 OPEN-x devices to create expanded LDEVs larger than standard OPEN-*x* disk devices. For example, an OPEN-x LUSE volume created from ten OPEN-x CVS volumes is designated as OPEN-x*10.
- **CVS devices (OPEN-x CVS):** Volume Size Configuration (VSC) defines custom volumes (CVS) that are smaller than normal fixed-sized logical disk devices (volumes). (OPEN-V is a CVS-based custom disk size that you determine. OPEN-L does not support CVS.)
- **LUSE (expanded) CVS devices (OPEN-x*n CVS):** LUSE CVS combines CVS devices to create an expanded device. This is done by first creating CVS custom-sized devices and then using LUSE to combine from 2 to 36 CVS devices. For example, if three OPEN-9 CVS volumes are combined to create an expanded device, this device is designated as OPEN-9*3-CVS.

Failover

The disk arrays support many standard software products that provide host, application, or I/O path failover and logical volume (storage) management. TruCluster is one of the supported applications.

Caution *Tru64 verson 5 fully supports dynamic load balancing and failover when multiple HBAs are connected to the same LUN. This feature is ONLY available with version 5.*

On Tru64 version 4, DO NOT connect multiple HBAs to the same LUN and DO NOT allow the same LUN to be accessible from multiple HBAs on Tru64 version 4 hosts. ON the version 4 system, this will allow multiple device names to access the same LUN. If both device names are used, data corruption will result.

SNMP configuration

The disk arrays support standard Simple Network Management Protocol (SNMP) for remotely managing the disk array from the host. The SNMP agent on the remote console PC or Command View can provide status and Remote Service Information Message (R-SIM) reporting to the SNMP manager on the host for up to eight disk arrays. To configure the SNMP manager on the host, refer to the operating system documentation.

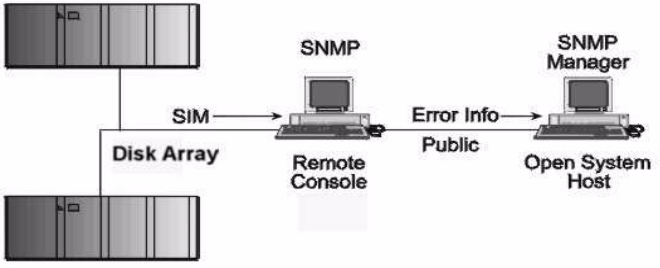

RAID Manager command devices

RAID Manager manages Business Copy (BC) and/or Continuous Access (CA) operations from a server host. To use RAID Manager with BC or CA, you must use Command View or LUN Configuration Manager to designate at least one LDEV as a command device. Refer to the Command View or LUN Configuration Manager user guide for information about how to designate a command device.

Installation procedures

The HP representative and you perform the following procedures:

1. Install and configure the disk array (page 17).
   - Setting the System Option Modes
   - Configuring the Fibre Channel ports
   - Setting the Host Mode for the disk array ports
2. Install and configure the host (page 21).
   - Loading the OS and software
   - Installing and configuring the HBAs
   - Clustering and Fabric zoning
   - Fabric zoning and LUN security for multiple operating systems
3. Connect the disk array (page 25).
   - Defining the paths
   - Verifying disk array device recognition
4. Configure disk devices (page 28).
   - Writing the partition labels
   - Writing the partition labels
   - Creating the file systems
   - Creating mount directories
   - Mounting the file systems
   - Verifying the file systems
   - Setting and verifying automatic mounting at bootup

Install and configure the disk array

The HP service representative performs the following tasks:

- Assembling hardware and installing software
- Loading the microcode updates
- Installing the channel adapters (CHAs) and cabling
- Installing and formatting devices

You perform the additional tasks below. If you do not have Command View or LUN Configuration Manager, your HP service representative can perform these tasks for you.

Setting the System Option Modes

The HP representative sets the System Option Mode(s) based on the operating system and software configuration of the host.

Configuring the Fibre Channel ports

Configure the disk array Fibre Channel ports by using Command View (shown) or the Fibre Parameter window in LUN Configuration Manager. Select the settings for each port based on your storage area network topology. Use switch zoning if you connect different types of hosts to the array through the same switch.

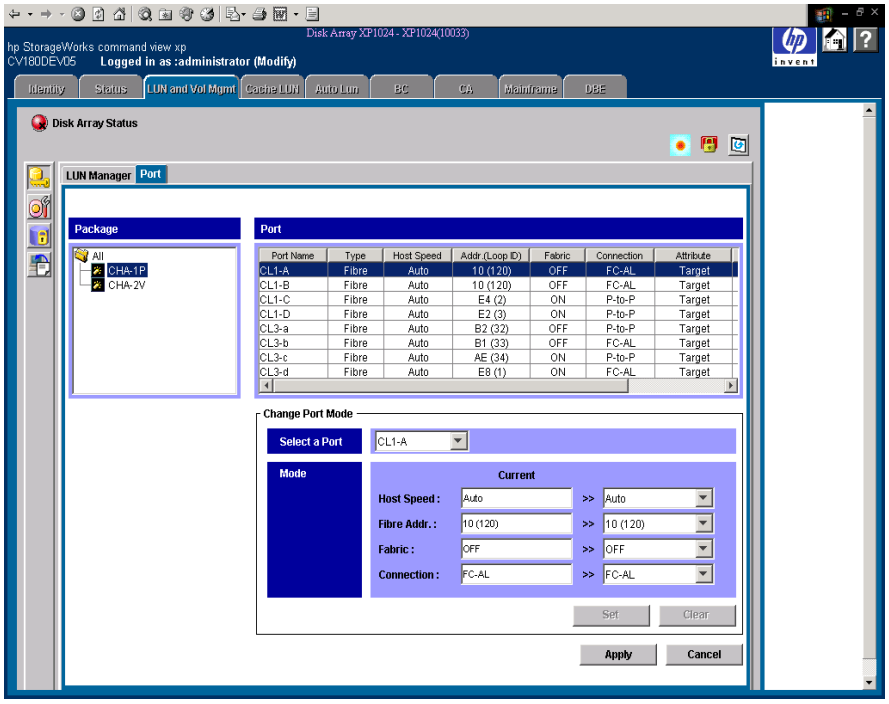

Fibre Address

In fabric environments, the port addresses are assigned automatically. In arbitrated loop environments, you set the port addresses by selecting a unique arbitrated loop physical address (AL-PA) or loop ID for each port.

Fabric and Connection parameter settings

You can set each array port to FABRIC ON or OFF with connections of POINT-TO-POINT or FC-AL as shown in the following table and figures. For detailed topology information, refer to the *HP StorageWorks SAN Design Reference Guide* on the hp.com website.

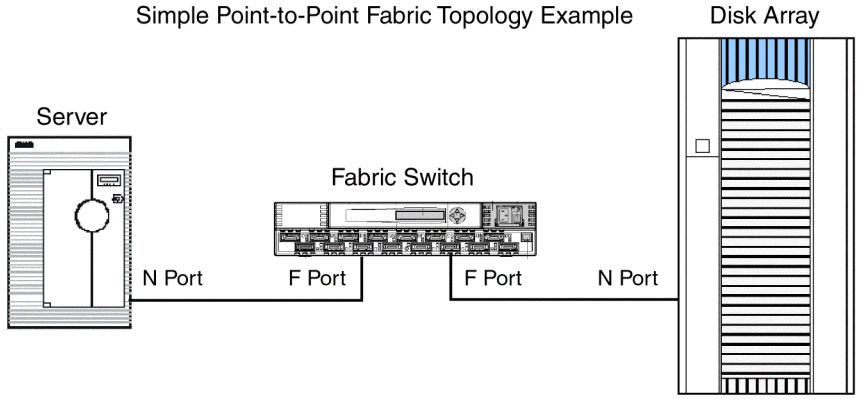

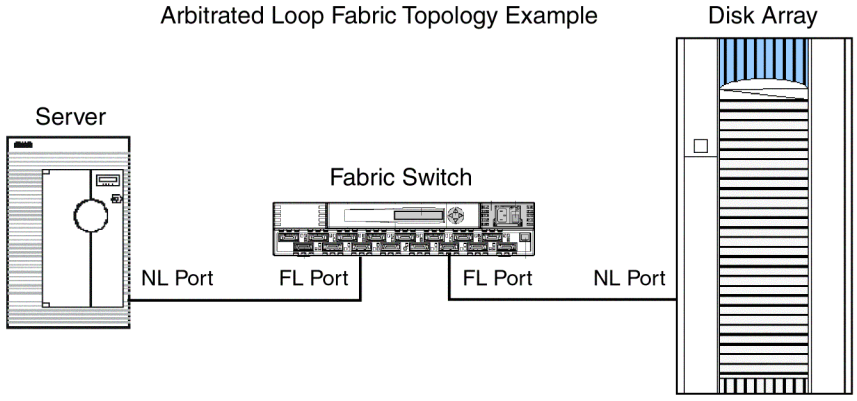

Fabric Parameter	Connection Parameter	Provides
ON	FC-AL	NL-port (fabric loop port)
ON	Point-to-Point	N-port (fabric port)
OFF	FC-AL	NL-port (private arbitrated loop)
OFF	Point-to-Point	Not supported

Setting the Host Mode for the disk array ports

The disk array ports have Host Modes that you must set depending on the host you use. After the disk array is installed, use Command View (shown) or LUN Configuration Manager to set the Host Mode for each port.

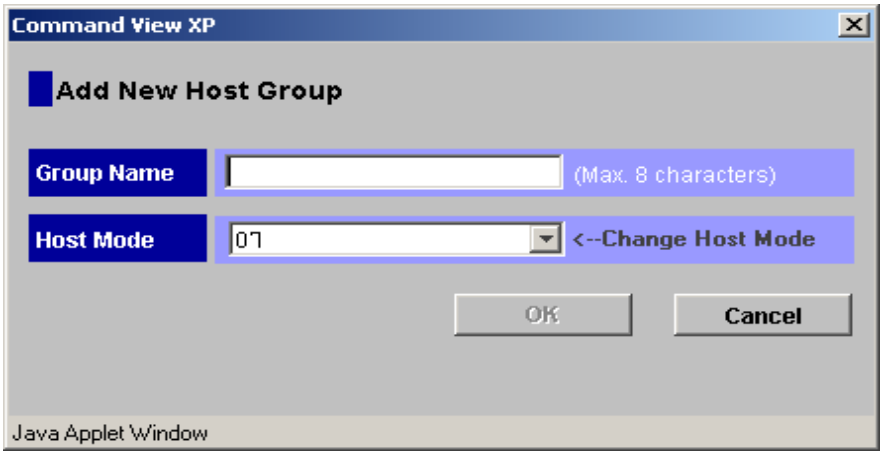

The required Host Mode setting for **Tru64 is 07**.

Install and configure the host

Install and configure the host and host bus adapters (HBAs) that connect the host to the disk array.

Loading the OS and software

Follow the manufacturer's instructions to load the operating system and software onto the host. Load all OS patches and configuration utilities supported by HP and the HBA manufacturer.

Installing and configuring the HBAs

Install and configure the host bus adapters using the HBA manufacturer's instructions.

Set HBA loop or fabric mode

Set the HBAs to run in arbitrated loop mode or in fabric mode as follows.

1. Set the Tru64 UNIX console to diagnostic mode.

   ```
   P00>>>set mode diag
   ```

2. Display the HBA configuration:

   ```
   P00>>>wwidmgr -show adapter
   ```

3. Set the HBA topology to fabric or loop. The 9999 qualifier sets all adapters with one command.

   ```
   P00>>>wwidmgr -set adapter -item 9999 -topology loop
   ```

   OR

   ```
   P00>>>wwidmgr -set adapter -item 9999 -topo fabric
   ```

Configuring system files

Configure the Tru64 system to recognize the HBA as explained in the Tru64 Installation Guide or New Hardware Delivery (NHD) kit. This usually consists of booting **genvmunix** and executing **doconfig** (as explained in the Installation Guide). This can also sometimes be done by editing the system kernel configuration file as explained below.

1. Edit the system kernel file to add a line defining the adapter:

Example

```
bus          emx0   at   pci0   slot3
controller   scsi5  at   emx0   slot0    <--- Add this line
```

2. **(Tru64 v4 only)** Edit the **/sys/data/cam_data.c** system file to change the I/O timeout value from 10 to 60:

   u_long cdisk_to_def = 60.

3. **(Tru64 v4 only)** Edit the TID in the **emx_data.c** file. The WWN that the host recognizes is displayed in a configuration file, such as **/etc/emx.info**. In this file, if the target ID for a WWN is equal to or greater than 8, copy it into the **emx_data.c** file and change the target ID to a value between 0 and 7.

4. Use the **doconfig -c** *config_file* command to reconfigure the kernel, where *config_file* is usually the system name.

5. Shut down and reboot the system.

   If system boot fails, reboot with a backup copy of the kernel and check the edited files for errors.

(Tru64 v5 only) When you add a new HBA after an HBA of the same type has already been installed, a simple reboot causes the host to recognize the new HBA. No rebuilding or reconfiguration is required.

Clustering and Fabric zoning

If you plan to use clustering, install and configure the clustering software on the servers. Clustering is the organization of multiple servers into groups. Within a cluster, each server is a node. Multiple clusters compose a multi-cluster environment. The following example shows a multi-cluster environment with three clusters, each containing two nodes. The nodes share access to the disk array.

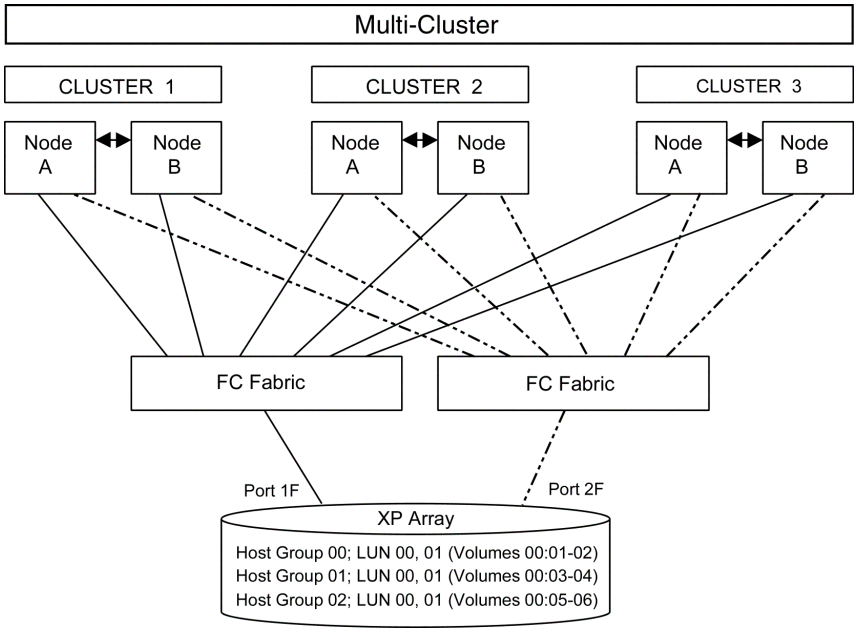

Within the Storage Area Network (SAN), the clusters may be homogeneous (all the same operating system) or they may be heterogeneous (mixed operating systems). How you configure LUN Security and fabric zoning depends on the operating system mix and the SAN configuration.

Fabric zoning and LUN security for multiple operating systems

By using appropriate zoning and LUN security, you can connect multiple clusters of various operating systems to the same switch and fabric:

- Host zones must contain only homogeneous operating systems.
- Storage port zones may overlap if more than one operating system needs to share an array port.
- Heterogeneous operating systems may share an XP array port if you use Secure Manager and set the appropriate host group and mode; all others must connect to a dedicated XP array port.
- Use Secure Manager for LUN isolation when multiple hosts connect through a shared array port. Secure Manager provides LUN security by allowing you to restrict which LUNs each host can access.

Environment	OS Mix	Fabric Zoning	LUN Security
Standalone SAN (non-clustered)	homogeneous (a single OS type present in the SAN)	Not required	Must be used when multiple hosts connect through a shared port
	heterogeneous (more than one OS type present in the SAN)	Required	
Clustered SAN	homogeneous (a single OS type present in the SAN)	Not required	Must be used when multiple cluster nodes connect through a shared port
	heterogeneous (more than one OS type present in the SAN)	Required	
Multi-Cluster SAN	homogeneous (a single OS type present in the SAN)	Required	Must be used when multiple cluster nodes connect through a shared port
	heterogeneous (more than one OS type present in the SAN)	Required	

Connect the disk array

Connect the disk array to the host as follows:

1. The HP service representative verifies operational status of the disk array channel adapters, LDEVs, and paths.
2. The HP representative connects the Fibre Channel cables between the disk array and the fabric or host.
3. Verify the ready status of the disk array and peripherals.

Defining the paths

Use Command View (shown) or LUN Configuration Manager to create paths (LUNs) between hosts and volumes in the disk array, also called LUN mapping. LUN mapping includes these tasks:

- Configuring ports
- Setting LUN security
- Creating host groups by operating system and setting their host modes
- Assigning host bus adapter WWNs to host groups.
- Mapping volumes to host groups (by assigning LUNs).

For details, see the Command View or LUN Configuration Manager guide. HP recommends that you note LUNS and their ports, WWNs, nicknames, and LDEVs for later use in verifying host and device configuration.

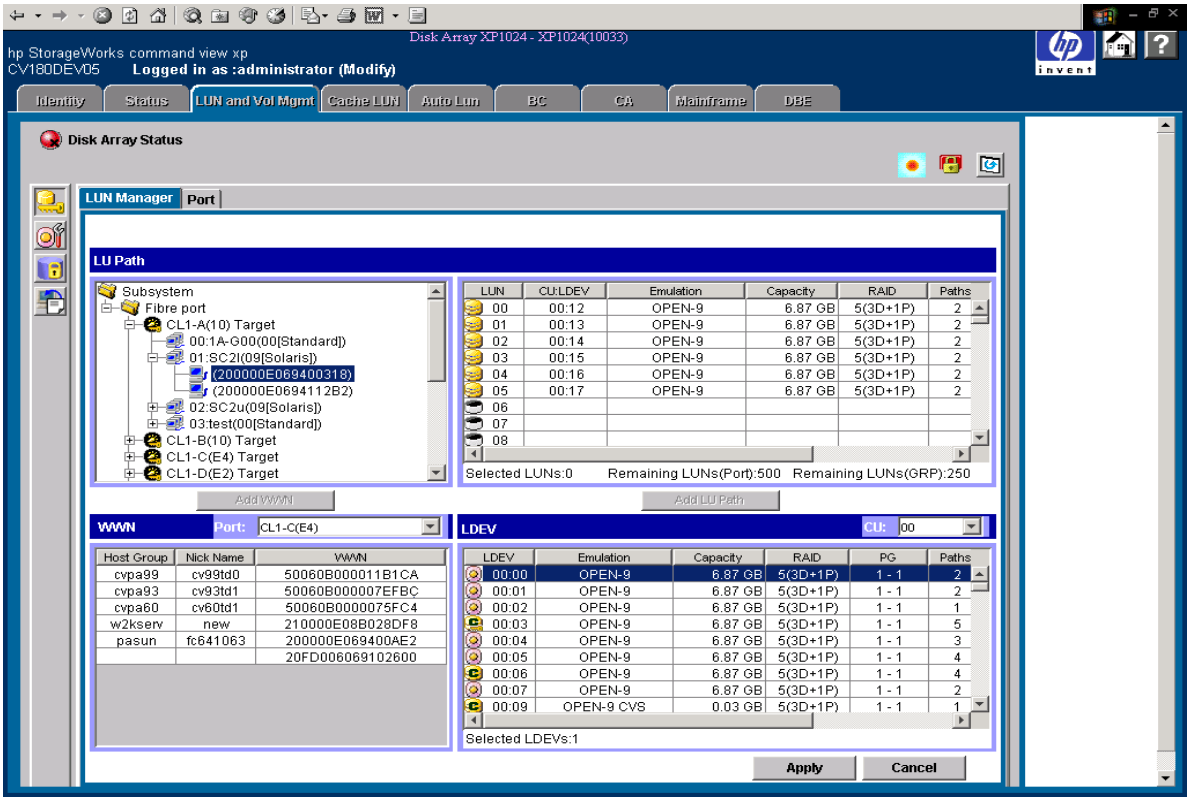

Verifying disk array device recognition

Use the **scu show edt** command at the UNIX prompt to see the list of new disk array devices.

The device files are created automatically in Tru64 UNIX during system startup. Device files are created for each logical unit.

1. Verify the character-type device files have been automatically created:

Tru64 v4 example

file /dev/rrz*

Tru64 v5 example

file /dev/rdisk/dsk*

2. Verify the block-type device files have been automatically created:

Tru64 v4 example

file /dev/rz*

Tru64 v5 example

file /dev/disk/dsk*

(*Tru64 v4 only)* You can verify the allocation for the device files by calculating as follows:

Block type : rz*XYZ*
Character type : rrz*XYZ*
X = b through h = LUN1 through LUN7 (no letter is used for LUN0)
Y = SCSI bus number × 8 + SCSI ID
Z = partition (a through h)

Note: Target ID > 8 is not allowed for hosts.

Example of Calculation

In this example, OPEN-3 of LUNs 0 through 7 with SCSI ID = 0 is connected to the card containing SCSI bus number = 1. With this calculation, Y=1x8+0=8, the following device files are allocated to each partition of LUN1 (X = b):

```
rzb8a,rrzb8a,rzb8b,rrzb8b,rzb8c,rrzb8c,
rzb8d,rrzb8d,rzb8e,rrzb8e,rzb8f,rrzb8f,
rzb8g,rrzb8g,rzb8h,rrzb8h
```

Configure disk devices

Configure the disk array devices in the same way you would configure any new disk on the host. Creating scripts to configure all devices at once may save you considerable time.

Configuring devices typically requires these steps:

1. Partition and label each device using the **disklabel** command.
2. Create a file system for each device using the **newfs** command.
3. Create a mount directory for each device using the **mkdir** command.
4. Enter each device into the mount table by editing **/etc/fstab**.
5. Reboot the host. After bootup use the **df** command to verify the devices auto-mounted.
6. Verify file system operation by copying a file to each device and then deleting it.

Writing the partition labels

Use the **disklabel** command to label the partition for each logical unit. Partition c specifies the entire area in the logical unit. Check that no errors are found in the partition settings after the labeling.

You can edit the disk partition size using the **disklabel** command with option **–e**. When the **disklabel –e** command is executed, the **vi** editor for the environment in which you are working starts up. After completing the editing, save the file and execute the **disklabel** command again; the partition setting is renewed.

To write the partition label:

1. Enter **disklabel –rw**.

Tru64 v4 **# disklabel –rw rz8 OPEN-3**

Specify the device file in the rz*XY* format (*X* = LUN, *Y* = SCSI bus number × 8 + SCSI ID).

Specify the disk name that is registered for that device file.

Tru64 v5 # **disklabel –rw dsk10 OPEN-3**

2. Enter **disklabel –r** to verify labeling and partition settings.

Tru64 v4 **# disklabel –r rz8**

Specify the device file in the rz*XY* format (*X* = LUN, *Y* = SCSI bus number × 8 + SCSI ID).

Tru64 v5 # **disklabel –r dsk10**

Creating the file systems

Create a file system for each new OPEN-*x* device. Optionally, you can create and use an advanced file system (AdvFS) to overcome the size and speed limitations of the file system. If you are not sure which file system is right for your setup, contact HP customer support.

To create a file system for each device:

 # newfs *device_file_name*

To create an advanced file system:

You can create a new advanced file system domain, or you can add a new fileset to an existing domain.

If you allocate multiple disk partitions to a domain, the advanced file system utilities must be installed. Installing the utilities will enable you to allocate multiple disk partitions to a domain using the **addvol** command, for example:

addvol /dev/rzb8c domain1

1. Create a new domain.

   # mkfdmn *device_file_name domain_name*

Tru64 v4 Example

mkfdmn /dev/rz12c domain1.

The example uses the block-type device file name (rz*XYZ*) to create an advanced file system domain for device rz12c.

Tru64 v5 Example

newfs /dev/rdisk/dsk10c OPEN-3

2. Create a new fileset in the new or existing advanced file system domain.

   # mkfset *domain_name fileset_name*

   For example, to create fileset1 in domain1:

   # **mkfset domain1 fileset1**

Creating mount directories

Create a mount directory for each device. Assign each mount directory a unique name that identifies the device being mounted.

1. Create a mount directory:

   # **mkdir** */mount_directory_name*

   For example, to create a mount directory for LUN 2 (partition c) on the disk array, enter **# mkdir /HP5700_LU2c**.

2. Verify the new mount directory:

   # ls /

Mounting the file systems

After the file systems and mount directories have been created, you can mount the file system for each new device.

To mount a Tru64 file system:

1. Mount device:

   # mount *device_file_name mount_directory*

   For example, to mount device rz12c with mount directory name HP5700_LU2c, enter **# mount /dev/rz12c /HP5700_LU2c**.

Tru64 v5

mount /dev/disk/dsk10c mnt

2. Assign the appropriate ownership and permissions:

   # chown *owner*:*group *device_file_name**

   For example, to assign ownership to rz12c with owner Oracle, group dba, enter **# chown oracle:dba *rz12c***.

To mount an advanced file system:

1. Mount the file system:

   # mount –t advfs *domain_name#fileset_name mount_directory*

   For example, to mount the file system with mount directory **HP5700_LU2c**, enter **mount –t advfs domain1#fileset1 /HP5700_LU2c**.

2. Assign the appropriate ownership and permissions:

   # chown *owner*:*group *device_file_name**

   For example, to assign ownership to rz12c with owner Oracle, group dba, enter **# chown oracle:dba *rz12c***.

Verifying the file systems

Verify that the new file systems were created correctly and are functioning properly.

1. Display all mounted file systems.

   # **df**

   The default display for drive capacity is 512-byte blocks. To view the capacity in KB rather than in 512-byte blocks, enter **df –k**.

2. Go to a new device directory:

   # cd /*mount_directory*

   For example, enter **cd /HP5700_LU2c**.

3. Copy a file from the root directory to the new device:

   # cp /*filename file_name***.back1**

   For example, to copy file vmunix from the root directory to the HP5700_LU2c device, enter **cp /vmunix vmunix.back1**.

4. Copy a file to the new device again:

   # cp /*filename file_name***.back2**

   For example, to copy the same file again, enter **cp /vmunix vmunix.back2**.

5. List the files in the current directory:

   # **ls –l**

   The **vmunix.back1** and **vmunix.back2** files should be shown.

6. Delete the files you copied:

   # rm *file_name*

   For example, enter **rm vmunix.back1** to remove the file **vmunix.back1**.

Setting and verifying automatic mounting at bootup

The **/etc/fstab** file contains boot time mounting parameters for disk devices.

1. Edit the **/etc/fstab** file after first making a backup copy of the file. Add a line for each new device to be mounted.

Example

```
#vi /etc/fstab
1               2             3       4     5     6
/dev/rz4a       /             ufs     rw    1     1
/proc           /proc         procfs  rw    0     0
/dev/rz12a      /usr          ufs     rw    1     2
/dev/rz12b      swap1         ufs     rw    0     2
/dev/rz12c      /HP5700_LU2c  ufs     rw    1     3
```

Parameter	Description	Enter
1	Device to mount	Block-type device file name
2	Mount point	Mount directory name
3	File system	Type of file system (for example, UFS, AdvFS)
4	Mount options	Options (that is, rw for read/write)
5	Frequency	The **dump** command controls file system back up. Specifying a 1 backs up the file system. Specifying a 0 prevents back up.
6	fsck order (File System Check—For UFS systems only, you must fsck disks that contain mount points before mounting other disks on those mount points.)	Block-type device file name.

2. Shut down and reboot the system.
3. Use the **df** or **df –k** command to verify file system auto mounting.

Checking path failover

The disk array supports Tru64 path failover (**Tru64 v5 only**). You can connect multiple HBAs to the disk array with shared LUNs. Confirm the existence of multiple paths for devices as follows:

1. Type **hwmgr -view device** to obtain the HWID for the device.
2. Type **hwmgr -show scsi -full -id HWID** to confirm the status of paths to the device.

If more than one path is currently connected, the status of each path shows as **Valid**. If you change the cabling configuration, the old paths will show as **Stale**. Use the **hwmgr -refresh** command to remove the stale paths.

Caution

Tru64 verson 5 fully supports dynamic load balancing and failover when multiple HBAs are connected to the same LUN. This feature is ONLY available with version 5.

On Tru64 version 4, DO NOT connect multiple HBAs to the same LUN and DO NOT allow the same LUN to be accessible from multiple HBAs on Tru64 version 4 hosts. ON the version 4 system, this will allow multiple device names to access the same LUN. If both device names are used, data corruption will result.

2

Troubleshooting

If you encounter an error condition, see “Error conditions” on page 38 for recommended actions.

If you are unable to resolve an error condition, ask your HP support representative for assistance. See “Calling the HP support center” on page 40.

Error conditions

Error Condition	Recommended Action
The logical devices are not recognized by the host.	Verify that the READY indicator lights on the disk array are ON. Verify that fiber cables are correctly installed and firmly connected. Verify that the target IDs are properly configured. The LUNs for each TID must start at 0 and continue sequentially without skipping any numbers. Verify that the TIDs/WWNs on each bus are unique. Do not install two devices with the same ID on the same bus. Recheck the buses for new devices. Verify that LUSE devices are not intermixed with normal LUNs on the same port. Verify that the maximum number of LUSE devices per port has not been exceeded. Verify that the disk array Host Mode is set correctly. For Tru64 v4, the emx_data.c mappings can cause devices not to be recognized. Verify the data in the emx_data.c file correctly maps each WWN to a TID between 0 and 7. For Tru64 v4, make sure TID 8 and above are not used.
The host does not reboot properly after hard shutdown.	If you power off the host without executing the shutdown process, wait three minutes to allow the disk array's internal timeout process to purge queued commands. If the host restarts while the disk array is processing queued commands, the host may not reboot successfully.
Physical volumes cannot be created.	Verify that the disk array logical devices are correctly formatted.

Error Condition	Recommended Action
Logical volumes cannot be created.	Verify that the volume capacity for OPEN-x volumes is not greater than the maximum capacity allowed. See Appendix B on page 43 for capacities. Verify that the capacity of the volume group is not less than the total capacity of the partitioned logical volume.
A file system cannot be created.	Verify that logical volume name is a character-type volume.
A file system is not mounted after rebooting.	Verify that the host was restarted correctly. Verify that the file system attributes are correct.
The disk array performs a self reboot because the disk array was busy or it logged a panic message.	Reboot the host.
The disk array responds “Not Ready” or the disk array has displayed “Not Ready” and timed out.	Contact HP.
The host detects a parity error.	Check the HBA and make sure it was installed properly. Reboot the host.
The host hangs or devices are declared and the host hangs.	Make sure there are no duplicate disk array TIDs and that disk array TIDs do not conflict with any host TIDs.

Calling the HP support center

If you need to call HP customer support, provide as much information about the problem as possible, including the circumstances of the error or failure and the exact content of any error messages.

Depending on your system configuration, you may be able to view error messages as follows:

- View SIMs in Command View (Device Health tab).
- View R-SIMs in Remote Control XP, including reference codes and severity levels of recent R-SIMs.
- View SIMs that generate SNMP traps on the host.

A

Worksheet

Path worksheet

LDEV (CU:LDEV) (CU = control unit)	Device Type	SCSI Bus Number	Path 1	Alternate Paths		
0:00				TID: LUN:	TID: LUN:	TID: LUN:
0:01				TID: LUN:	TID: LUN:	TID: LUN:
0:02				TID: LUN:	TID: LUN:	TID: LUN:
0:03				TID: LUN:	TID: LUN:	TID: LUN:
0:04				TID: LUN:	TID: LUN:	TID: LUN:
0:05				TID: LUN:	TID: LUN:	TID: LUN:
0:06				TID: LUN:	TID: LUN:	TID: LUN:
0:07				TID: LUN:	TID: LUN:	TID: LUN:
0:08				TID: LUN:	TID: LUN:	TID: LUN:
0:09				TID: LUN:	TID: LUN:	TID: LUN:
0:10				TID: LUN:	TID: LUN:	TID: LUN:
0:11				TID: LUN:	TID: LUN:	TID: LUN:
0:12				TID: LUN:	TID: LUN:	TID: LUN:
0:13				TID: LUN:	TID: LUN:	TID: LUN:

B

Disk array device emulations

This appendix provides information about disk array supported emulations and device type specifications. Some parameters may not be relevant to your array. Consult your HP representative for information about supported configurations for your system.

Supported emulations

XP Type	Emulation	OPEN-x	LUSE	CVS	LUSE & CVS
XP48 XP512	OPEN-3	Yes	Yes	Yes	Yes
	OPEN-8	Yes	Yes	Yes	Yes
	OPEN-9	Yes	Yes	Yes	Yes
	OPEN-E	Yes	Yes	Yes	Yes
	OPEN-K	Yes	Yes	Yes	Yes
	OPEN-L	Yes	Yes		
	OPEN-M	Yes	Yes		
	OPEN-V				
XP128 XP1024 XP12000	OPEN-3	Yes	Yes	Yes	Yes
	OPEN-8	Yes	Yes	Yes	Yes
	OPEN-9	Yes	Yes	Yes	Yes
	OPEN-E	Yes	Yes	Yes	Yes
	OPEN-K				
	OPEN-L	Yes	Yes		
	OPEN-M				
	OPEN-V	Yes	Yes		

Device type specifications

Device Type *(Note 1)*	Category *(Note 2)*	Blocks (512 bytes)	Sector Size (bytes)	# of Cylinders	Heads	Sectors per Track	Capacity MB* *(Note 3)*
OPEN-3	SCSI disk	4806720	512	3338	15	96	2347
OPEN-8	SCSI disk	14351040	512	9966	15	96	7007
OPEN-9	SCSI disk	14423040	512	10016	15	96	7042
OPEN-E	SCSI disk	28452960	512	19759	15	96	13893
OPEN-L	SCSI disk	71192160	512	49439	15	96	34761
OPEN-V	SCSI disk	max=125827200	512	Note 5	15	128	Note 6
LUSE							
OPEN-3*n	SCSI disk	4806720*n	512	3338*n	15	96	2347*n
OPEN-8*n	SCSI disk	14351040*n	512	9966*n	15	96	7007*n
OPEN-9*n	SCSI disk	14423040*n	512	10016*n	15	96	7042*n
OPEN-E*n	SCSI disk	28452960*n	512	19759*n	15	96	13893*n
OPEN-L*n	SCSI disk	71192160*n	512	49439*n	15	96	34761*n
OPEN-V*n	SCSI disk	max=125827200 Note 4	512	Note 5	15	128	Note 6
CVS							
OPEN-3 CVS	SCSI disk	Note 4	512	Note 5	15	96	Note 6
OPEN-8 CVS	SCSI disk	Note 4	512	Note 5	15	96	Note 6
OPEN-9 CVS	SCSI disk	Note 4	512	Note 5	15	96	Note 6
OPEN-E CVS	SCSI disk	Note 4	512	Note 5	15	96	Note 6
CVS LUSE							
OPEN-3*n CVS	SCSI disk	Note 4	512	Note 5	15	96	Note 6
OPEN-8*n CVS	SCSI disk	Note 4	512	Note 5	15	96	Note 6
OPEN-9*n CVS	SCSI disk	Note 4	512	Note 5	15	96	Note 6
OPEN-E*n CVS	SCSI disk	Note 4	512	Note 5	15	96	Note 6
OPEN-V*n	SCSI disk	Note 4	512	Note 5	15	128	Note 6

*Capacity = (512 x number of blocks) ÷ 1024^2

Note 1: The availability of a disk type depends on the disk array.

Note 2: The devices are defined to the host as SCSI disk devices, even though the interface is Fibre Channel.

Note 3 The device capacity can sometimes be changed by the BIOS or host adapter board. This may make actual capacity different from that listed in the table.

Note 4: The number of blocks for a CVS volume is calculated as follows:
# of blocks = (# of cylinders) × (# of heads) × (# of sectors per track)

Example 1: For an OPEN-3 CVS volume with capacity = 37 MB:
# of blocks = (53 cylinders–see Note 5) × (15 heads) × (96 sectors per track) = 76320

Example 2: For an OPEN-V CVS volume with capacity = 49 MB:
# of blocks = (53 cylinders–see Note 5) × (15 heads) × (128 sectors per track) = 101760

Note 5: The number of cylinders for a CVS volume is calculated as follows (↑ …↑ means that the value should be rounded up to the next integer):

OPEN-3/8/9/E: The number of cylinders for a CVS volume =
# of cylinders = ↑ (capacity (MB) specified by user) × 1024/720 ↑

Example: For an OPEN-3 CVS volume with capacity = 37 MB:
# of cylinders = ↑ 37 × 1024/720 ↑ = ↑ 52.62 ↑ (rounded up to next integer) = 53 cylinders

OPEN-V: The number of cylinders for a CVS volume =
# of cylinders = ↑ (capacity (MB) specified by user) × 16/15 ↑

Example: For an OPEN-V CVS volume with capacity = 49 MB:
# of cylinders = ↑ 49 × 16/15 ↑ = ↑ 52.26 ↑ (rounded up to next integer) = 53 cylinders

OPEN-3/8/9/E: The number of cylinders for a CVS LUSE volume =
# of cylinders = ↑ (capacity (MB) specified by user) × 1024/720 ↑ × n

Example: For a CVS LUSE volume with capacity = 37 MB and n = 4
# of cylinders = ↑ 37 × 1024/720↑ × 4 = ↑ 52.62↑ × 4 = 53 × 4 = 212

OPEN-V: The number of cylinders for a CVS LUSE volume =
# of cylinders = ↑ (capacity (MB) specified by user) × 16/15 ↑ × n

Example: For an OPEN-V CVS LUSE volume with capacity = 49 MB and n = 4
# of cylinders = ↑ 49 × 16/15↑ × 4 = ↑ 52.26↑ × 4 = 53 × 4 = 212

Note 6: The capacity of an OPEN-3/8/9/E CVS volume is specified in MB, not number of cylinders. The capacity of an OPEN-V CVS volume can be specified in MB or number of cylinders. You set the volume size using the LUN Configuration Manager or Command View software.

Glossary

AL	Arbitrated loop.
AL-PA	Arbitrated loop physical address.
BC	HP StorageWorks Business Copy XP. BC lets you maintain up to nine local copies of logical volumes on the disk array.
CA	HP StorageWorks Continuous Access XP. CA lets you create and maintain duplicate copies of local logical volumes on a remote disk array.
Command View	HP StorageWorks Command View XP, a software product for managing XP arrays. Command View runs on a Windows-based management workstation.
command device	An LDEV that transfers RAID commands to BC or CA logical volumes.
CVS	CVS devices (OPEN-x CVS) are custom volumes that are smaller than normal fixed-sized logical disk devices (volumes).
DKC (disk controller unit)	The array cabinet that houses the channel adapters and service processor (SVP).
DKU (disk cabinet unit)	The array cabinets that house the disk array physical disks.
emulation modes	Emulation modes can be assigned to LDEVs to make them operate like standard OPEN system disk drives. The emulation mode of an LDEV determines its capacity. Refer to the appendices for device capacities.
FC	Fibre Channel.

FC-AL	Fibre Channel arbitrated loop.
FCP	Fibre Channel Protocol.
HBA	Host bus adapter.
HP	Hewlett-Packard Company.
LDEV	Logical device. An LDEV is created when a RAID group is divided into sections using a selected host emulation mode (for example, OPEN-9 or OPEN-M). The number of resulting LDEVs depends on the emulation mode. "LDEV" and "volume" are synonyms.
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.
LUSE	Logical Unit Size Expansion, a feature which logically combines LDEVs so they appear as a larger LDEV. This allows a LUN to be associated with 2 to 36 LDEVs. LUSE allows applications to access data requiring large amounts of disk space.
OFC	Open Fibre Control.
OPEN-*x*	A general term describing any one of the supported OPEN emulation modes (for example, OPEN-L).
OS	Operating system.
PA	Physical address.
path	"Path" and "LUN" are synonymous. Paths are created by associating a port, a target, and a LUN ID with one or more LDEVs.
port	A connector on a channel adapter card in the disk array. A port passes data between the disk array and external devices, such as a host server. Ports are named using a port group and port letter, for example, CL1-A.

RAID	Redundant array of independent disks.
remote console PC	The PC running HP StorageWorks Remote Control XP.
Remote Control (RC)	HP StorageWorks Remote Control XP. A software product used for managing XP arrays.
R-SIM	Remote service information message.
SCSI	Small computer system interface.
SIM	Service information message.
SNMP	Simple Network Management Protocol.
SVP	Service processor. A notebook computer built into the disk array. The SVP provides a direct interface to the disk array and is used only by the HP service representative.
TID	Target ID.
VSC	Volume Size Configuration is a feature that defines custom volumes (CVS volumes) that are smaller than normal fixed-sized logical disk devices (volumes).
WWN	World Wide Name. A unique identifier assigned to a Fibre Channel device.

Index

W

X

Z