HP StorageWorks

XP Disk Array Configuration Guide for HP-UX

HP XP24000 Disk Array
HP XP20000 Disk Array
HP XP12000 Disk Array
HP XP10000 Disk Array
HP 200 Storage Virtualization System

Abstract

This guide provides requirements and procedures for connecting an XP disk array or SVS 200 to a host system, and for configuring the disk array for use with the HP-UX operating system. This document is intended for system administrators, HP representatives, and authorized service providers who are involved in installing, configuring, and operating the HP XP storage systems.

Part number: A5951-96083
Thirteenth edition: January 2010

Contents

Figures

Tables

1 Installing and configuring HP-UX for the XP disk array

You and the HP service representative each play a role in installation. The HP service representative is responsible for installing the disk array and formatting the disk devices. You are responsible for configuring the host server for the new devices with assistance from the HP service representative.

Features and requirements

The disk array has the following features:

- **Storage capacity:** The storage capacity for each model is listed below:

Model	Maximum drives	Maximum Capacity	Maximum FC ports
XP24000	1152	2.27TB	224 4GB or 24 8GB
XP20000	240	472TB	48 4GB or 24 8GB
XP12000	1152	454TB	224 4GB
XP10000	240	94.5TB	48 4GB
SVS 200	External	124TB	48

- **Server support:** HP-UX-supported processor
- For supported disk array microcode and OS versions, see the HP StorageWorks Single Point of Connectivity Knowledge (SPOCK):
  http://spock.corp.hp.com

Before installing the disk array, ensure the following requirements are met:

- **Fibre Channel Adapters (FCAs):** Install FCAs and all utilities and drivers. Refer to the adapter documentation for installation details.
- HP StorageWorks XP Remote Web Console, HP StorageWorks XP Command View Advanced Edition Software or HP StorageWorks XP Command View with LUN management feature for configuring disk array ports and paths.
- HP StorageWorks XP Array Manager.
- Check with your HP representative for other XP software available for your system.

Fibre Channel interface

The XP family of disk arrays and the SVS 200 support these Fibre Channel elements:

- Connection speeds of 1 Gbps, 2 Gbps, and 4 Gbps

- 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 direct fabric 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 emulation types

The XP family of disk arrays and the SVS 200 support these device emulation 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 based on a CVS device. 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) devices 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 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.) Although OPEN-V is a CVS-based device, the product name in the SCSI inquiry string is OPEN-V opposed to the fixed size OPEN-[389E] devices that appear as OPEN-x-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. OPEN-V devices are designated as OPEN-V*n (without CVS).

NOTE:

For the SVS 200, and the XP24000/XP20000/XP12000/XP10000 when connected to external storage devices, HP recommends using OPEN-V as the emulation the array makes visible to the host. This allows configuration of external storage LDEVs without losing data. Using any other emulation may cause data loss in the external storage LUNs. For new deployments, OPEN-V should be chosen because some features (such as HP StorageWorks XP Snapshot Software or HP StorageWorks XP Continuous Access Journal Software) are only supported with OPEN-V.

For detailed information, see “Emulation specifications” on page 37.

Failover

The disk arrays support many standard software products that provide host, application, or I/O path failover and management.

The following are supported for HP-UX:

- HP Multi-Computer/ServiceGuard (MC/ServiceGuard) software for application failover
- Alternate link for I/O path failover (included in HP-UX)
- Logical volume management (included in HP-UX)

SNMP configuration

The XP family of disk arrays and the SVS 200 support standard Simple Network Management Protocol (SNMP) for remotely managing arrays. The SNMP agent on the service processor (SVP) performs error-reporting operations requested by the SNMP manager. SNMP properties are usually set from the SVP but they can also be set remotely using XP Remote Web Console or XP Command View Advanced Edition Software. Refer to the applicable user's guide for procedures.

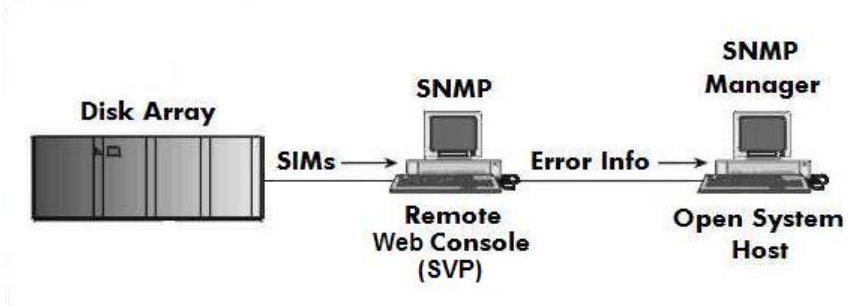

Figure 1 SNMP configuration

XP RAID Manager command devices

HP StorageWorks XP RAID Manager manages HP StorageWorks XP Business Copy Software or HP StorageWorks XP Continuous Access Software operations from a host server. To use RAID Manager you must designate at least one LDEV as a command device. This can be done with Remote Web Console or XP Command View Advanced Edition Software. Refer to the applicable user guide for information about how to designate a command device.

Creating scripts to configure all devices at once may save you considerable time.

Installation roadmap

Perform the following actions to install and configure the disk array:

1. Installing and configuring the disk array
   - Setting the host mode and host group mode for the disk array ports
   - Setting the system option modes
   - Configuring the Fibre Channel ports
2. Installing and configuring the host
   - Loading the operating system and software
   - Installing and configuring the FCAs
   - Clustering and fabric zoning
   - Fabric zoning and LUN security for multiple operating systems
3. Connecting the disk array
   - Defining the paths
   - Verifying FCA installation
   - Verifying device recognition

4. Configuring disk array devices
   - Verifying the device files and drivers
   - Creating the device files
   - Creating the physical volumes
   - Creating new volume groups
   - Creating logical volumes
   - Creating the file systems
   - Setting the I/O timeout parameter
   - Creating the mount directories
   - Mounting and verifying the file systems
   - Setting and verifying the auto-mount parameters

Installing and configuring the disk array

The HP service representative performs these tasks:

- Assembling hardware and installing software
- Loading the microcode updates
- Installing and formatting devices

After these tasks are finished, you will use XP Remote Web Console, XP Command View Advanced Edition Software, or XP Array Manager to complete the remaining tasks listed below. If you do not have these programs, your HP service representative can perform these tasks for you.

Setting the host mode and host group mode for the disk array ports

After the disk array is installed, you must set the host mode for each host group that is configured on a disk array port to match the host OS. Set the host mode using LUN Manager in XP Remote Web Console (shown) or XP Command View Advanced Edition Software. If these are not available, the HP service representative can set the host mode using the SVP. The host mode setting for HP-UX is **08**.

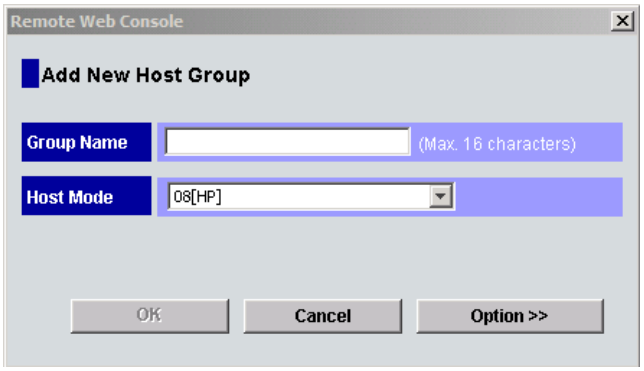

△ CAUTION:

The correct host mode must be set for all new installations (newly connected ports) to HP-UX hosts. Do not select a mode other than **08** for HP-UX. Changing a host mode after the host has been connected is disruptive and requires the server to be rebooted.

When a new host group is added, additional host group modes (options) may be configured. The storage administrator must verify if an additional host group mode is required for the host group.

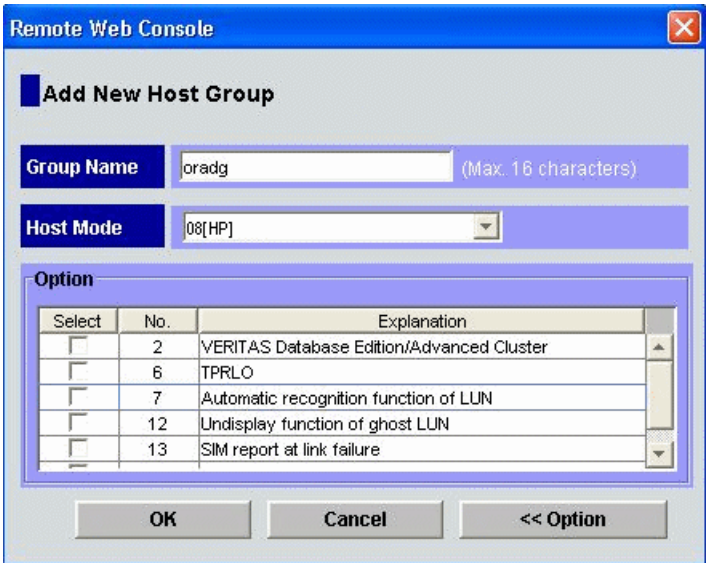

The following host group modes (options) are available for HP-UX:

Table 1 Host group modes (options)

Host Group Mode	Function	Default	Comments
12	Deletion of Ghost LUN	Inactive	Previously MODE280
33	Task retry ID enable	Inactive	*HP-UX 11.31 only*

△ CAUTION:

Changing host group modes for ports where servers are already installed and configured is disruptive and requires the server to be rebooted.

Setting the system option modes

The HP service representative sets the system option mode(s) based on the operating system and software configuration of the host. Notify your HP representative if you install storage agnostic software (such as backup or cluster software) that may require specific settings.

Configuring the Fibre Channel ports

Configure the disk array Fibre Channel ports by using XP Remote Web Console (shown) or XP Command View Advanced Edition Software. 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.

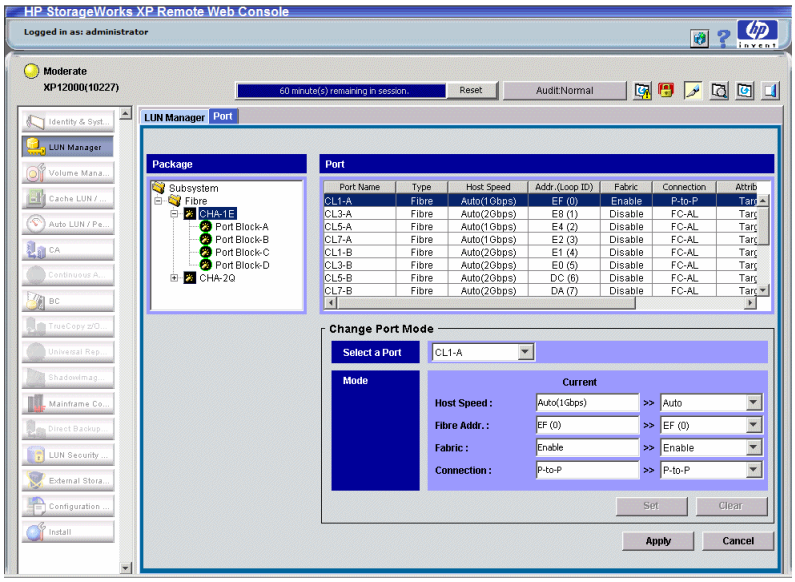

Fibre address

In fabric environments, the port addresses are assigned automatically. In arbitrated loop environments, set the port addresses by selecting a unique arbitrated loop physical address (AL-PA) or loop ID for each port. For specific values, refer to the *HP StorageWorks XP LUN Manager User's Guide* applicable to your array.

Fabric and connection parameter settings

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* at:

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

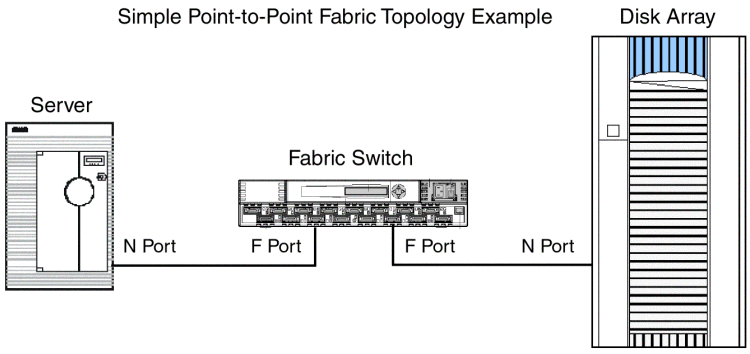

Figure 2 Point-to-point fabric topology example

Table 2 Fabric topology settings

Fabric parameter	Connection parameter	Provides
OFF	FC-AL	NL-port (private arbitrated loop)
ON	Direct Fabric Attach	F-port (fabric port)

Fabric parameter	Connection parameter	Provides
ON	FC-AL	Not supported
OFF	Direct Fabric Attach	Not supported

Installing and configuring the host

This section explains how to install and configure Fibre Channel adapters (FCAs) that connect the host to the disk array.

Loading the operating system 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 FCA manufacturer.

Installing and configuring the FCAs

Install and configure the Fibre Channel adapters using the FCA manufacturer's instructions.

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.

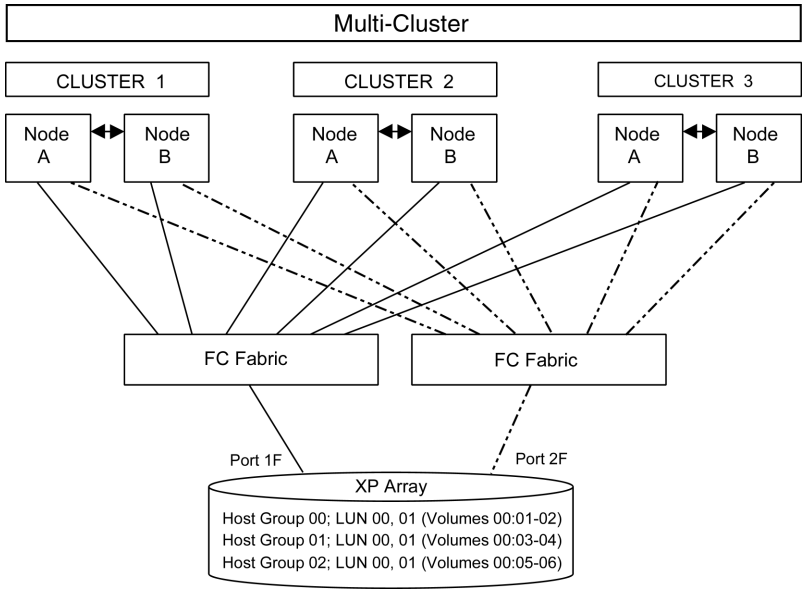

Figure 3 Multi-cluster environment

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.

⚠ WARNING!

For OpenVMS — HP recommends that a volume be presented to one OpenVMS cluster or stand alone system at a time. Volumes should not be presented to allow them to move between stand alone systems and/or OpenVMS clusters, as this can lead to corruption of the OpenVMS volume and data loss.

Fabric zoning and LUN security for multiple operating systems

You can connect multiple clusters of various operating systems to the same switch and fabric using appropriate zoning and LUN security as follows:

- 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.

Table 3 Fabric zoning and LUN security settings

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

Connecting the disk array

The HP service representative connects the disk array to the host by:

1. Verifying operational status of the disk array channel adapters, LDEVs, and paths.
2. Connecting the Fibre Channel cables between the disk array and the fabric switch or host.
3. Verifying the ready status of the disk array and peripherals.

Defining the paths

Use XP Remote Web Console (shown) or XP Command View Advanced Edition Software to define paths (LUNs) between hosts and volumes in the disk array.

This process is also called “LUN mapping.” In XP Remote Web Console and XP Command View, LUN mapping includes:

- Configuring ports
- Enabling LUN security on the ports
- Creating host groups
- Assigning Fibre Channel adapter WWNs to host groups
- Mapping volumes (LDEVs) to host groups (by assigning LUNs)

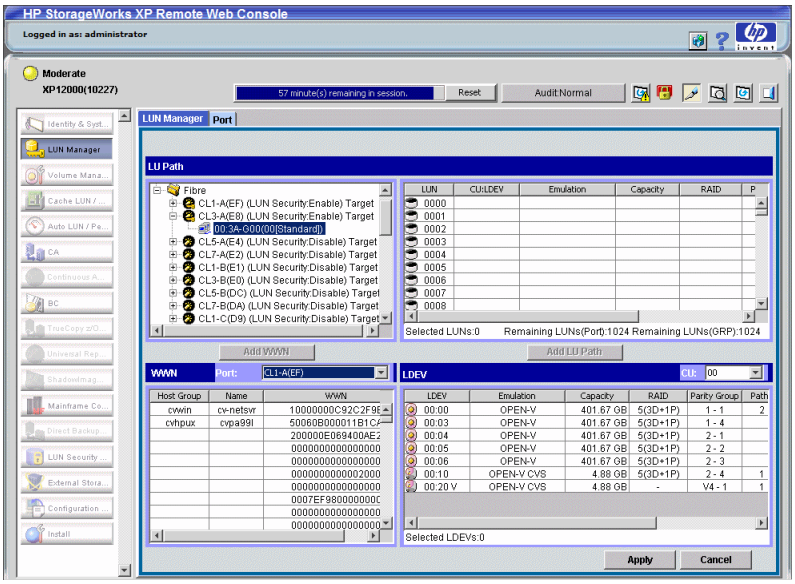

In XP Command View Advanced Edition Software, LUN mapping includes:

- Configuring ports
- Creating storage groups
- Mapping volumes and WWN/host access permissions to the storage groups

For details see the *HP StorageWorks XP LUN Manager User's Guide* or the *HP StorageWorks XP Command View Advanced Edition Software Device Manager Web Client User's Guide*. Note the LUNs and their ports, WWNs, nicknames, and LDEVs for later use in verifying host and device configuration.

Verifying FCA installation

After configuring the ports on the disk array, verify that the FCAs are installed properly.

Use the `ioscan -f` command, and verify that the rows shown in the example are displayed. If these rows are not displayed, check the host adapter installation (hardware and driver installation) or the host configuration.

Example

```
# ioscan -f
Class     I H/W Path             Driver     S/W State H/W Type  Description
...
fc        0 8/12                 fcT1       CLAIMED   INTERFACE HP Fibre ...
lan       1 8/12.5               fcT1_cntl  CLAIMED   INTERFACE HP Fibre ...
fcp       0 8/12.8               fcp        CLAIMED   INTERFACE FCP Proto...
ext bus   2 8/12.8.0.255.0       fcpdev     CLAIMED   INTERFACE FCP Devic...
```

Verifying device recognition

Verify that the HP-UX system recognizes the new devices on the disk array.

If the SCSI paths were defined after the system is powered on, you must halt and restart the system to allow the system to recognize the new devices.

To verify device recognition:

1. Log in to the system as root.
2. Display the device data to verify that the system recognizes the newly installed devices on the disk array. Use the `ioscan -fn` command to display the device data.

   On a system with a large LUN configuration, HP-UX may not build device files on all LUNs. Enter `insf -e` to build all missing device files.

   Example

```
# ioscan -fn
Class     I H/W Path             Driver     S/W State H/W Type...
bc        6 14                   ccio       CLAIMED   BUS_NEXUS...
fc        1 14/12                fcT1       CLAIMED   INTERFACE...
lan       2 14/12.5              fcT1_cntl  CLAIMED   INTERFACE...
fcp       1 14/12.8              fcp        CLAIMED   INTERFACE...
ext_bus   6 14/12.8.0.0.0        fcpmux     CLAIMED   INTERFACE...
disk      4 14/12.8.0.0.0.0.0    sdisk      CLAIMED   DEVCE...
disk      5 14/12.8.0.0.0.0.1    sdisk      CLAIMED   DEVICE...
ext_bus   7 14/12.8.0.255.0      fcpdev     CLAIMED   INTERFACE...
target   10 14/12.8.0.255.0.0    tgt        CLAIMED   DEVICE...
ctl       5 14/12.8.0.255.0.0.0  sctl       CLAIMED   DEVICE...
```

   In the example:

   - HP OPEN-9 device: SCSI bus number = 14/12, bus instance = 6, SCSI target ID = 0, LUN = 0.
   - HP OPEN-9*2 device: SCSI bus number = 14/12, bus instance = 6, SCSI target ID = 0, LUN = 1.
   - If UNKNOWN is displayed for a disk, the HP 9000 system may not be configured properly. Refer to the HP documentation or contact HP customer support for assistance with the HP 9000 system or the HP-UX operating system.
3. Enter the device data for each disk array device in a table. See Appendix A on page 35.

4. Construct the device file name for each device, using the device information, and enter the file names in your table. Use the following formula to construct the device file name:

   `cxtydz`

   where:

   - x = SCSI bus instance number
   - y = SCSI target ID
   - z = LUN
   - c stands for controller
   - t stands for target ID
   - d stands for device

   The numbers x, y, and z are hexadecimal.

   Table 4 Device file name example

SCSI bus instance number	Hardware path	SCSI TID	LUN	File name
00	14/12.6.0	6	0	c6t0d0
00	14/12.6.1	6	2	c6t0d1

5. Verify that the SCSI TIDs correspond to the assigned port address for all connected ports (see mapping tables in Appendix B on page 37, for values). If so, the logical devices are recognized properly.

   If the logical devices are not recognized properly:

   - Check the AL-PA for each port using the LUN Manager software.
   - If the same port address is set for multiple ports on the same loop (AL with HUB), all port addresses except one changed to another value, and the relationship between AL-PA and TID does not correspond to the mapping given in Appendix B on page 37, set a different address for each port, reboot the server, and then verify new device recognition again.
   - If unused device information remains, the TID-to-AL-PA mapping will not correspond to the mapping given in Appendix B on page 37. Renew the device information, and then verify new device recognition again.

Configuring disk array devices

Disk arrays are configured using the same procedure for configuring any new disk on the host. This includes the following procedures:

1. Verifying the device files and drivers
2. Creating the device files
3. Creating the physical volumes
4. Creating new volume groups
5. Creating logical volumes
6. Creating the file systems
7. Setting the I/O timeout parameter
8. Creating the mount directories

9. Mounting and verifying the file systems
10. Setting and verifying the auto-mount parameters

The HP-UX system uses the Logical Volume Manager (LVM) to manage the OPEN-*x* devices on the disk array. The instructions in this section do not explicitly cover all LVM configuration issues. For further information on LVM configuration, see the HP-UX user documentation.

HP System Administrator Manager (SAM) can be used instead of UNIX commands to configure SCSI disk devices. See "Reference information for the HP System Administrator Manager SAM" on page 45 for further information.

Verifying the device files and drivers

The device files for new devices are usually created automatically during HP-UX startup. Each device must have a block-type device file in the `/dev/dsk` directory and a character-type device file in the `/dev/rdsk` directory.

However, some HP-compatible systems do not create the device files automatically. If verification shows that the device files were not created, follow the instructions in "Creating the device files" on page 19.

The following procedure verifies both types of device files:

1. Display the block-type device files in the `/dev/dsk` directory using the `ls -l` command with the output piped to `more`. Verify there is one block-type device file for each disk array device.

   Example

   ```
   # ls -l /dev/dsk | more

   Total 0

   brw-r - - - - - 1 bin sys 28 0x006000 Dec 6 15:08 c6t0d0
   brw-r - - - - - 1 bin sys 280 0x06100 Dec 6 15:08 c6t0d1
   ```

2. Verify that the block-type device file name for each device is correct.
3. Display the character-type device files in the `/dev/rdsk` directory using the `ls -l` command with the output piped to `more`. Verify that there is one character-type device file for each disk array device.

   Example

   ```
   # ls -l /dev/rdsk | more

   Total 0

   crw-r - - - - - 1 bin sys 177 0x006000 Dec 6 15:08 c6t0d0
   crw-r - - - - - 1 bin sys 177 0x006100 Dec 6 15:08 c6t0d1
   ```

4. Use the device data table you created to verify that the character-type device file name for each device is correct.

   This task can also be accomplished with the `lssf` command.

5. After verifying the block-type and character-type device files, verify the HP-UX driver for the disk array using the `ioscan -fn` command.

   Example

```
# ioscan -fn
Class   I H/W Path           Driver S/W State H/W Type  Desc
-------------------------------------------------------------
bc      0                   root   CLAIMED   BUS_NEXUS...
bc      1 8                 bc     CLAIMED   BUS_NEXUS...
fc      0 8/12              fcT1   CLAIMED   INTERFACE...
fcp     0 8/12.8            fcp    CLAIMED   INTERFACE...
ext_bus 2 8/12.8.0.255.0    fcpdev CLAIMED   INTERFACE...
disk    3 8/12.8.8.255.0.6.0 sdisk  CLAIMED   DEVICE...
                /dev/dsk/c2t6d0    /dev/rdsk/c2t6d0
disk    4 8/12.8.8.255.0.6.1 sdisk  CLAIMED   DEVICE...
                /dev/dsk/c2t6d1    /dev/rdsk/c2t6d1
disk    5 8/12.8.8.255.0.8.0 sdisk  CLAIMED   DEVICE...
                /dev/dsk/c2t8d0    /dev/rdsk/c2t8d0
```

Creating the device files

If the device files were not created automatically when the system was restarted, use the `insf -e` command in the `/dev` directory to create the device files. After this command is executed, repeat the procedures in "Verifying device recognition" on page 15 to verify new device recognition and the device files and driver.

Example

```
# insf -e

insf: Installing special files for mux2 instance 0 address 8/0/0
  :               :               :               :
  :               :               :               :
#
```

Failure of the `insf -e` command indicates a SAN problem.

If the device files for the new disk array devices cannot be created automatically, you must create the device files manually using the `mknodcommand` as follows:

1. Retrieve the device information you recorded earlier.
2. Construct the device file name for each device, using the device information, and enter the file names in your table. Use the following formula to construct the device file name:

   `cxtydz`

   where:

   - x = SCSI bus instance number
   - y = SCSI target ID
   - z = LUN
   - c stands for controller
   - t stands for target ID
   - d stands for device

   The numbers x, y, and z are hexadecimal.

3. Construct the minor number for each device, using the device information, and enter the file names in your table. Use the following formula to construct the minor number:

   `0xxxyz00` where

   xx = SCSI bus instance number, y = SCSI target ID, and z = LUN.

4. Display the driver information for the system using the `lsdev` command.

   Example

```
# lsdev

Character   Block   Driver   Class
   :          :       :        :
  188        31     sdisk    disk
 #
```

5. Enter the major numbers for the device drivers into the table. You should now have all required device and driver information in the table.

6. Create the device files for all disk array devices (SCSI disk and multiplatform devices) using the `mknodcommand`. Create the block-type device files in the `/dev/dsk` directory and the character-type device files in the `/dev/rdsk` directory.

   Example

```
# cd /dev/dsk                              Go to /dev/dsk directory.

# mknod /dev/dsk/c2t6d0 b 31 0x026000   Create block-type file.
                                         File name, b=block-type,
                                         31=major #, 0x026000= minor #

# cd /dev/rdsk                             Go to /dev/rdsk directory.

# mknod /dev/rdsk/c2t6d0 c 188 0x026000 Create character-type file.
                                         File name, c=character-type,
                                         177=major #, 0x026000=minor #
:
#
```

   The character-type device file is required for volumes used as raw devices (for example, 3390-3A/B/C). The block-type device file is not required for volumes used as raw devices.

   If you need to delete a device file, use the `rm –i` command.

Table 5 Device information example

Bus	Inst	Disk	HW path	Driver	Dev type	TID	LUN	Dev file	Minor no.	Major no. char. files	Major no. block files
8/12	02	3	8/12.8.8.255.0.6.0	sdisk	OPEN-9	6	0	c2t6d0	0x026000	188	31
8/12	02	4	8/12.8.8.255.0.6.1	sdisk	OPEN-9	6	1	c2t6d1	0x026100	188	31
8/12	02	5	8/12.8.8.255.0.8.0	sdisk	3390-3B	8	0	c2t8d0	0x028000	188	31

Creating the physical volumes

A physical volume must be created for each new SCSI disk device.

To create the physical volumes:

1. Use the `pvcreate` command to create the physical volumes with the character-type device file as the argument. Specify the `/dev/rdsk` directory.

   Example

   ```
   # pvcreate /dev/rdsk/c6t0d0

   Physical volume "/dev/rdsk/c6t0d0" has been successfully created.
   :
   # pvcreate /dev/rdsk/c6t0d1

   Physical volume "/dev/rdsk/c6t0d1" has been successfully created.
   ```

   Do not use the `-f` option with the `pvcreate` command. This option creates a new physical volume forcibly and overwrites the existing volume. If you accidentally enter the character-type device file for an existing volume, you will lose the data on that volume.

2. Repeat step 1 for each OPEN-*x* device on the disk array.

Creating new volume groups

You must create new volume groups for the new physical volumes. If desired, you can also add any of the volumes on the disk array to existing volume groups using the `vgextend` command. The physical volumes that make up one volume group can be located either in the same disk array or in other disk arrays.

To allow more volume groups to be created, use SAM to modify the HP-UX system kernel configuration. See "Reference information for the HP System Administrator Manager SAM" on page 45 for details.

To create volume groups:

1. Use the `vgdisplay` command to display the existing volume groups.
2. Choose a unique name for the new volume group (for example: vg06).
3. Create the directory for the new volume group.

   Example

   ```
   # mkdir /dev/vg06
   ```

4. Use the `ls -l` command (with the output piped to `grep` to display only the files containing "group") to display the minor numbers for the existing group files.

   Example

   ```
   # ls -1 /dev/vg* | grep group

   crw-rw-rw 1 root root 64 0x0000000 Nov7 08:13 group
   :
   ```

5. Choose a unique minor number for the new group file in sequential order (for example, when existing volume groups are vg00-vg05 and the next group name is vg06, use minor number 06 for the vg06 group file).

   The minor numbers are hexadecimal (for example, the 10th minor number is 0x0a0000).

6. Use `mknod` to create the group file for the new directory. Specify the volume group name, major number, and minor number. The major number for all group files is 64.

   Example

   In this example: group name = vg06, major number of group file = 64, minor number of existing group file = 06 (which must be unique for each volume group), and c = character.

   ```
   # mknod /dev/vg06/group c 64 0x060000
   :
   ```

7. Create the volume group.

   To allocate more than one physical volume to the new volume group, add the other physical volumes, separated by a space.

   Example

   ```
   # vgcreate /dev/vg06 /dev/dsk/c6t0d0

   Volume group "/dev/vg06" has been successfully created.

   Volume group configuration for /dev/vg06 has been saved in
   /etc/1vmconf/vg06.conf.
   ```

   For Logical Unit Size Expansion (LUSE) volumes that contain more than 17 OPEN-8/9 LDEVs or more than 7043 MB (OPEN-8/9*n-CVS), use the `-s` and `-e` physical extent (PE) parameters of the `vgcreate` command. See "LUSE device parameters" on page 40.

   If you need to delete a volume group, use the `vgremove` command (for example, `vgremove /dev/vgnn`). If the `vgremove` command does not work because the volume group is not active, use the `vgexport` command (for example, `vgexport /dev/vgnn`).

8. Use the `vgdisplay` command to verify that the new directory was created.

9. Use `vgdisplay –v` to verify that the volume group was created correctly. The `–v` option displays the detailed volume group information.

   Example

```
# vgdisplay –v /dev/vg06

- - - Volume groups - - -
VG Name                     /dev/vg06
VG Write Access             read/write
VG Status                   available
Max LV                      255
Cur LV                      0
Open LV                     0
Max PV                      16
Cur PV                      1
Act PV                      1
Max PE per PV               1016
VGDA 2
PE Size (Mbytes)            4
Total PE                    586
Alloc PE                    0
Free PE                     586
Total PVG                   0

- - Physical Volumes - - -
PV Name /dev/dsk/c6t0d0
PV Status available
Total PE 586
Free PE 586
```

Creating logical volumes

Use these commands for logical volume configuration:

- `lvremove`

  Deletes a logical volume. Any file system attached to the logical volume must be unmounted before executing the `lvremove` command.

  Example

```
lvremove /dev/vgnn/lvolx
```

- `lvextend`

  Increases the size of an existing logical volume.

  Example

```
lvextend –L size /dev/vgnn/lvolx
```

- `lvreduce`

  Decreases the size of an existing logical volume. Any file system attached to the logical volume must be unmounted before executing the `lvreduce` command.

  Example

```
lvreduce –L size /dev/vgnn/lvolx
```

△ CAUTION:

Data within the file system can be lost after execution of `lvreduce`.

Create logical volumes after you create volume groups. A logical volume must be created for each new SCSI disk device.

To create logical volumes:

1. Use the `lvcreate –L` command to create a logical volume.

   Specify the volume size (in megabytes) and the volume group for the new logical volume. HP-UX assigns the logical volume numbers automatically (lvol1, lvol2, lvol3). Use the following capacity values for the size parameter:

   OPEN-K = 1740

   OPEN-3 = 2344

   OPEN-8 = 7004

   OPEN-9 = 7004

   OPEN-E = 13888

   OPEN-L = 34756

   OPEN-V = 61432

   To calculate S1 for CVS, LUSE, and CVS LUSE volumes, first use the vgdisplay command to display the physical extent size (PE Size) and usable number of physical extents (Free PE) for the volume. Calculate the maximum size value (in MB) as follows:

   S1 = (PE Size) × (Free PE)

   Logical volumes can span multiple physical volumes. Use the `diskinfo` command for extended LUNs.

2. Create an OPEN-3 logical volume the size of the physical volume, using 2344 for the size parameter. An OPEN-9 volume uses 7040 for the size parameter to create a logical volume the size of the physical volume.

   Example

```
# lvcreate –L 2344 /dev/vg06
Logical volume "/dev/vg06/lvol1" has been successfully created with
   character device "/dev/vg06/rlvol1".
Logical volume "/dev/vg06/lvol1" has been successfully extended.
Volume Group configuration for /dev/vg06 has been saved in
/etc/1vmconf/vg06.conf.
```

3. Use the `lvdisplay` command to verify that the logical volume was created correctly.

   Example

```
# lvdisplay /dev/vg06/lvol1

- - - Logical volume - - -
LV Name                /dev/vg06/lvol1
VG Name                /dev/vg06
LV Permission          read/write
LV Status              available/syncd
Mirror copies          0
Consistency Recovery   MWC
Schedule               parallel
LV Size (Mbytes)       2344
Current LE             586
Allocated PE           586
Stripes                0
Stripe Size (Kbytes)   0
Bad block              on
Allocation             strict
```

4. Repeat above steps for each logical volume to be created.

   You can create only one logical volume at a time. However, you can verify multiple logical volumes at a time.

Creating the file systems

Create the file system for each new logical volume on the disk array. The default file system types are:

- HP-UX OS version 10.20 = `hfs` or `vxfs`, depending on entry in the `/etc/defaults/fs` file.
- HP-UX OS version 11.0 = `vxfs`
- HP-UX OS version 11.i = `vxfs`

To create file systems:

1. Use the newfs command to create the file system using the logical volume as the argument.

   Example 1

```
# newfs /dev/vg06/rlvol1
newfs: /etc/default/fs    determine the file system type
mkfs  (hfs): Warning -272 sectors in the last cylinder are not
allocated.
mkfs  (hfs): /dev/vg06/rlvol1 - 2400256 sectors in 3847 cylinders
of 16 tracks,
2547.9MB in 241 cyl groups (16 c/g, 10.22Mb/g,  1600 ig/g)
Super block backups (for fsck -b) at:
16, 10040, 20064, 30038, 40112, 50136, 60160, 70184, 80208, 90232, ...
2396176
```

   Example 2

```
# newfs /dev/vg06/rlvol1  create file system
newfs: / etc/default/fs   determine the file system type
mkfs (hfs): ...
:
7188496, 7198520, 7208544
#
```

   Example 3

```
# newfs -F vxfs /dev/vg06/rlvol1   Specify file system type
:
# newfs -F hfs /dev/vg06/rlvol2
```

2. Repeat step 1 for each logical volume on the disk array.

Setting the I/O timeout parameter

Set the I/O timeout value for each disk device to 60 seconds.

1. Verify the current I/O time-out value using the `pvdisplay` command:

   Example

```
# pvdisplay /dev/dsk/c0t6d0

- - - Physical volumes - - -
PV Name                 /dev/dsk/c0t6d0
VG Name                /dev/vg06
PV Status              available
Allocatable            yes
VGDA                   2
Cur LV                 1
PE Size (Mbytes)       4
Total PE               586
Free PE                0
Allocated PE           586 [OPEN-9]
Stale PE               0
IO Timeout (Seconds)   default [I/O timeout value]
```

2. If the I/O timeout value is not 60, change the value to 60 using the `pvchange -t` command:

   Example

```
# pvchange -t 60 /dev/dsk/c0t6d0

Physical volume "/dev/dsk/c0t6d0" has been successfully changed.

Volume Group configuration for /dev/vg06 has been saved in
/etc/lvmconf/vg06.conf.
```

3. Verify that the new I/O timeout value is 60 seconds using the `pvdisplay` command:

   Example

```
# pvdisplay /dev/dsk/c0t6d0

--- Physical volumes ---
PV Name               /dev/dsk/c0t6d0
VG Name               /dev/vg06
PV Status             available
:
Stale PE              0
IO Timeout (Seconds)  60   [New I/O timeout value]
```

4. Repeat the above steps for each new disk connected to the system.

Creating the mount directories

Create a mount directory for each logical volume. Choose a unique name for each mount directory that identifies the logical volume.

To create a mount directory for each logical volume:

1. Use `mkdir` with the new mount directory name as the argument to create the mount directory.

   Example

   ```
   # mkdir /AHPMD-LU00
   ```

2. Use the `ls -x` command to verify the new mount directory.

   Example

   The following example shows the root directory as the location for the mount directories.

   ```
   # ls -x
     AHPMD-LU00    bin     dev         device    etc       export
     floppy        home    hstsboof    kadb      kernel    lib
   ```

3. Repeat above steps for each logical volume on the disk array.

Mounting and verifying the file systems

After the mount directories have been created, mount and verify the file system for each logical volume.

To mount and verify the file systems:

1. Use `mount` to mount the file system for the volume.

   Example

   ```
   # mount /dev/vg06/lvol1  /AHPMD-LU00
   ```

2. Repeat step 1 for each logical volume on the disk array.

   If you need to unmount a file system, use the `unmount` command.

3. Use the `bdf` command to verify that the file systems are correct. The capacity is listed under Kbytes.

   Example

   ```
   # bdf

   Filesystem          Kbytes    used    avail     %used  Mounted on
   /ldev/vg00/lvol1    59797     59364   0         100%   /
   :
   /ldev/vg06/lvol1    2348177   9       2113350   0%     /AHPMD-LU00
   ```

4. As a final verification, perform some basic UNIX operations (for example file creation, copying, and deletion) on each logical device to make sure that the devices on the disk array are fully operational.

   Example

```
#cd /AHPMD-LU00

#cp /bin/vi /AHPMD-LU00/vi.back1

#ls -l

drwxr-xr-t 2 root root 8192 Mar 15 11:35 lost+found
-rwxr-xr-x 1 root sys 217088 Mar 15 11:41 vi.back1

#cp vi.back1 vi.back2

#ls -l drwxr-xr-t 2 root root 8192 Mar 15 11:35 lost+found
-rwxr-xr-x 1 root sys 217088 Mar 15 11:41 vi.back1
-rwxr-xr-x 1 root sys 217088 Mar 15 11:52 vi.back2
```

Setting and verifying the auto-mount parameters

Set up and verify the auto-mount parameters for each new volume. The `/etc/checklist` file (which can also be called the `/etc/fstab` file) contains the auto-mount parameters for the logical volumes.

To set up and verify the auto-mount parameters:

1. Edit the `/etc/checklist` (`/etc/fstab`) file to add a line for each OPEN-*x* device on the disk array. This example and the following table show the auto-mount parameters.

   Example

```
#cp -ip /etc/checklist /etc/checklist.standard
#vi /etc/checklist

/dev/vg00/lvol1  /             hfs    rw        0  1   # root
/dev/vg00/lvol2  swap          ignore rw        0  0   # primary swap
:
/dev/vg06/lvol1  /AHPMD-LU00   hfs    defaults  0  2   # AHPMD-LU00
/dev/vg06/lvol2  /AHPMD-LU01   hfs    defaults  0  2   # AHPMD-LU01

P1               P2            P3     P4        P5 P6  P7
```

Table 6 Auto-mount parameters

Parameter	Name	Enter
P1	Device to mount	Block-type device file name
P2	Mount point	Mount directory name
P3	File system	Type of file system (for example, hfs, vxfs)
P4	Mount options	"defaults" or other appropriate mount options
P5	Enhance	0
P6	File system check (fsck pass)	Order for performing file system checks
P7	Comments	Comment statement

2. Reboot the system.
3. Use the `bdf` command to verify the file system again.

2 Troubleshooting

This section includes resolutions for various error conditions you may encounter.

If you are unable to resolve an error condition, ask your HP support representative for assistance.

Error conditions

Depending on your system configuration, you may be able to view error messages (R-SIMS) as follows:

- In XP Remote Web Console (Status tab)
- In XP Command View Advanced Edition Software (Alerts window)
- In XP Command View (Event History or Event Notification windows) (XP arrays only)

Table 7 Error conditions

Error condition	Recommended action
The logical devices are not recognized by the host.	Verify the following: • The READY indicator lights on the disk array are ON. • Fiber cables are correctly installed and firmly connected. • 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. • The maximum number of LUSE devices per port has not been exceeded. • The disk array host mode is set correctly.
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.
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 37. Verify that the capacity of the volume group is not less than the total capacity of the partitioned logical 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.

Error condition	Recommended action
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 FCA 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.

3 Support and other resources

Related documentation

The following documents provide related information:

- *HP StorageWorks XP10000 Disk Array: Owner's Guide*
- *HP StorageWorks XP12000 Disk Array: Owner's Guide*
- *HP StorageWorks XP24000/XP20000 Disk Array Owner's Guide*
- *HP StorageWorks 200 Storage Virtualization System: Owner's Guide*
- *HP StorageWorks XP LUN Manager User's Guide*
- *HP StorageWorks SAN Design Reference Guide*
- *HP StorageWorks XP Command View Advanced Edition Software Device Manager Web Client User's Guide*

You can find these documents on the HP Manuals website:

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

In the Storage section, click **Disk Storage Systems** and then select a product.

Conventions for storage capacity values

HP XP storage systems use the following values to calculate physical storage capacity values (hard disk drives):

- 1 KB (kilobyte) = 1,000 bytes
- 1 MB (megabyte) = $1,000^2$ bytes
- 1 GB (gigabyte) = $1,000^3$ bytes
- 1 TB (terabyte) = $1,000^4$ bytes
- 1 PB (petabyte) = $1,000^5$ bytes

HP XP storage systems use the following values to calculate logical storage capacity values (logical devices):

- 1 KB (kilobyte) = 1,024 bytes
- 1 MB (megabyte) = $1,024^2$ bytes
- 1 GB (gigabyte) = $1,024^3$ bytes
- 1 TB (terabyte) = $1,024^4$ bytes
- 1 PB (petabyte) = $1,024^5$ bytes
- 1 block = 512 bytes

HP technical support

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

http://www.hp.com/support

Before contacting HP, collect the following information:

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

Subscription service

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

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

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

HP websites

For additional information, see the following HP websites:

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

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.

A Path worksheet

Worksheet

Table 8 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:

B Disk array supported emulations

This appendix provides information about 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

HP recommends using OPEN-V as the emulation for better performance and features that may not be supported with the legacy emulations (OPEN-[389LE]).

Table 9 Supported emulations

XP model	Emulation	Fixed Size	LUSE	CVS	LUSE & CVS
	OPEN-3	Yes	Yes	Yes	Yes
XP10000	OPEN-8	Yes	Yes	Yes	Yes
XP12000	OPEN-9	Yes	Yes	Yes	Yes
XP20000	OPEN-E	Yes	Yes	Yes	Yes
XP24000	OPEN-K				
SVS 200	OPEN-L	Yes	Yes		
	OPEN-M				
	OPEN-V			Yes	Yes

NOTE:

For the SVS 200, and the XP24000/XP20000/XP12000/XP10000 when connected to external storage devices, HP recommends using OPEN-V as the emulation the array makes visible to the host. This allows external storage LDEVs to be configured without losing data. Using any other emulation may cause data loss in the external storage LUNs.

Emulation specifications

Table 10 Emulation specifications

Emulation *(Note 1)*	Category *(Note 2)*	Product name *(Note 3)*	Blocks (512 bytes)	Sector size (bytes)	# of cylinders	Heads	Sectors per track	Capacity MB* *(Note 4)*
OPEN-3	SCSI disk	OPEN-3	4806720	512	3338	15	96	2347

Emulation *(Note 1)*	Category *(Note 2)*	Product name *(Note 3)*	Blocks (512 bytes)	Sector size (bytes)	# of cylinders	Heads	Sectors per track	Capacity MB* *(Note 4)*
OPEN-8	SCSI disk	OPEN-8	14351040	512	9966	15	96	7007
OPEN-9	SCSI disk	OPEN-9	14423040	512	10016	15	96	7042
OPEN-E	SCSI disk	OPEN-E	28452960	512	19759	15	96	13893
OPEN-L	SCSI disk	OPEN-L	71192160	512	49439	15	96	34761
LUSE								
OPEN-3*n	SCSI disk	OPEN-3*n	4806720*n	512	3338*n	15	96	2347*n
OPEN-8*n	SCSI disk	OPEN-8*n	14351040*n	512	9966*n	15	96	7007*n
OPEN-9*n	SCSI disk	OPEN-9*n	14423040*n	512	10016*n	15	96	7042*n
OPEN-E*n	SCSI disk	OPEN-E*n	28452960*n	512	19759*n	15	96	13893*n
OPEN-L*n	SCSI disk	OPEN-L*n	71192160*n	512	49439*n	15	96	34761*n
CVS								
OPEN-3 CVS	SCSI disk	OPEN-3-CVS	Note 5	512	Note 6	15	96	Note 7
OPEN-8 CVS	SCSI disk	OPEN-8-CVS	Note 5	512	Note 6	15	96	Note 7
OPEN-9 CVS	SCSI disk	OPEN-9-CVS	Note 5	512	Note 6	15	96	Note 7
OPEN-E CVS	SCSI disk	OPEN-E-CVS	Note 5	512	Note 6	15	96	Note 7
OPEN-V	SCSI disk	OPEN-V	Note 5	512	Note 6	15	128	Note 7
CVS LUSE								
OPEN-3*n CVS	SCSI disk	OPEN-3*n-CVS	Note 5	512	Note 6	15	96	Note 7
OPEN-8*n CVS	SCSI disk	OPEN-8*n-CVS	Note 5	512	Note 6	15	96	Note 7
OPEN-9*n CVS	SCSI disk	OPEN-9*n-CVS	Note 5	512	Note 6	15	96	Note 7
OPEN-E*n CVS	SCSI disk	OPEN-E*n-CVS	Note 5	512	Note 6	15	96	Note 7
OPEN-V*n	SCSI disk	OPEN-V*n	Note 5	512	Note 6	15	128	Note 7

General notes:

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

The value *n* is the number of volumes combined together. For example, with 8 combined volumes: OPEN-V*8.

Note 1:

The availability of an emulation 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 command device (used for XP Raid Manager) is distinguished by -CM on the product name (for example, OPEN-3-CM, OPEN-3-CVS-CM).

Note 4:

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 5:

The number of blocks for a CVS volume is calculated as follows:

# of blocks = (# of cylinders) × (# of heads) × (# of sectors per track)

Example

```
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

```
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 6:

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 7:

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 XP Remote Web Console or XP Command View Advanced Edition Software.

LUSE device parameters

Table 11 LUSE device parameters

Device type		Physical extent size (PE)	Max physical extent size (MPE)
OPEN-K/3/8/9/E OPEN-3/K*n (n= 2 to 36) OPEN-3/K-CVS OPEN-3/K*n-CVS (n = 2 to 36)		default	default
OPEN-8/9*n	n = 2 to 17	default	default
	n = 18	8	15845
	n = 19	8	16725
	n = 20	8	17606
	n = 21	8	18486
	n = 22	8	19366
	n = 23	8	20247
	n = 24	8	21127
	n = 25	8	22007
	n = 26	8	22888
	n = 27	8	23768
	n = 28	8	24648
	n = 29	8	25529
	n = 30	8	26409
	n = 31	8	27289
	n = 32	8	28170
	n = 33	8	29050
	n = 34	8	29930

Device type		Physical extent size (PE)	Max physical extent size (MPE)
	n = 35	8	30810
	n = 36	8	31691
OPEN-E*n	n = 2 to 9	default	default
	n = 10	8	17366
	n = 11	8	19102
	n = 12	8	20839
	n = 13	8	22576
	n = 14	8	24312
	n = 15	8	26049
	n = 16	8	27786
	n = 17	8	29522
	n = 18	8	31259
	n = 19	8	32995
	n = 20	8	34732
	n = 21	8	36469
	n = 22	8	38205
	n = 23	8	39942
	n = 24	8	41679
	n = 25	8	43415
	n = 26	8	45152
	n = 27	8	46889
	n = 28	8	48625
	n = 29	8	50362
	n = 30	8	52098
	n = 31	8	53835
	n = 32	8	55572
	n = 33	8	57308
	n = 34	8	59045
	n = 35	8	60782

Device type		Physical extent size (PE)	Max physical extent size (MPE)
	n = 36	8	62518
OPEN-L*n	n = 2 to 3	default	default
OPEN-8/9/E-CVS OPEN-V		default	default
OPEN-8/9/E*n-CVS OPEN-V*n (n = 2 to 36)	70 to 119731(MB) × N1	8	default
	119732 to (MB) × N1	8	N2
N1 = [VCS volume capacity (in MB) from Remote Console PC] × nN2 = ↑ N1 / PE ↑ (↑ means round up to next integer) Example: CVS volume capacity is 6000 MB for OPEN-9*22-CVS volume:N1 = 6000 × 22 = 132000N2 = ↑ 132000/8 ↑ = 16500			

SCSI TID map for Fibre Channel adapters

When an arbitrated loop (AL) is established or reestablished, the port addresses are assigned automatically to prevent duplicate TIDs. With the SCSI over Fibre Channel protocol (FCP), there is no longer a need for target IDs in the traditional sense.

SCSI is a bus-oriented protocol requiring each device to have a unique address because all commands go to all devices. For Fibre Channel, the AL-PA is used instead of the TID to direct packets to the desired destination.

Unlike traditional SCSI, when control of the loop is acquired, a point-to-point connection is established from initiator to target. To enable transparent use of FCP, the operating system maps a TID to each AL-PA.

The host maps SCSI protocol to Fibre Channel protocol and detects and accesses Fibre Channel-connected devices using device files (`/dev/dsk/c*t*d*` and `/dev/rdsk/c*t*d*`) in the same way as for SCSI-connected devices. The device files for Fibre Channel-connected devices are configured in a different way from SCSI-connected devices, because Fibre Channel supports 126 addresses per path while SCSI supports 16 TIDs per path.

The following table identifies the fixed mappings between the TID (drive) values assigned by the operating system and the Fibre Channel native addresses (AL_PA/SEL_ID) for Fibre Channel adapters. The controller number (the dks value in `/dev/dsk/dks*d*l*s*`) depends on the server configuration, and a different value is assigned per each column.

The mapping cannot be done when these conditions exist:

- Disk array devices and other types of devices are connected in the same loop
- Information for unused devices remains in the server system
- Multiple ports participate in the same arbitrated loop

Table 12 SCSI TID map

t value	AL-PA	AL-PA	AL-PA	AL-PA	AL-PA	AL-PA	AL-PA	AL-PA
0	EF	CD	B2	98	72	55	3A	25
1	E8	CC	B1	97	71	54	39	23

t value	AL-PA	AL-PA	AL-PA	AL-PA	AL-PA	AL-PA	AL-PA	AL-PA
2	E4	CB	AE	90	6E	53	36	1F
3	E2	CA	AD	8F	6D	52	35	1E
4	E1	C9	AC	88	6C	51	34	1D
5	E0	C7	AB	84	6B	4E	33	1B
6	DC	C6	AA	82	6A	4D	32	18
7	DA	C5	A9	81	69	4C	31	17
8	D9	C3	A7	80	67	4B	2E	10
9	D6	BC	A6	7C	66	4A	2D	0F
10	D5	BA	A5	7A	65	49	2C	08
11	D4	B9	A3	79	63	47	2B	04
12	D3	B6	9F	76	5C	46	2A	02
13	D2	B5	9E	75	5A	45	29	01
14	D1	B4	9D	74	59	43	27	
15	CE	B3	9B	73	56	3C	26	

C Reference information for the HP System Administration Manager (SAM)

The HP System Administration Manager (SAM) is used to perform HP-UX system administration functions, including:

- Setting up users and groups
- Configuring the disks and file systems
- Performing auditing and security activities
- Editing the system kernel configuration

This appendix provides instructions for:

- Using SAM to configure the disk devices
- Using SAM to set the maximum number of volume groups

Configuring the devices using SAM

The SAM Areas window displays the system administration functions and allows you to select the desired function. The **Disks and File Systems** function allows you to configure new disk devices for LVM operations.

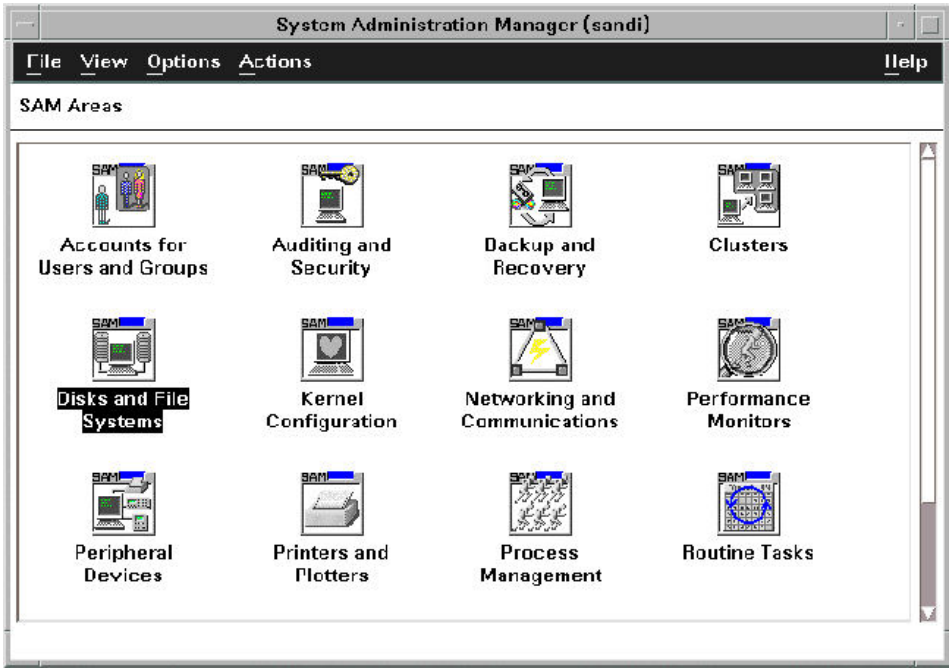

To configure the newly-installed disk array devices:

1. Select **Disks and File Systems**, then select **Disk Devices**.

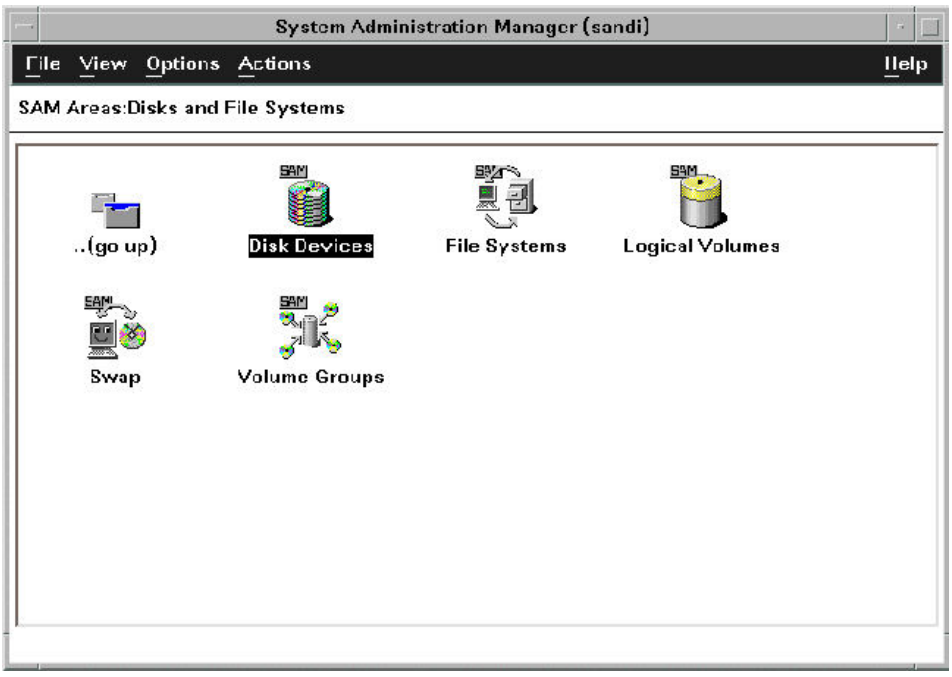

2. Verify that the new disk array devices are displayed in the **Disk Devices** window.
3. Select the device to configure, select the **Actions** menu, select **Add**, and then select **Using the Logical Volume Manager**.
4. In the **Add a Disk Using LVM** window, select **Create...** or **Extend a Volume Group**.
5. In the **Create a Volume Group** window, enter the name of the new or existing volume group to assign the new device to, and then click **OK** twice. The **Add a Disk Using LVM** window now displays the volume group name.
6. Select **Add New Logical Volume** to open the **Create New Logical Volumes** window.
7. In the **Create New Logical Volumes** window, enter the name, size in megabytes, and mount directory for the new logical volume. Select **Add**, and then click **OK** twice.
8. Repeat steps 3–7 for each new disk device on the disk array.

Setting the maximum number of volume groups using SAM

The HP-UX kernel specifies the maximum number of volume groups that can be created. The default is 10. You may need to change this number to accommodate new devices on the disk array. To change this number, use the Kernel Configuration function.

To change the maximum number of volume groups:

1. Select the **Kernel Configuration** function, then select **Configurable Parameters**.
2. In the **Configurable Parameters** window, select the **maxvgs** parameter, select the **Actions** menu, and then select **Modify Configurable Parameter**.
3. In the **Modify Configurable Parameter** window, enter the desired maximum number of volume groups in the **Formula/Value** field, and then click **OK**.

4. In the **Configurable Parameters** window, make sure that none of the parameters are selected. Then select the **Actions** menu, and select **Create New Kernel**.
5. When the configuration window opens, click **Yes** to create the new kernel (or click **No** to return to the **Configurable Parameters** window.)
6. When the **Reboot the System** window opens, click **OK** to move the new kernel into place and reboot the system.

Glossary

AL-PA	Arbitrated loop physical address.
array group	A group of 4 or 8 physical hard disk drives (HDDs) installed in an XP disk array and assigned a common RAID level. RAID1 array groups consist of 4 (2D+2D) or 8 HDDs (4D+4D). RAID5 array groups include a parity disk but also consist of 4 (3D+1P) or 8 HDDs (7D+1P). All RAID6 array groups are made up of 8 HDDs (6D+2P).
command device	A volume on the disk array that accepts XP Continuous Access Software or XP Business Copy Software control operations which are then executed by the disk array.
CU	Control Unit. Contains LDEVs and is approximately equivalent to SCSI Target ID.
CVS	Custom volume size. CVS devices (OPEN-x CVS) are custom volumes configured using array management software to be smaller than normal fixed-size OPEN system volumes. Synonymous with volume size customization (VSC). OPEN-V is a CVS-based volume.
emulation modes	The logical devices (LDEVs) associated with each RAID group are assigned an emulation mode that makes them operate like OPEN system disk drives. The emulation mode determines the size of an LDEV: OPEN-3: 2.46 GB OPEN-8: 7.38 GB OPEN-9: 7.42 GB OPEN-E: 13.56 GB OPEN-L: 36 GB OPEN-V: User-defined custom size
failover	Using an alternate unit or path instead of a failed unit or path in order to continue functioning.
FC	Fibre Channel.
FCA	Fibre Channel Adapter.
FC-AL	Fibre Channel arbitrated loop.
FCP	Fibre Channel Protocol.
HBA	Host bus adapter.
host mode	Each port can be configured for a particular host type. These modes are represented as two-digit hexadecimal numbers. For example, host mode 08 represents an HP-UX host.
LDEV	Logical device. An LDEV is created when a RAID group is carved into pieces according to the selected host emulation mode (that is, OPEN-3, OPEN-8,

	OPEN-E). The number of resulting LDEVs depends on the selected emulation mode. The term LDEV is often used synonymously with the term volume.
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	A LUN is normally associated with only a single LDEV. The LUSE feature allows a LUN to be associated with 1 to 36 LDEVs. Essentially, LUSE makes it possible for applications to access a single large pool of storage. The LUSE feature is available when the HP StorageWorks XP Array Manager product is installed.
OFC	Open Fibre Control.
OPEN-*x*	A general term describing any one of the supported OPEN emulation modes (for example, OPEN-E. There are two types of OPEN-x devices: legacy OPEN-x devices with a fixed size (such as, OPEN-3, OPEN-8, OPEN-9, OPEN-E), and OPEN-V, which has a variable size and is a CVS-based volume.
PA	Physical address.
path	"Path" and "LUN" are synonymous. Paths are created by associating a port and a LUN ID with an LDEV.
port	A physical connection that allows data to pass between a host and the disk array. The number of ports on an XP disk array depends on the number of supported I/O slots and the number of ports available per I/O adapter. The XP family of disk arrays supports Fibre Channel (FC) ports as well as other port types. Ports are named by port group and port letter, such as CL1-A. CL1 is the group, and A is the port letter.
RAID	Redundant array of independent disks.
R-SIM	Remote service information message.
SCSI	Small computer system interface.
SIM	Service information message.
SNMP	Simple Network Management Protocol.
SVP	Service processor, which is the PC built into the disk array. The SVP provides a direct interface into the disk array. SVP use is reserved for HP support representatives only.
SVS	Storage Virtualization System, an appliance that manages multiple disk array storage systems.
TID	Target ID.
Volume	On the XP array, a volume is a uniquely identified virtual storage device composed of a control unit (CU) component and a logical device (LDEV) component separated by a colon. For example 00:00 and 01:00 are two uniquely identified volumes; one is identified as CU = 00 and LDEV = 00, and the other as CU = 01 and LDEV = 00; they are two unique separate virtual storage devices within the XP array.

VSC	Volume size customization. Synonymous with CVS.
WWN	World Wide Name. A unique identifier assigned to a Fibre Channel device.
XP Command View	HP StorageWorks XP Command View, a software product for managing XP arrays. XP Command View runs on a Windows-based management workstation.
XP Command View Advanced Edition Software	HP StorageWorks XP Command View Advanced Edition Software, installs on the user-provided Device Manager server and provides a browser-based platform from which you can manage the SVS 200 and the XP family of disk arrays—even globally distributed arrays.
XP Remote Web Console	HP StorageWorks XP Remote Web Console. A browser-based program installed on the SVP that allows you to configure and manage the disk array.

Index

G

H

I

L

M

O

P

R

S

T

U

V

W

X

Z