HITACHI
DATA SYSTEMS

Hitachi Freedom Storage™
Lightning 9900™

Compaq Tru64™ UNIX™ Configuration Guide

Document Revision Level

Revision	Date	Description
MK-90RD021-P	May 2000	Preliminary Release
MK-90RD021-0	November 2000	Initial Release, supersedes and replaces MK-90RD021-P

Source Document Revision Level

The following source document was used to produce this 9900 configuration guide: Hitachi DKC310 Disk Subsytem SCSI/Fibre Installation Manual, Digital UNIX V4.0, Tru64™ UNIX 4.0F, Tru64™ UNIX 5.0.

Preface

The *Hitachi Freedom Storage™ Lightning 9900™ Compaq Tru64™ UNIX Configuration Guide* describes and provides instructions for installing and configuring the devices on the Lightning 9900™ array subsystem for operation with the Tru64™ UNIX operating system. This configuration guide assumes that:

- the user has a background in data processing and understands direct-access storage device subsystems and their basic functions,
- the user is familiar with the Hitachi Freedom Storage™ 9900 array subsystem,
- the user is familiar with the Tru64® UNIX® operating system.

Note: The term "9900" refers to the entire Hitachi Freedom Storage(TM) Lightning 9900(TM) subsystem family, unless otherwise noted. Please refer to the Hitachi Lightning 9900(TM) User and Reference Guide (MK-90RD008) for further information on the 9900 RAID disk array subsystems.

Note: In this document, "Tru64™ " refers to the Compaq Tru64™ UNIX operating system.

For further information on the 9900 array subsystem, please refer to the *Lightning 9900™ User and Reference Guide* (MK-90RD008), or contact your Hitachi Data Systems account team. The Hitachi Data Systems worldwide web site (http://www.hds.com) also provides information on the Hitachi Lightning 9900™ subsystem and its features and options.

For further information on Tru64™, please consult the Compaq user documentation, or contact Compaq technical support.

Contents

List of Figures

List of Tables

Chapter 1 Overview of 9900 Tru64™ Configuration

1.1 Tru64™ UNIX Configuration

This document describes the requirements and procedures for connecting the 9900 subsystem to a Tru64® system and configuring the new 9900 devices for operation with the Tru64™ UNIX operating system. The Hitachi Data Systems Customer Service representative performs the physical installation of the 9900 subsystem. The user prepares for 9900 subsystem installation and configures the new 9900 devices with assistance as needed from the Hitachi Data Systems representative.

Configuration of the 9900 SCSI disk devices for Tru64™ operations includes:

- Setting the host mode and configuring the 9900 fibre-channel ports (see sections 2.3.1. and 2.3.2),
- Configuring the 9900 fibre-channel ports (see section 2.3.2)
- Verifying new device recognition (see section 2.5),
- Creating the file systems (if necessary) (see section 3.4),
- Mounting and verifying the file systems (see section 3.6 and 3.7).

Note: The term "9900" refers to any 9900 subsystem configured for UNIX and/or PC server operations, including the all-open and multiplatform models. The term "9900" also refers to the entire Lightning 9900™ subsystem family, including the 9960 and 9910 RAID storage subsystems, unless otherwise noted. Please refer to the *Hitachi Lightning 9900™ User and Reference Guide* (MK-90RD008) for further information on the 9960 and 9910 disk array subsystems. The Hitachi Data Systems worldwide web site (http://www.hds.com) also provides information on the Hitachi Lightning 9900 subsystem and its features and options.

Note: For further information on the Compaq Tru64® UNIX® operating system, please consult the Tru64® UNIX® user documentation, or contact Compaq technical support.

1.2 Hitachi Lightning 9900™ Array Subsystem

The Lightning 9900™ array subsystem supports concurrent attachment to multiple PC server and UNIX-based platforms (please contact your Hitachi Data Systems account team for the latest information on platform support). The 9900 subsystem provides continuous data availability, high-speed response, scaleable connectivity, and expandable capacity for PC server and open-system storage. The 9900 subsystem can operate with multihost applications and host clusters, and is designed to handle very large databases as well as data warehousing and data mining applications that store and retrieve terabytes of data.

The Hitachi Lightning 9900™ subsystem provides the following host connectivity options:

- **All open**. The all-open 9900 subsystem is configured with all fibre-channel (FC) ports (no serial ports).
- **Multiplatform**. The multiplatform 9900 subsystem is configured with fibre ports as well as serial (ESCON®) channel ports.
- **All mainframe**. The all-mainframe 9900 subsystem does not have any fibre ports and can only be connected to enterprise server (S/390® mainframe) hosts.

For further information on the 9900 subsystem, please refer to the *Hitachi Lightning 9900™ User and Reference Guide* (MK-90RD008), or contact your Hitachi Data Systems account team. The Hitachi Data Systems worldwide web site (http://www.hds.com) also provides information on the 9900 array subsystem and its features and options.

1.3 Device Types and Configuration Procedures

The 9900 subsystem allows the following types of logical devices (LDEVs) to be installed and configured for operation with the Tru64™ operating system. Table 1.1 lists the device specifications for the 9900 devices. Table 1.3 shows the volume usage (i.e., file system or raw device) for the 9900 devices.

OPEN-x Devices. The OPEN-x logical units (LUs) are SCSI disk devices that are used exclusively by the open-systems host. The 9900 currently supports OPEN-3, OPEN-8, OPEN-9, and OPEN-K and OPEN-E devices. Please contact your Hitachi Data Systems account team for the latest information on supported LU types.

LUSE Devices (OPEN-x*n). The LUSE devices are SCSI disk devices that are used exclusively by the open-systems host. The LU size expansion (LUSE) feature of the 9900 subsystem enables you to configure custom-size LUs which can be from 2 to 36 times larger than standard OPEN-x LUs. LUSE devices are designated as OPEN-x*n, where x is the type of open system, and n is the number of LDEVs that were combined into the LUSE device ($2 \leq n \leq 36$). For example, an OPEN-3 LUSE device that was created from 10 LDEVs would be designated as OPEN-3*10. This capability enables the PC server and open-system hosts to access the data stored on the 9900 subsystem using fewer devices. For more information about the LUSE feature, see *Hitachi Lightning 9900™ LUN Manager User's Guide* MK-90RD006.

CVS Devices (OPEN-x CVS). The CVS devices are SCSI disk devices that are used exclusively by the open-systems host. The custom volume size (CVS) feature of the 9900 subsystem enables you to configure custom-size devices that are smaller than standard OPEN-x devices. This capability enables you to configure a single LU into several custom-sized LUs (usually smaller than the standard LUs) in order to improve host access to frequently accessed files. For more information about the CVS feature, see *Hitachi Lightning 9900™ Virtual LVI/LUN User's Guide* MK-90RD005.

CVS LUSE Devices (OPEN-x*n CVS). The CVS LUSE devices are SCSI disk devices that are used exclusively by the open-systems host. With this combination of features, the user first creates custom-size CVS volumes, and then uses the LUSE feature to combine (concatenate) those volumes into a single LUSE volume. The user can combine from 2 to 36 CVS volumes into a CVS LUSE volume. CVS LUSE volumes are designated as OPEN-x*n CVS, where x is the type of open system and n is the number of CVs that were combined into the LUSE volume ($2 \leq n \leq 36$). For example, an OPEN-3 LUSE volume that was created from 10 CVS volumes would be designated as OPEN-3*10 CVS.

Configuration of the 9900 SCSI disk devices for Tru64™ operations includes:

- Setting the host mode and configuring the 9900 fibre-channel ports (see sections 2.3.1. and 2.3.2),
- Configuring the 9900 fibre-channel poets (see section 2.3.2)
- Verifying new device recognition (see section 2.6),
- Creating the file systems (if necessary) (see section 3.3),
- Mounting and verifying the file systems (see section 3.5 and 3.6).

HMDE Devices (3390-3A/B/C, 3380-KA/B/C, OPEN-x-HMDEoto). The Hitachi Multiplatform Data Exchange (HMDE) feature of the 9900 subsystem enables user data to be shared across mainframe, PC server, and open-system platforms using special multiplatform volumes. The CVS feature can also be applied to HMDE devices for maximum flexibility in volume size. For further information on HMDE, please refer to the *Hitachi Multiplatform Data Exchange User's Guide* (MK-90RD008), or contact your Hitachi Data Systems account team.

The HMDE devices are not SCSI disk devices. The HMDE devices must be installed as raw devices. UNIX/PC server hosts must use HMDE to access the HMDE devices as raw devices (no disk partition, no file system, no mount operation).

Note: The 3390-3B and 3380-KB devices are write-protected from PC server and open-system access. The 9900 subsystem will reject all UNIX/PC server write operations (including SCSI initiators) for the 3390-3B and 3380-KB devices.

WARNING: The 3390-3A/C, 3380-KA/C, and OPEN-x-HMDEoto devices are ***not*** write-protected for open-system or PC server access. Do not execute any write operation by the SCSI initiators. Do not create a file system or use the **newfs** UNIX command on a 3390-3A or 3380-KA device. This will overwrite the mainframe data on the device.

Table 1.1 9900 Device Specifications for Tru64™ Operations (continues on next page)

Device Type *(Note 1)*	Category *(Note 2)*	Vendor Name	Product Name	# of Blocks (512-byte blk)	Sector Size (bytes)	# of Data Cylinders	# of Heads	# of Sectors per Track	Capacity MB *(Note 3)*
OPEN-3	SCSI disk	HITACHI	OPEN-3	4806720	512	3338	15	96	2347
OPEN-9	SCSI disk	HITACHI	OPEN-9	14423040	512	10016	15	96	7042
OPEN-K	SCSI disk	HITACHI	OPEN-K	3661920	512	2543	15	96	1788
OPEN-8	SCSI disk	HITACHI	OPEN-8	14351040	512	9966	15	96	7007
OPEN-E	SCSI disk	HITACHI	OPEN-E	28452960	512	19759	15	96	13888
OPEN-3*n	SCSI disk	HITACHI	OPEN-3*n	4806720*n	512	3338*n	15	96	2347*n
OPEN-9*n	SCSI disk	HITACHI	OPEN-9*n	14423040*n	512	10016*n	15	96	7042*n
OPEN-K*n	SCSI disk	HITACHI	OPEN-K*n	3661920*n	512	2543*n	15	96	1788*n
OPEN-8*n	SCSI disk	HITACHI	OPEN-8*n	14351040*n	512	9966*n	15	96	7007*n
OPEN-E*n	SCSI disk	HITACHI	OPEN-E*n	28452960*n	512	19759*n	15	96	13888*n
OPEN-3 CVS	SCSI disk	HITACHI	OPEN-3-CVS	***Note 4***	512	***Note 5***	15	96	***Note 6***
OPEN-9 CVS	SCSI disk	HITACHI	OPEN-9-CVS	***Note 4***	512	***Note 5***	15	96	***Note 6***
OPEN-K CVS	SCSI disk	HITACHI	OPEN-K-CVS	***Note 4***	512	***Note 5***	15	96	***Note 6***
OPEN-8 CVS	SCSI disk	HITACHI	OPEN-8-CVS	***Note 4***	512	***Note 5***	15	96	***Note 6***
OPEN-E CVS	SCSI disk	HITACHI	OPEN-E-CVS	***Note 5***	512	***Note 6***	15	96	***Note 6***
OPEN-3*n CVS	SCSI disk	HITACHI	OPEN-3*n-CVS	***Note 4***	512	***Note 5***	15	96	***Note 6***
OPEN-9*n CVS	SCSI disk	HITACHI	OPEN-9*n-CVS	***Note 4***	512	***Note 5***	15	96	***Note 6***
OPEN-K*n CVS	SCSI disk	HITACHI	OPEN-K*n-CVS	***Note 4***	512	***Note 5***	15	96	***Note 6***
OPEN-8*n CVS	SCSI disk	HITACHI	OPEN-8*n-CVS	***Note 4***	512	***Note 5***	15	96	***Note 6***

Table 1.2 9900 Device Specifications for Tru64™ Operations (continued)

OPEN-E*n CVS	SCSI disk	HITACHI	OPEN-E*n- CVS	***Note 5***	512	***Note 6***	15	96	***Note 6***
3390-3A	HMDE otm/mto	HITACHI	3390-3A	5825520	512	3345	15	116	2844
3380-KA	HMDE otm/mto	HITACHI	3380-KA	3836160	512	2662	15	96	1873
3390-3B	HMDEmto	HITACHI	3390-3B	5822040	512	3343	15	116	2842
3380-KB	HMDEmto	HITACHI	3380-KB	3833280	512	2660	15	96	1871
3390-3C	OPEN-x-HMDEoto	HITACHI	OP-C-3390-3C	5820300	512	3345	15	116	2844
3380-KC	OPEN-x-HMDEoto	HITACHI	OP-C-3380-KC	3836160	512	2662	15	96	1873
HMDE OPEN-3	OPEN-x-HMDEoto	HITACHI	OPEN-3	4806720	512	3338	15	96	2347
3390-3A CVS	HMDE otm/mto	HITACHI	3390-3A-CVS	***Note 4***	512	***Note 5***	15	116	***Note 6***
3380-KA CVS	HMDE otm/mto	HITACHI	3380-KA-CVS	***Note 4***	512	***Note 5***	15	96	***Note 6***
3390-3B CVS	HMDEmto	HITACHI	3390-3B-CVS	***Note 4***	512	***Note 5***	15	116	***Note 6***
3380-KB CVS	HMDEmto	HITACHI	3380-KB-CVS	***Note 4***	512	***Note 5***	15	96	***Note 6***
3390-3C CVS	OPEN-x-HMDEoto	HITACHI	OP-C-3390-3C-CVS	***Note 4***	512	***Note 5***	15	116	***Note 6***
3380-KC CVS	OPEN-x-HMDEoto	HITACHI	OP-C-3380-KC-CVS	***Note 4***	512	***Note 5***	15	96	***Note 6***
HMDE OPEN-3 CVS	OPEN-x-HMDEoto	HITACHI	OPEN-3-CVS	***Note 4***	512	***Note 5***	15	96	***Note 6***

Note 1: The availability of a specific 9900 device type depends on the level of microcode installed on the 9900 subsystem.

Note 2: The category of a device (SCSI disk or HMDE) determines its volume usage. Table 1.3 shows the volume usage for SCSI disk devices and HMDE devices. The SCSI disk devices (OPEN-x, CVS, LUSE, CVS LUSE) require partitions and file systems for Tru64® operations. The HMDE devices (3390-3A/B/C, 3380-KA/B/C, OPEN-x-HMDEoto) must be installed as raw devices and can only be accessed using HMDE. Do not write a signature or create a partition or file system on any device used for HMDE operations.

Table 1.3 Volume Usage for Device Categories

Category	Device Type	Volume Usage
SCSI Disk	OPEN-x, OPEN-x CVS, OPEN-x*n LUSE, OPEN-x*n CVS LUSE	File System
HMDE	3390-3A/B/C, 3380-KA/B/C 3390-3A/B/C CVS, 3380-KA/B/C CVS OPEN-x-HMDEoto	Raw Device

Note 3: The 9900 command device (used for the Hitachi Command Control Interface operations) is distinguished by –CM in 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. These device capacities are calculated based on 1 MB = 1024^2 bytes rather than 1000^2 bytes.

Note 5: The number of blocks for a CVS volume is calculated as follows:

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

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

Note 6: The number of data cylinders for a CVS volume (including CVS LUSE volumes) is calculated as follows:

(↑…↑ means that the value should be rounded up to the next integer)

- The number of data cylinders for an OPEN-* CVS volume =
  # of cylinders = ↑ (capacity (MB) specified on the Remote Console PC) × 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
- The number of data cylinders for a 3390-3A/C or 3380-KA/C CVS volume =
  # of cylinders = (number of cylinders specified on the Remote Console PC) + 9
- The number of data cylinders for a 3390-3B or 3380-KB CVS volume =
  # of cylinders = (number of cylinders specified on the Remote Console PC) + 7

Note 7: The size of an OPEN-x CVS volume is specified by capacity in MB, not by number of cylinders. The user specifies the volume size using the Remote Console PC.

Chapter 2 Preparing for New Device Configuration

Before you can configure the 9900, you must take the following steps:

- Installing the 9900 subsystem (see section 2.2),
- Setting the host mode for the fibre-channel ports (see section 2.3.1),
- Configuring the fibre-channel ports (see section 2.3.2),
- Connecting the 9900 subsystem to the Tru64 system (see section 2.4),
- Verifying new device recognition (see section 2.6).

2.1 Configuration Requirements

Note: In this document the term "9900" refers to the entire Hitachi Lightning 9900™ subsystem family, including the 9960 and 9910 RAID storage subsystems, unless otherwise noted. Please refer to the Hitachi Lightning 9900 User and Reference Guide (MK-90RD008) for further information on the 9960 and 9910 disk array subsystems.

The requirements for 9900 Tru64™ configuration are:

- Hitachi Lightning 9900™ subsystem.
  - The 9900 Remote Console PC and LUN Manager software are used to configure the fibre-channel (FC) ports. If the remote LUN Manager feature is not installed, please contact your Hitachi Data Systems account team for information on LUN Manager configuration services.
  - ***Note:*** The availability of 9900 features and devices (e.g., CVS, 3390-3C, OPEN-8) depends on the level of microcode installed on the 9900 subsystem.
- Fibre-channel adapters. Make sure to install all utilities, tools, and drivers that come with the adapter(s). For information on driver requirements for the adapters, please refer to the user documentation for the adapter or contact the vendor.
  - The 9900 supports full-speed (100 MB/s) fibre-channel interface, including shortwave non-OFC (open fibre control) optical interface, multimode optical cables with SC connectors, and gigabit link modules (GLMs).
  - For information on supported FC adapters, optical cables, hubs, and fabric switches, please contact your Hitachi Data Systems account team or the Hitachi Data Systems Support Center (see section 5.2).

  Note: LUSE devices cannot be configured on a DKC-410I-4GS/8GS port connected to a QLogic 2100 FC adapter (error code (EC) = B503, invalid LUN).

  Note: Do not connect any OFC-type connector to the 9900 subsystem.
- Tru64® operating system, version 4.0f.

  Note: Hitachi Data Systems plans to support future releases of Tru64™ . This document will be updated as needed to cover version-specific information. For further information on Tru64™ version support, please contact your Hitachi Data Systems account team.

- Superuser (root) login access to the host system.
- Fibre-channel adapters. Make sure to install all utilities, tools, and drivers that come with the adapter(s). For information on driver requirements for the adapters, please refer to the user documentation for the adapter or contact the vendor.
  - The 9900 subsystem supports full-speed (100 MB/s) fibre-channel interface, including shortwave non-OFC (open fibre control) optical interface, and multimode optical cables with SC connectors. Do not connect any OFC-type fibre-channel interface to the 9900 subsystem. For information on supported FC adapters, optical cables, hubs, and fabric switches, please contact your Hitachi Data Systems account team or the Hitachi Data Systems Support Center (see section 5.2).

    Note: LUSE devices cannot be configured on a DKC-410I-4GS/8GS port connected to a QLogic 2100F FC adapter (error code (EC) = B503, invalid LUN).

    Note: If you plan to connect different types of servers to the 9900 subsystem via the same fabric switch, you must use the **zoning** function of the fabric switch.

2.2 Installing the 9900 Subsystem

The 9900 subsystem comes with all hardware and cabling required for installation. Installation of the 9900 subsystem involves the following activities:

1. **Hardware installation**. The Hitachi Data Systems representative performs this activity, which includes:
   - Assembling all hardware and cabling.
   - Loading the latest microcode and SVP updates for full fibre support.
   - Installing and formatting the logical devices (LDEVs) using the SVP. (This requires complete LDEV configuration information from the user, including the desired number of OPEN-x, LUSE, CVS, and multiplatform (HMDE) devices.)
   - Installing the fibre-channel adapters and cabling. ***Note:*** The total fibre cable length attached to each fibre-channel adapter must not exceed 500 meters (1,640 feet). Do not install/de-install fibre-channel cabling while Tru64™ is active. This can cause the system to hang. Always confirm that Tru64™ is shut down before installing/de-installing fibre cabling.
2. Remote console PC and LUN Manager installation. The user or Hitachi Data Systems representative can perform this activity. For instructions on installing the Remote Console PC for the 9900, please refer to *Hitachi Lightning 9900™ LUN Manager User's Guide* MK-90RD003. For more information on the LUN Manager feature, please refer to *Hitachi Lightning 9900™ LUN Manager User's Guide* MK-90RD006. ***Note:*** If the remote LUN Manager feature is not installed, please contact your the Hitachi Data Systems account team for information on LUN Manager configuration services.

2.3 Preparing to Connect the 9900 Subsystem

Before the 9900 is connected to the host system, you must perform the following tasks:

- Set the host mode for the 9900 fibre-channel port(s) (see section 2.3.1), and
- Configure the 9900 fibre-channel ports (see section 2.3.2).

2.3.1 Setting the Host Mode for the 9900 Ports

The 9900 ports have special modes which must be set for the connected operating system. The required host mode setting for Compaq Tru64 operations is **00** (standard mode, default setting). Use the LUN Manager remote console software to ensure that the host mode for each port is connected to the Compaq Tru64 (see Figure 2.1).

Note: If you plan to connect different types of servers to the 9900 subsystem via the same fabric switch, you must use the **zoning** function of the fabric switch.

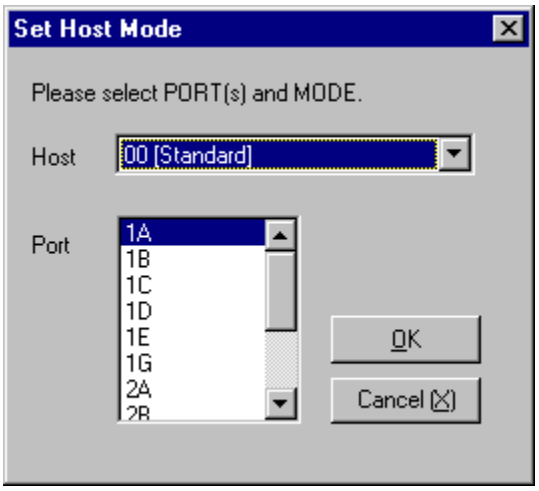

Figure 2.1 Setting the Host Mode

2.3.2 Configuring the 9900 Fibre-Channel Ports

You need to configure the 9900 FC ports to define the fibre parameters (see Figure 2.4 and Table 2.1) and port addresses (see Table 2.2). You will use the LUN Manager remote console software to configure the 9900 FC ports. For instructions on using LUN Manager, please refer to the *Hitachi Freedom Storage™ Lightning 9900™ LUN Manager User's Guide* (MK-90RD006).

Note: The 9900 subsystem supports up to 32 LUs for Tru64 operations. For further information, see your Hitachi Data Systems representative on the number of LUs supported.

Fibre topology. Figure 2.4 shows the Fibre Parameter panel (part of the LUN Manager software), and Table 2.1 explains the settings on this panel. You will select the appropriate settings for each 9900 FC port based on the device to which the port is connected. Determine the topology parameters supported by the device, and set your topology accordingly. The type of 9900 port is also important. ***Note:*** If you plan to connect different types of servers to the 9900 via the same fabric switch, you must use the **zoning** function of the fabric switch.

You may be able to set an HBA to run in arbitrated loop mode or in fabric mode after selecting diagnostic mode. Ask your Hitachi Data Systems representative if your system supports this.

Note: If the HBA is set to loop topology, console commands such as Show Device and Quick Assign will not display any HBAs or their attached device.

Display the HBA configuration as shown in Figure 2.2.

```
P00>>>wwidmgr-show adapter

Item                              adapter            WWN        Cur.Topo    Next Topo
kgpsaa0.0.0.4.6 - Nvram read failed
[0]                            kgpsaa0.0.0.4.6  1000-0000-c921-0c22  fabric        unavail
kgpsab0.0.0.8.6 - Nvram read failed
[1]                            kgpsaa0.0.0.8.6  1000-0000-c921-0c7a  loop          unavail
```

Figure 2.2 Displaying the HBA Configuration

NOTE: The message "Nvram Read Failed" indicates that the NVRAM on the HBA has not been initialized and formatted. This is normal.

Set the HBA configuration as shown in Figure 2.3.

```
P00>>>wwidmgr-set adapter -item 9999 -topo loop
Item  adapter          WWN                  Cur.Topo      Next Topo
[0]   Kgpsaa0.0.0.4.6  1000-0000-c921-0c22  fabric        loop

P00>>>wwidmgr-set adapter -item 9999 topo fabric
Item  adapter          WWN                  Cur.Topo      Next Topo
[1]   kgpsaa0.0.0.4.6  1000-0000-c921-0c22  loop          fabric
```

Figure 2.3 Setting the HBA Configuration

Port address. In fabric environments, the port addresses are assigned automatically by fabric switch port number and are not controlled by the 9900 port settings. In arbitrated loop environments, the port addresses are set by entering an AL-PA (arbitrated-loop physical address, or loop ID). Table 2.2 shows the available 9900 AL-PA values ranging from 01 to EF. Fibre-channel protocol uses the AL-PAs to communicate on the fibre-channel link, but the software driver of the platform host adapter translates the AL-PA value assigned to the 9900 port to a SCSI TID. See Appendix B for a description of the AL-PA-to-TID translation.

Note on loop ID conflict: The Compaq Tru64 system assigns port addresses from lowest (01) to highest (EF). To avoid loop ID conflict, assign the port addresses from highest to lowest (i.e., starting at EF). The AL-PAs should be unique for each device on the loop to avoid conflicts. Do not use more than one port address with the same TID in same loop (e.g., addresses EF and CD both have TID 0, refer to Appendix B for the TID-to-AL-PA mapping).

Table 2.1 Fibre Parameter Settings on the 9900 Remote Console PC

Fabric Parameter	Connection Parameter	Provides:
ON	FC-AL	FL-port (public arbitrated loop)
ON	Point-to-Point	F-port (fabric port)
OFF	FC-AL	AL-port (private arbitrated loop)
OFF	Point-to-Point	*Not supported*

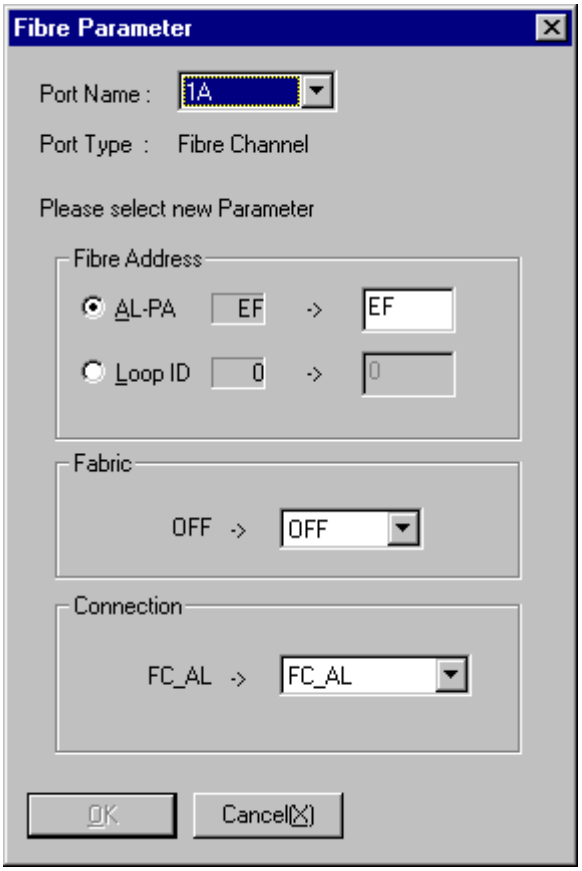

Figure 2.4 Fibre Parameter Panel (from the LUN Manager Software)

Table 2.2 Available AL-PA Values

EF	CD	B2	98	72	55	3A	25
E8	CC	B1	97	71	54	39	23
E4	CB	AE	90	6E	53	36	1F
E2	CA	AD	8F	6D	52	35	1E
E1	C9	AC	88	6C	51	34	1D
E0	C7	AB	84	6B	4E	33	1B
DC	C6	AA	82	6A	4D	32	18
DA	C5	A9	81	69	4C	31	17
D9	C3	A7	80	67	4B	2E	10
D6	BC	A6	7C	66	4A	2D	0F
D5	BA	A5	7A	65	49	2C	08
D4	B9	A3	79	63	47	2B	04
D3	B6	9F	76	5C	46	2A	02
D2	B5	9E	75	5A	45	29	01
D1	B4	9D	74	59	43	27	
CE	B3	9B	73	56	3C	26	

2.4 Connecting the 9900 Subsystem to the Tru64™ System

After you have configured the 9900 fibre-channel ports and verified the host fibre-channel adapter installation, you are ready to connect the 9900 subsystem to the Tru64™ system. The 9900 subsystem comes with all the hardware and cabling required for connection to the host system(s).

To connect the 9900 subsystem to the host system:

1. **Verify subsystem installation**. The Hitachi Data Systems representative verifies that the status of the fibre ports and LDEVs is normal. The Hitachi Data Systems representative should also check the fibre device parameters to make sure that all 9900 LDEVs are unique for each host system.
2. **Shut down and power off the host system**. The user should perform this activity. You must shut down and power off the host system before connecting the 9900:
   a) Shut down the host system.
   b) When shutdown is complete, power off the Tru64™ display.
   c) Power off all peripheral devices except for the 9900 subsystem.
   d) Power off the Tru64™ system. You are now ready to connect the 9900 subsystem.
3. **Connect the 9900 to the host system**. The Hitachi Data Systems representative installs the fibre cables between the 9900 and the host system. ***Note:*** The Hitachi Data Systems representative must use the 9900 maintenance manual during all installation activities. Follow all precautions and procedures in the maintenance manual, and always check all specifications to ensure proper installation and configuration.
4. **Power on and boot up the host system**. The user should perform this activity. To power on the host system after connecting the 9900:
   a) Power on the host system display.
   b) Power on all peripheral devices. The 9900 should already be on. The fibre-channel ports should already be configured. If not, the host system may need to be restarted in order to recognize the new devices.
   c) Confirm the ready status of all peripheral devices, including the 9900.
   d) Power on the Tru64™ system. **Do not boot the system yet.**

2.5 Verifying New Device Recognition

After connecting the 9900 subsystem to the Tru64® UNIX® system, you need to verify that the system recognizes the new devices. The devices must be installed and formatted and the fibre-channel ports configured before the host system is powered on. If not, the user must shut down and restart the system to allow the system to recognize the new devices.

To verify that the Tru64® UNIX® system recognizes the devices on the newly installed 9900:

1. If desired, you can verify new 9900 device recognition before booting the Tru64® UNIX® system using the **show device |more** command (see Figure 2.5). This listing of the device information is not as easy to read as the listing produced by the **scu show edt** UNIX command (see step 4).
2. Boot the Tru64® UNIX® system as usual (e.g., **boot**).
3. Log into the Tru64® UNIX® system as **root**.
4. If you did not verify device recognition for all 9900 devices before booting the system, then you need to verify new device recognition using the **scu show edt** UNIX command (see Figure 2.6). Make sure all new devices, including OPEN-x, LUSE, CVS, and HMDE, are recognized by the Tru64® UNIX® system.
5. Record the device filename in your SCSI Device Information Worksheet. See Table 2.3. You will need this device file information when you configure the new devices

Note: The sample screen shown here may not be exactly the same for all Tru64® UNIX® systems.

```
>>> show device↵
polling for units on kzpsa0, slot0, bus0, hose0...
kzpsa0.7.0.0.0 → SCSI board dka → SCSI bus r 1 Fast 1 Bus ID=7 F01 A10 → ID =7 Initiator
ID


dka0.0.0.0.0        Dka0                                DEC RZ29B
              Indicates the SCSI ID.
dka200.2.0.0.0     Dka200                               DEC RZ29B
dka300.3.0.0.0     Dka300                               DEC RZ29B
dka400.4.0.0.0     Dka400                               DEC RZ29B
dka500.5.0.0.0     Dka500                               DEC RZ29B
polling for units on kzpsa1, Cslot2, Cbus0, Chose0...
kzpsa 1.7.0.2.0     dkb     TPwr 1 Fast 1 Bus ID=7 N01 A10
:
:
polling kgpsa0(KGPSA-B) slot3, bus 0 PCI, hose0
kgpsa0.0.0.3.3                PGA0     WWN1000-0000-c921-3a54   → Fibre host adapter
:
```

Figure 2.5 Verifying 9900 Device Recognition (from the boot prompt)

```
# scu show edt↵
Device: RZ29B    Bus: 0, Target: 0, Lun: 0, Type: Direct Access → Bus = Bus Number
                                                                  → Target = SCSI ID
                                                                   → LUN = LU Number

→ (ID=0).
Device: RZ29B    Bus: 0, Target: 0, Lun: 1, Type: (not present)
Device: RZ29B    Bus: 0, Target: 0, Lun: 2, Type: (not present)
Device: RZ29B    Bus: 0, Target: 0, Lun: 3, Type: (not present)
Device: RZ29B    Bus: 0, Target: 0, Lun: 4, Type: (not present)
Device: RZ29B    Bus: 0, Target: 0, Lun: 5, Type: (not present)
Device: RZ29B    Bus: 0, Target: 0, Lun: 6, Type: (not present)
Device: RZ29B    Bus: 0, Target: 0, Lun: 7, Type: (not present)
Device: RZ29B    Bus: 0, Target: 2, Lun: 0, Type: Direct Access  ← TID =2
Device: RZ29B    Bus: 0, Target: 2, Lun: 1, Type: (not present)
Device: RZ29B    Bus: 0, Target: 2, Lun: 2, Type: (not present)
Device: RZ29B    Bus: 0, Target: 2, Lun: 3, Type: (not present)
Device: RZ29B    Bus: 0, Target: 2, Lun: 4, Type: (not present)
Device: RZ29B    Bus: 0, Target: 2, Lun: 5, Type: (not present)
Device: RZ29B    Bus: 0, Target: 2, Lun: 6, Type: (not present)
Device: RZ29B    Bus: 0, Target: 2, Lun: 7, Type: (not present)
Device: RZ29B    Bus: 0, Target: 3, Lun: 0, Type: Direct Access  ← TID = 3
Device: RZ29B    Bus: 0, Target: 3, Lun: 1, Type: (not present)
Device: RZ29B    Bus: 0, Target: 3, Lun: 2, Type: (not present)
Device: RZ29B    Bus: 0, Target: 3, Lun: 3, Type: (not present)
Device: RZ29B    Bus: 0, Target: 3, Lun: 4, Type: (not present)
Device: RZ29B    Bus: 0, Target: 3, Lun: 5, Type: (not present)
Device: RZ29B    Bus: 0, Target: 3, Lun: 6, Type: (not present)
Device: RZ29B    Bus: 0, Target: 3, Lun: 7, Type: (not present)
Device: RZ29B    Bus: 0, Target: 4, Lun: 0, Type: Direct Access  ← TID = 4
Device: RZ29B    Bus: 0, Target: 4, Lun: 1, Type: (not present)
Device: RZ29B    Bus: 0, Target: 4, Lun: 2, Type: (not present)
Device: RZ29B    Bus: 0, Target: 4, Lun: 3, Type: (not present)
Device: RZ29B    Bus: 0, Target: 4, Lun: 4, Type: (not present)
Device: RZ29B    Bus: 0, Target: 4, Lun: 5, Type: (not present)
Device: RZ29B    Bus: 0, Target: 4, Lun: 6, Type: (not present)
Device: RZ29B    Bus: 0, Target: 4, Lun: 7, Type: (not present)
Device: RZ28M    Bus: 0, Target: 5, Lun: 0, Type: Direct Access ← TID = 5
Device: RZ28M    Bus: 0, Target: 5, Lun: 1, Type: (not present)
Device: RZ28M    Bus: 0, Target: 5, Lun: 2, Type: (not present)
Device: RZ28M    Bus: 0, Target: 5, Lun: 3, Type: (not present)
Device: RZ28M    Bus: 0, Target: 5, Lun: 4, Type: (not present)
Device: RZ28M    Bus: 0, Target: 5, Lun: 5, Type: (not present)
Device: RZ28M    Bus: 0, Target: 5, Lun: 6, Type: (not present)
#
```

Figure 2.6 Verifying 9900 Device Recognition from UNIX prompt for Tru64 UNIX 4.0 and 4.0F

Table 2.3 Sample SCSI Device Information Worksheet

LDEV (CU:LDEV)	LU Type	LUSE (✓)	CVS (✓)	Device Number	Bus Number	Path 1	Alternate Path(s)		
0:00						TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____
0:01						TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____
0:02						TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____
0:03						TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____
0:04						TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____
0:05						TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____
0:06						TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____
0:07						TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____
0:08						TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____
0:09						TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____
0:0a						TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____
0:0b						TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____
0:0c						TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____
0:0d						TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____
0:0e						TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____
0:0f						TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____
0:10						TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____
and so on...									

Chapter 3 Configuring the New Devices

After 9900 subsystem installation and connection are complete, the devices on the 9900 subsystem are ready to be configured for use. Configuration of the 9900 devices is performed by the user and requires superuser (root) access to the Tru64® UNIX® system. The activities involved in 9900 Tru64® UNIX® configuration are:

- Create the device files (see section 3.1),
- Adding the 9900 device information to the device definition file (see section 3.2),
- Create the file systems (see section 3.3),
- Create the mount directories (see section 3.4),
- Mount the file systems (see section 3.5),
- Verify the file systems (see section 3.6), and
- Set and verify the auto-mount parameters (see section 3.7).

3.1 Verifying the Device Files

The Tru64® UNIX® system should create the device files for all new devices automatically during startup. Please check to make sure these files were created. When the device files are created, Tru64® UNIX® automatically reads the disk partitioning information directly from the disks.

Note: The Tru64® UNIX® device file naming convention is different than the device file naming convention used by other UNIX-based platforms. The Tru64® UNIX® character-type device file name has the form **rrzXYZ** (**rzXYZ** for block-type device files), where:

X = b through h = LUN1 through LUN7 (no letter is used for LUN0)
Y = SCSI bus number × 8 + SCSI TID
Z = partition = a through h

For example: **rrzc18a** = SCSI bus 2, SCSI TID 2, LUN 2, partition a
rrz12c = SCSI bus 1, SCSI TID 4, LUN 0, partition c

To verify that the device files for the new 9900 devices were successfully created:

1. Go to the **/dev** directory by entering: `cd /dev`
2. List all device files by entering: `ls -l *rz*`
   The character-type device files are listed first by bus number, and then the block-type device files are listed.
3. Check the device file list to make sure block-type and character-type device files were generated for each newly installed 9900 device. For example, if you configured a LUN at TID 2, LUN 2 on bus 2, make sure that both **rrzc18(a-h)** and **rzc18(a-h)** are present.

```
# file /dev/rrz* | grep OPEN-3↵ ←Disk type
                 └
/dev/rrz8a:    character special (8/16384) SCSI #1 OPEN-3 disk #64 (SCSI ID #0) (SCSI LUN #0)
(File name)
/dev/rrz8b:           character special (8/16385) SCSI #1 OPEN-3 disk #64 (SCSI ID #0) (SCSI
LUN #0)
/dev/rrz8c:           character special (8/16386) SCSI #1 OPEN-3 disk #64 (SCSI ID #0) (SCSI
LUN #0)
/dev/rrz8d:           character special (8/16387) SCSI #1 OPEN-3 disk #64 (SCSI ID #0) (SCSI
LUN #0)
/dev/rrz8e:           character special (8/16388) SCSI #1 OPEN-3 disk #64 (SCSI ID #0) (SCSI
LUN #0)
/dev/rrz8f:           character special (8/16389) SCSI #1 OPEN-3 disk #64 (SCSI ID #0) (SCSI
LUN #0)
/dev/rrz8g:           character special (8/16390) SCSI #1 OPEN-3 disk #64 (SCSI ID #0) (SCSI
LUN #0)
/dev/rrz8h:           character special (8/16391) SCSI #1 OPEN-3 disk #64 (SCSI ID #0) (SCSI
LUN #0)
/dev/rrzb8a:  character special (8/16448) SCSI #1 OPEN-3 disk #65 (SCSI ID #0) (SCSI LUN #1)
/dev/rrzb8b:  character special (8/16449) SCSI #1 OPEN-3 disk #65 (SCSI ID #0) (SCSI LUN #1)
:
:
```

Figure 3.1 Verifying the Device Files for Tru64 UNIX 4.0 and 4.0F

3.2 Enabling Command Tag Queuing

To enable SCSI command tag queuing for the 9900 devices, you must add the 9900 device information to the device definition file (**/etc/ddr.dbase**). Table 3.1 lists the queue depth requirements for 9900 devices.

To enable SCSI command tag queuing for the 9900 devices:

1. Before editing, make a backup copy of the **/etc/ddr.dbase** file.
2. Change the file permission to allow write access to the file by entering:
   ```
   chmod u+rw /etc/ddr.dbase
   ```
3. Edit the file using the UNIX **vi** editor to add the 9900 device information. Add the text shown in Figure 3.2
4. Exit the UNIX **vi** editor and save your changes.
5. Rebuild the kernel by entering: `ddr_config c.`
6. Verify the new 9900 device information as shown in Figure 3.5:
   ```
   /sbin/ddr_config -s disk HITACHI 9900
   ```

The 9900 device information you just entered should be displayed.

```
# ddr_config -s disk "HITACHI" "OPEN-3"↵                    ← Specify the disk type.

Building Device Information for:
  Type = disk
  Vendor ID: "HITACHI"   Product ID: "OPEN-3"

Applying Modifications from Device Record for :
  Stype: 0x0   Vendor ID: "HITACHI"   Product ID: "OPEN-3"
    TypeSubClass = hard_disk, RAID
    BadBlockRecovery = disabled
    DisperseQueue = false
    TagQueueDepth = 0x8                                     ← See Table 3.1.
    InquiryLength = 0x24
    RequestSenseLength = 0x60


The resulting SCSI Device information looks as follows:

SCSIDEVICE
    Type = disk
    Name = "HITACHI" "OPEN-3"
    PARAMETERS:
       TypeSubClass                 = hard_disk, RAID
       BlockSize                    = 0x200
       MaxTransferSize              = 0x1000000
       BadBlockRecovery             = disabled
       SyncTransfers                = enabled
       DynamicGeometry              = false
       Disconnects                  = enabled
       TaggedQueuing                = enabled
       CmdReordering                = enabled
       LongTimeoutRetry             = disabled
       DisperseQueue                = false
       WideTransfers                = enabled
       WCE_Capable                  = true
       PwrMgmt_Capable              = true
       Additional_Flags             = 0x0
       TagQueueDepth                =                      ← See Table 3.1.
       ReadyTimeSeconds             = 0x2d
       CMD_PreventAllow             = notsupported
       CMD_ExtReserveRelease        = notsupported
       CMD_WriteVerify              = notsupported
       Additional_Cmds              = 0x0
       InquiryLength                = 0x24
       RequestSenseLength           = 0x60

#
```

Figure 3.2 Verifying the New 9900 Device Information for Tru64 UNIX 4.0 and 4.0F

Table 3.1 Queue Depth Requirements for the 9900 Device

Parameter	Required Value
Queue depth per LU (TagQueueDepth)	≤ 32 per LU
Queue depth per port	≤ 256 per port

Note: You can adjust the queue depth for the 9900 devices later as needed (within the specified range) to optimize the I/O performance of the 9900 devices.

3.3 Setting I/O Timeout Value

In Compaq Tru64 UNIX 4.0, the non-I/O timeout is 10 seconds. To change this to 60 seconds, edit the /sys/data/cam_data.c file as shown in Figure 3.3 and change u_long cdisk-to-def to **60.**

```
# cp /sys/data/cam_data.c /sys/data/cam_data.c.backup↵          ← Back up the file.
# vi /sys/data/cam_data.c↵
      :
      :
/*
 * Changeable disk driver timeouts. Cdisk_to_def for non read/write
 * commands to the disks (test unit ready, mode select etc.)
 * cdisk_io_def Hard fixed disk timeout value for i/o, and
 * cdisk_io_rmb for removable media disks since they are slow.
 */
u_long cdisk_to_def = 60;     /* 60 seconds */        ← Change 10 [sec] to 60 [sec].
u_long cdisk_io_def = 60;     /* 60 seconds (Tagged commands) */
u_long cdisk_io_rmb = 120;    /* 120 seconds (slow removables) */
      :
      :
 : wq                                                 ← Enter : wq! to save and exit.
#
```

Figure 3.3 Changing the I/O Timeout Value

```
# cp /sys/data/cam_data.c /sys/data/cam_data.c.backup↵  ····· ← Back up the file.


# vi /sys/data/cam_data.c↵
                                              :
                                              :
/*
 * Changeable disk driver timeouts. Cdisk_to_def for non read/write
 * commands to the disks (test unit ready, mode select etc.)
 * cdisk_io_def Hard fixed disk timeout value for i/o, and
 * cdisk_io_rmb for removable media disks since they are slow.
 */
u_long cdisk_to_def = 60;     /* 60 seconds */          ← Change 10 {sec} to 60 {sec}.
 u_long cdisk_io_def = 60;    /* 60 seconds (Tagged commands) */
u_long cdisk_io_rmb = 120;    /* 120 seconds (slow removables) */
                                              :
                                              :
                                        ·····
#:wq                                                  ← Enter : wq! To save and exit.
```

Figure 3.4 Reconfiguring the Kernel

3.4 Creating the File Systems

Now you can create a file system for the new 9900 SCSI disk devices. If desired, you can create and use an advanced file system (AdvFS) to overcome the size and speed limitations of the UNIX file system. The file system is application-dependent. If you are not sure which file system is right for your setup, contact Hitachi Data Systems or Compaq customer support.

Note: Do not create file systems for the multiplatform devices (e.g., 3390-3A/B/C).

3.4.1 UNIX File System (UFS)

To create a UNIX file system:

1. Create a UNIX file system for a new device by entering: **`newfs <device_file_name>`**
   Note: Use the character-type device file name (rrzXYZ).

   For example, to create the file system for device **rrz12c** (SCSI bus 1, TID 4, LUN 0, partition c), enter: **`newfs /dev/rrz12c`**

2. Repeat step (1) for each UNIX file system you need to create.

3.4.2 Advanced File System (AdvFS)

To create an advanced file system:

1. You can either create a new AdvFS domain, or you can add a new fileset to an existing domain. To create a new domain, enter: `mkfdmn <device_file_name> <domain_name>`
   Note: **Use** the block-type device file name (rzXYZ).

   For example, to create an AdvFS domain for device **rz12c**, enter:
   `mkfdmn /dev/rz12c domain1`

2. Create a new fileset in the new or existing AdvFS domain by entering:
   `mkfset <domain_name> <fileset_name>`

   For example, to create fileset1 in domain1 enter: `mkfset domain1 fileset1`

3. Repeat steps (1) and (2) for each advanced file system you need to create.

3.5 Creating the Mount Directories

After you have created the file systems, the next step in configuring the 9900 devices is to create a mount directory for each SCSI disk device. Make sure to choose a unique name for each mount directory, which identifies the device being mounted.

Note: Do not create mount directories for the multiplatform devices (e.g., 3390-3A/B/C).

To create the mount directories for the new devices:

1. Create a mount directory by entering: **`mkdir /<mount_directory_name>`**

   For example, to create a mount directory for LUN 2 (partition c) on the 9900, enter: **`mkdir /9900_LU2c`**

Verify the new mount directory by entering (see Figure 3.5): **`ls -x`**

2. Repeat steps (1) and (2) for each new SCSI disk device on the 9900.

```
#ls -x                                                    ← Verify new mount directory.
.mrg                ..DXsession         9900_LU2c         ← New mount directory appears.
CDROM               DXclock             DXprint
DXsessionMail       a                   bin
cdfsdev             df300               df300-2
df300-3df300-4      df300-5             etc
homel               lib                 lnn
mdecmnt             opt                 osf_boot
picassoproc         real.profile        sbin
subsyssys           tcb                 tmp
usrvar              vmunix              vmunix.stdrd
```

Figure 3.5 Verifying the Mount Directories

3.6 Mounting the File Systems

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

To mount a UNIX file system:

1. Mount the file system by entering:
   ```
   mount <device_file_name> <mount_directory>
   ```
   For example, to mount device **rz12c** with mount directory name **9900_LU2c**, enter:
   ```
   mount /dev/rz12c /9900_LU2c
   ```
2. Assign the appropriate ownership and permissions by entering:
   ```
   chown <owner>:<group> *<device_file_name>*
   ```
   For example, to assign ownership to **rz12c** with owner Oracle, group dba, enter:
   ```
   chown oracle:dba *rz12c*
   ```
3. Repeat steps (1) and (2) for each new UNIX file system.

3.6.1 Mounting the Advanced File Systems

To mount an advanced file system:

1. Mount the advanced file system by entering:
   ```
   mount -t advfs <domain_name>#<fileset_name> <mount_directory>
   ```
   For example, to mount the file system with mount directory 9900_LU2c enter:
   ```
   mount  -t advfs domain1#fileset1 /9900_LU2c
   ```
2. Assign the appropriate ownership and permissions by entering:
   ```
   chown <owner>:<group> *<device_file_name>*
   ```
   For example, to assign ownership to **rz12c** with owner Oracle, group dba, enter:
   ```
   chown oracle:dba *rz12c*
   ```
3. Repeat steps (1) and (2) for each new advanced file system.

3.7 Verifying the File Systems

Now you need to verify that the new file systems were created correctly and are functioning properly.

To verify the new file systems:

1. Verify that the new file systems were created correctly by displaying all mounted file systems. To display all mounted file systems, enter (see Figure 3.6): **df**

   Note: The default display for drive capacity is 512-byte blocks. To view the capacity in kilobytes rather than in 512-byte blocks, enter: `df -k`

2. Verify the operation of each new file system as follows (see Figure 3.7):

   a) Go to the new device directory by entering: `cd /<mount_directory>`
   For example, `cd /9900_LU2c`

   b) Copy a file from the root directory to the new device by entering:
   `cp /<file_name> <file_name>.back1`

   For example, to copy file **vmunix** from the root directory to the 9900_LU2c device, enter: `cp /vmunix vmunix.back1`

   c) Copy a file to the new device again. For example, to copy the same file again, enter:
   `cp /vmunix vmunix.back2`

   d) List the files in the current directory by entering: `ls -l`

   e) Remove the files you copied by entering: `rm <file_name>`
   For example, enter **rm vmunix.back1** to remove the file **vmunix.back1**.

3. Repeat step (2) for each new file system.

```
# df -k↵
Filesystem       1024-blocks        Used        Available Capacity   Mounted on
root_domain#root       65536       53652             4824       92%            /
/proc                      0           0                0      100%            /proc
usr_domain#usr       1968128      437777          1512776       23%            /usr
/dev/rz8c            2328100           1          2095289        0%            /mnt
dmn11#set1           2403360          16          2398816        1%            /mnt1
#
```

Note: This example shows an OPEN-3 file system. An OPEN-9 file system is not shown.

Figure 3.6 Verifying the File Systems for Tru64 UNIX 4.0 and 4.0F

```
# mount /dev/rz10c /mnt↵
  (#mount /dev/disk/dsk10c /mnt)
# cd /mnt↵
# cp /genvmunix /mnt/genvmunix.back1↵
# ls -l↵
total 7240
-rwxr-xr-x   1   root  system  7397832 Nov 14 15:08  genvmunix.back1
# cp genvmunix.back1 genvmunix.back2↵
# ls -l↵
total 14480
-rwxr-xr-x   1   root  system  7397832 Nov 14 15:08  genvmunix.back1
-rwxr-xr-x   1   root  system  7397832 Nov 14 15:12  genvmunix.back2
# rm genvmunix.back1 genvmunix.back2↵
#
```

Figure 3.7 Verifying File System Operation

3.8 Setting and Verifying the Auto-Mounting Parameters

The final step in configuring the new 9900 devices is to set up and verify the auto-mount parameters for each new device. The **/etc/fstab** file contains the auto-mount parameters for the logical devices.

To set and verify the auto-mount parameters:

1. Edit the **/etc/fstab** file to add a line for each new OPEN device to be auto-mounted (see Figure 3.8). Table 3.2 shows the auto-mount parameters needed for this procedure.

   Note: If you make a mistake while editing, exit the **vi** editor without saving the file, and then begin editing again. See Appendix B for reference information on using the **vi** editor.

2. Reboot the Tru64® UNIX® system after you are finished editing the **/etc/fstab** file.

3. Verify that the device file systems were auto-mounted using the **df** (or **df -k**) command.

```
cp /etc/fstab /etc/fstab.standard                                  ← Make backup
file.
#vi /etc/fstab                                                     ← Start vi editor.
/root_domain#root          /          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     /9900_LU2c  ufs        rw        1       3

                                                                   ← See Table 3.1.
```

Figure 3.8 Setting the Auto-Mount Parameters

Table 3.2 Auto-Mount Parameters

Parameter #	Name	Enter
	Device to mount	Block type device filename
	Mount point	Mount directory name
	File system	Type of file system (e.g., UFS, AdvFS)
	Mount options	Options (i.e., rw for read-write)
	Frequency dump in days	# of days (e.g., 1, 2, 3)
	File system check (fsck)	Order of performing file system checks

Chapter 4 Middleware and SNMP Configuration

The 9900 subsystem supports many industry-standard middleware products, which provide host failover, I/O path failover, and logical volume management functions. For Tru64™, the 9900 supports the following middleware products and functions:

- HACMP and HAGEO for host failover (see section 4.1),
- TruCluster™ for path failover

The 9900 disk subsystem also supports the industry-standard simple network management protocol (SNMP) for remote subsystem management from the open-system host. SNMP is used to transport management information between the 9900 SVP and the SNMP manager on the host. The SNMP agent on the SVP sends status information to the host(s) when requested by the host or when a significant event occurs.

4.1 Host Failover

The user must make sure to configure the Tru64™ software and any other middleware products as needed to recognize and operate with the newly attached 9900 devices. For assistance with Tru64™ operations, please refer to the Compaq user documentation, or contact Compaq technical support. For assistance with specific configuration issues related to the 9900 subsystem, please contact the Hitachi Data Systems Technical Support Center (see section 4.2). ***Note:*** TruCluster™ does not provide a complete disaster recovery or backup solution, and is not a replacement for standard disaster recovery planning and backup/recovery.

4.2 SNMP Remote Subsystem Management

SNMP is a part of the TCP/IP protocol suite that supports maintenance functions for storage and communication devices. The 9900 subsystem utilizes SNMP to transfer status and management commands to the open-system host via the 9900 SVP (see Figure 4.1). When the SNMP manager requests status information or when a service information message (SIM) occurs, the SNMP agent on the 9900 SVP notifies the SNMP manager on the open-system server. Notification of 9900 error conditions is made in real time, providing the open-system user with the same level of monitoring and support available to the mainframe user. The SIM reporting via SNMP enables the user to monitor the 9900 subsystem from the open-system host without having to check the Remote Console PC for SIMs.

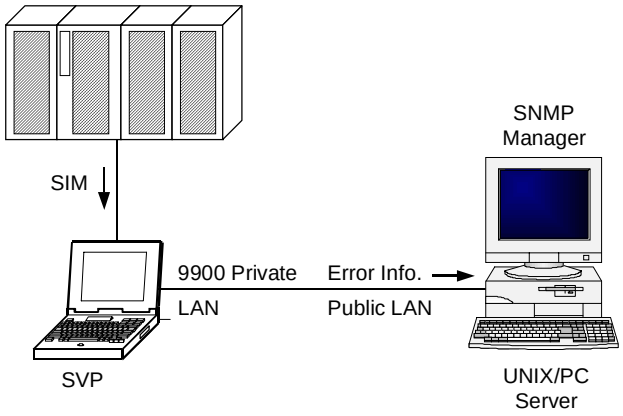

Figure 4.1 9900 SNMP Environment

When a SIM occurs, the 7700E SNMP agent initiates trap operations, which alert the SNMP manager of the SIM condition. The SNMP manager receives the SIM traps from the SNMP agent, and can request information from the SNMP agent at any time.

For assistance with SNMP manager configuration on the open-system host, please refer to the user documentation for the SNMP management software, or contact the vendor's technical support. For assistance with specific SNMP configuration issues related to the 9900 subsystem, please contact your Hitachi Data Systems representative.

Chapter 5 Troubleshooting

5.1 Troubleshooting

For troubleshooting information on the 9900 subsystem, please refer to the *Hitachi Freedom 9900 User and Reference Guide* (MK-90RD008). Table 5.1 lists potential error conditions during 9900 Tru64™ configuration and provides instructions for resolving each condition. If you are unable to resolve an error condition, please ask your Hitachi Data Systems Customer Service representative for help, or call the Hitachi Data Systems Support Center for assistance.

Table 5.1 Troubleshooting

Error Condition	Recommended Action
The logical devices are not recognized by the system.	Make sure that the READY indicator lights on the 9900 subsystem are ON. The LUNs for each SCSI TID must start at 0 and continue sequentially without skipping any numbers. Make sure that LUSE devices are not intermixed with normal devices or with HMDE devices on the same Fibre channel port.
The file system is not mounted after rebooting.	Make sure the system was restarted properly. Make sure that the values listed under file system are correct.
System hangs, or devices are declared and then the system hangs.	Make sure the target IDs are set 0 through 6 and 8 through 15 and target ID 7 has been reserved for the SCSI controller card.

5.2 Calling the Hitachi Data Systems Technical Support Center

If you need to call the Hitachi Data Systems Technical Support Center, make sure to provide as much information about the problem as possible. Include the circumstances surrounding the error or failure, the exact content of any messages displayed on the remote console PC, and the severity levels and reference codes of the R-SIMs on the R-SIM panel. The worldwide Hitachi Data Systems Technical Support Centers are:

- Hitachi Data Systems North America/Latin America
  San Diego, California, USA
  1-800-348-4357
- Hitachi Data Systems Europe
  Contact Hitachi Data Systems Local Support
- Hitachi Data Systems Asia Pacific
  North Ryde, Australia
  011-61-2-9325-3300

Appendix A Acronyms and Abbreviations

AL-PA	arbitrated-loop physical address
CVS	custom volume size
ESCON®	Enterprise System Connection (IBM® trademark for optical channels)
FC	fibre channel
HACMP	High Availability Cluster Multi-Processing
HAGEO	High Availability Geographic
HMDE	Hitachi Multiplatform Data Exchange
I/O	input/output
LDEV	logical device
LU	logical unit
LUN	logical unit number
LUSE	LU Size Expansion
PC	personal computer system
r/w	read/write
SCSI	small computer system interface
SIM	service information message
SMIT®	System Management Interface Tool
SNMP	simple network management protocol
SVP	service processor
TID	target ID
VOLSER	volume serial number

Appendix B SCSI TID Maps for Fibre-Channel Adapters

When an arbitrated loop (AL) is established or re-established, the port addresses are assigned automatically to prevent duplicate TIDs. When using 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 since 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, once control of the loop is acquired, a point-to-point connection is established from initiator to target. To enable transparent use of FCP, the TRU64™ operating system "maps" a TID to each AL-PA.

Note: For specific information on AL-PA to TID mapping, please refer to the Compaq Tru64 user documentation and/or fibre channel adapter.

Appendix C Online Installation/Deinstallation of D evices

After initial installation and configuration of the 9900 subsystem, additional devices can be installed and de-installed online without having to restart the system.

Note: For additional instructions regarding online installation and deinstallation of devices, please refer to the *Hitachi Lightning 9900™ LUN Manager User's Guide* (MK-90RD006).

C.1 Online Device Installation

This section provides instructions for installing and configuring new 9900 SCSI disk devices without rebooting the Tru64® UNIX® system.

Note: If a new fibre-channel connection is being installed, you must use the normal disruptive device configuration procedure. New fibre-channel connections can only be installed when the host system is powered off. New devices under existing fibre-channel ports can be installed and configured nondisruptively.

To perform online device installation and configuration:

1. Verify that the new devices on the 9900 subsystem are ready to be configured. The Hitachi Data Systems representative should have completed hardware installation and verified the normal status of the new devices (refer to section 2.2).
2. Make sure that you are logged in as **root**.
3. Perform online device recognition as shown in Figure C.1.
4. Verify that the Tru64® UNIX® system recognizes the new devices as described in section 2.5 (refer to Figure 3.2).
5. Configure the new 9900 disk devices for Tru64® operations as described in Chapter 3. For new HMDE devices, you only need to verify the device files. Do not partition or create a file system on any HMDE device.
6. Remember to configure the application failover, path failover, and/or SNMP software on the Tru64® UNIX® system as needed to recognize the new devices.

```
# scsimgr -scan_all ↵
# sizer -n test ↵                ← test,test.devs. will be made on /tmp.
# cd /tmp ↵                      ← Change tmp directory.
# vi test.devs ↵                 ← Add /dev/ to head of test.devs file.
# chmod +x test.devs ↵           ← Add executable attribute to test.devs file.
# cd /dev ↵                      ← Change dev directory.
# /tmp/test.devs ↵
```

Figure C.1 Online Volume Recognition