HITACHI
DATA SYSTEMS

Hitachi Freedom Storage™
Lightning 9900™

Windows NT® Configuration Guide

Document Revision Level

Revision	Date	Description
MK-90RD015-P	May 2000	Preliminary Release.
MK-90RD015-P1	June 2000	Rev. 1 of Prelim. Release, supersedes and replaces MK-90RD015-P.
MK-90RD015-0	July 2000	Rev. 0 supersedes and replaces MK-90RD015-P1.
MK-90RD015-1	November 2000	Rev. 1 supersedes and replaces MK-90RD015-0.

Source Document Revision Level

The following source document was used to produce this Lightning 9900 configuration guide: DKC310 Disk Subsystem SCSI/Fibre Installation Manual, Windows NT 4.0, rev. 4.

Preface

The *Hitachi Lightning 9900 Windows NT Configuration Guide* describes and provides instructions for configuring the devices on the Hitachi Lightning 9900 array subsystem for operation with the Microsoft Windows NT 4.0 operating system (OS). This configuration guide assumes that:

- the user has a background in data processing and understands direct-access storage device (DASD) subsystems and their basic functions,
- the user is familiar with the Hitachi Lightning 9900 array subsystems, and
- the user is familiar with the Microsoft Windows NT Server 4.0 and/or Windows NT Workstation 4.0 operating systems, the NT server/workstation computer, and the fibre-channel adapters.

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

For further information on the Lightning 9900 array subsystem, please refer to the *Hitachi Lightning 990™0 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 Windows NT 4.0, please consult the Windows NT 4.0 online help and/or user documentation, or contact Microsoft technical support.

Contents

List of Figures

List of Tables

Chapter 1 Overview of Lightning 9900™ Windows NT® Configuration

1.1 Windows NT® Configuration

This document describes the requirements and procedures for connecting the 9900 subsystem to a Novell® server and configuring the new 9900 devices for operation with the Windows NT® server operating system. The Hitachi Data Systems representative performs the physical installation of the 9900 subsystem. The user prepares for 9900 subsystem installation, and then configures the new 9900 devices with assistance as needed from the Hitachi Data Systems representative.

Configuration of the 9900 disk devices for Windows NT® operations includes:

- Writing the signatures (see section 3.1),
- Creating and formatting the partitions (see section 3.2),
- Verifying file system operations (see section 3.3), and
- Verifying client operations (see section 3.4).

Note on the term "SCSI disk": The 9900 logical devices are defined to the host as SCSI disk devices, whether the interface is SCSI or fibre-channel.

1.2 Hitachi Lightning 9900™ Array Subsystem

The Hitachi Lightning 9900(tm) RAID subsystem supports concurrent attachment to multiple UNIX(r)-based and PC-server 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 can be configured with fibre-channel ports and/or serial interface ports (compatible with ESCON® protocol) to provide connectivity with S/390(r) mainframe hosts as well as UNIX®/PC-server hosts. For further information on the 9900 subsystem, please refer to the Hitachi Freedom Storage(tm) Lightning 9900™ User and Reference Guide (MK-90RD008), or contact your Hitachi Data Systems account team

Note on the term "SCSI disk": The 9900 logical devices are defined to the host as SCSI disk devices, even though the interface is fibre-channel.

1.3 Device Types and Configuration Procedures

The 9900 subsystem allows the following types of logical devices (also called volumes) to be installed and configured for operation with the Windows NT operating system. Table 1.1lists the device specifications for the 9900 devices. Table 1.2 shows the volume usage (i.e., file system or raw device) for the 9900 devices.

OPEN-x Devices. The OPEN-x logical units (LUs) (e.g., OPEN-3, OPEN-9) are disk devices of predefined sizes. The 9900 subsystem currently supports OPEN-3, OPEN-8, OPEN-9, 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 combined LUs which can be from 2 to 36 times larger than standard OPEN-x LUs. The Logical Unit Size Expansion (LUSE) feature of the 9900 subsystem enables you to configure these custom-size devices. LUSE devices are designated as OPEN-*x**n, where *x* is the LU type (e.g., OPEN-9*n) and $2 \leq n \leq 36$. For example, a LUSE device created from ten OPEN-3 LUs would be designated as an OPEN-3*10 disk device. This capability enables the server host to combine logical devices and access the data stored on the 9900 subsystem using fewer LU numbers (LUNs) For further information on the LUSE feature, please refer to the *Hitachi Lightning 9900™ LUN Manager User's Guide*, (MK-90RD006).

CVS Devices (OPEN-x CVS). The CVS devices are disk devices which are used exclusively by the NT host system. The Custom Volume Size (CVS) feature of the 9900 subsystem (also called Virtual LUN) enables you to configure custom-size LUs that are smaller than standard OPEN-x LUs. This capability enables you to "slice up" a single LU into several smaller LUs to best fit the application needs and improve host access to frequently used files. For further information on the Virtual LUN/CVS feature, please refer to the *Hitachi Lightning 9900™ Virtual LVI/LUN User's Guide* (MK-90RD005).

CVS LUSE Devices (OPEN-x*n CVS). The CVS LUSE devices combine CVS devices (instead of standard OPEN-x LUs) into LUSE devices. The CVS feature is used to create CVS devices, and then the LUSE feature is used to combine (concatenate) these CVS devices. The user can combine from 2 to 36 CVS devices into one CVS LUSE device. For example, an OPEN-3 LUSE volume that was created from ten OPEN-3 CVS volumes would be designated as an OPEN-3*10 CVS device.

Configuration of the 9900 disk devices for Windows NT operations includes:

- Verifying new device recognition (see section 2.7),
- Writing the signatures on the new disks (see section 3.1),
- Creating and formatting disk partitions (see section 3.2),
- Verifying file system operations (see section 3.3), and
- Verifying auto-mount (see section 3.4).

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 S/390®, UNIX, and PC server 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-90RD020), or contact your Hitachi Data Systems account team.

The HMDE devices are not SCSI disk devices. The HMDE devices must be installed and accessed 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 UNIX/PC server access. The 9900 subsystem will reject all UNIX/PC server write operations (including fibre-channel adapters) 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 UNIX/PC server access. Do not execute any write operation by the fibre-channel adapters on these devices. Do not create a partition or file system on these devices. This will overwrite any data on the HMDE device and also prevent the HMDE software from accessing the device. Do not write a signature on the HMDE devices, unless the devices will be operated in the Microsoft Cluster Server (MSCS) environment.

Configuration of the 9900 HMDE devices for operation with Windows NT includes:

- Verifying new device recognition (see section 2.7).

For Microsoft Cluster Server (MSCS) environments only, you must also write signatures (see section 3.1) on the HMDE devices. For non-MSCS environments, **DO NOT** write signatures on the HMDE devices.

WARNING: After a signature has been written on an HMDE device (MSCS environment only), there is no way to distinguish the HMDE device from a SCSI disk device. The user must exercise extreme caution not to accidentally partition and format an HMDE device. This will overwrite any data on the HMDE device and also prevent the HMDE software from accessing the device.

Table 1.1 9900 Device Specifications for Windows NT Operations

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-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***
OPEN-E*n	SCSI disk	HITACHI	OPEN-E*n	28452960*n	512	19759*n	15	96	13888*n
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	5820300	512	3345	15	116	2844
3380-KA	HMDE otm/mto	HITACHI	3380-KA	3833280	512	2662	15	96	1873
3390-3B	HMDEmto	HITACHI	3390-3B	5816820	512	3343	15	116	2842
3380-KB	HMDEmto	HITACHI	3380-KB	3830400	512	2660	15	96	1871
3390-3C	HMDEotm	HITACHI	OP-C-3390-3C	5820300	512	3345	15	116	2844
3380-KC	HMDEotm	HITACHI	OP-C-3380-KC	3833280	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	HMDEotm	HITACHI	OP-C-3390-3C-CVS	***Note 4***	512	***Note 5***	15	116	***Note 6***

3380-KC CVS	HMDEotm	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 1.2 shows the volume usage for SCSI disk devices and HMDE devices. The SCSI disk devices (OPEN-x, CVS, LUSE, and CVS LUSE) require partitions and file systems for Windows NT 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 create a partition or file system on any device used for HMDE operations. Do not write a signature on an HMDE device unless it is used in an MSCS environment.

Note on the term "SCSI disk": The 9900 logical devices are defined to the host as SCSI disk devices, even though the interface is fibre-channel.

Table 1.2 Volume Usage for Device Categories

Category	Device Type	Volume Usage
SCSI Disk	OPEN-x, OPEN-x CVS, OPEN-x*n LUSE, OPEN-3*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, OPEN-x CVS for OPEN-x-HMDEoto:	Raw Device

Note 3: The device capacity can sometimes be changed by the BIOS or host adapter board. Also, different capacities may be due to variations such as 1 MB = 1000^2 or 1024^2 bytes.

Note 4: 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 5) × (15 heads) × (96 sectors per track) = 76320

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

- The number of data cylinders for an OPEN-x CVS volume =
  # of cylinders = ↑ (capacity (MB) specified on 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 CVS LUSE volume =
  # of cylinders = ↑ (capacity (MB) specified on Remote Console PC) × 1024/720 ↑ × n
  Example: For an OPEN-3 CVS LUSE volume with capacity = 37 MB and n = 4
  # of cylinders = ↑37 × 1024/720↑ × 4 = ↑52.62↑× 4 = 53 × 4 = 212
- The number of data cylinders for a 3390-3A/C or 3380-KA/C CVS volume =
  # of cylinders = (number of cylinders specified on 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 Remote Console PC) + 7

Note 6: 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

2.1 Configuration Requirements

The requirements for 9900 Windows NT 4.0 configuration are:

- Hitachi Lightning 9900™ subsystem, all-open or multiplatform configuration.
  - The 9900 Remote Console PC and LUN Manager software are used to define the LUN mapping and configure the fibre-channel (FC) ports. If the remote LUN Manager feature is not installed, the Hitachi Data Systems representative can configure the LUN mapping and configure the fibre channel ports using the 9900 service processor (SVP). For information on LUN configuration services, please contact your Hitachi Data Systems account team.

  Note: The availability of 9900 features and devices (e.g., OPEN CVS, 3390-3C, OPEN-8) depends on the level of microcode installed on the 9900 subsystem.

- Windows NT server/workstation. Please refer to the Microsoft user documentation for PC server hardware requirements.

- 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).

- Windows NT Workstation or Windows NT Server operating system. ***Important:*** Please contact Microsoft to make sure that the most current OS patches are installed.

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

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-channel support.
   - Installing and formatting the logical devices (LDEVs) using the SVP. Make sure to get the desired LDEV configuration information from the user, including the desired number of OPEN-x, LUSE, CVS, CVS LUSE, and multiplatform (HMDE) devices.
   - Installing the fibre-channel adapters and cabling: The total fibre cable length attached to each fibre-channel adapter must not exceed 500 meters (1,640 feet). Do not connect any OFC-type connector to the 9900 subsystem. Do not connect/disconnect fibre-channel cabling that is being actively used for I/O. This can cause the Windows NT system to hang. Always confirm that the devices on the fibre cable are offline before connecting/disconnecting the fibre cable.

     9900 FC Port: The fibre topology parameters for each 9900 fibre-channel port depend on the type of device to which the 9900 port is connected. Determine the topology parameters supported by the device, and set your topology accordingly (see section 2.3.4). The type of 9900 port is also important.

   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.

2. **Remote console PC and LUN Manager installation**. The user or Hitachi Data Systems representative can perform this activity. You will use the LUN Manager software on the 9900 Remote Console PC to define the fibre channel ports for the 9900 devices. For instructions on installing the Remote Console PC and LUN Manager remote console software, please refer to the *9900 Remote Console User's Guide* (MK-90RD003) and the *9900 LUN Manager User's Guide* (MK-90RD006) respectively.

   Note: If the remote LUN Manager feature is not installed, the Hitachi Data Systems representative can define the fibre channel ports for you using the SVP of the subsystem. Please contact your Hitachi Data Systems account team for further information on LUN configuration services.

2.3 Preparing to Connect the 9900 Subsystem

Before the 9900 is connected to your NT 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 9900 Windows NT ® operations is **0C.** Use the LUN Manager remote console software to ensure that the host mode for each port is connected to the Windows NT operation (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.

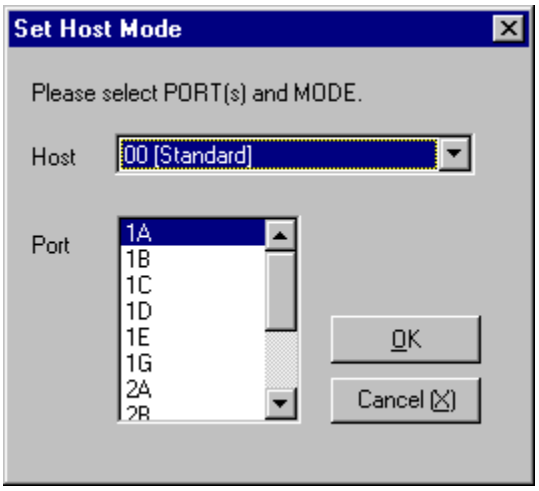

Figure 2.1 Setting the Host Mode

2.3.2 Configuring the 9900 Fibre-Channel Ports

You need to configure the 9900 FC ports to define the fibre parameters (see Figure 2.2 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 Windows NT® operations. Please see your Hitachi Data Systems account team regarding the number of supported LUs.

Fibre topology. Figure 2.2 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.

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 Windows NT® 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	*Not supported*
OFF	FC-AL	AL-port (private arbitrated loop)
OFF	Point-to-Point	*Not supported*

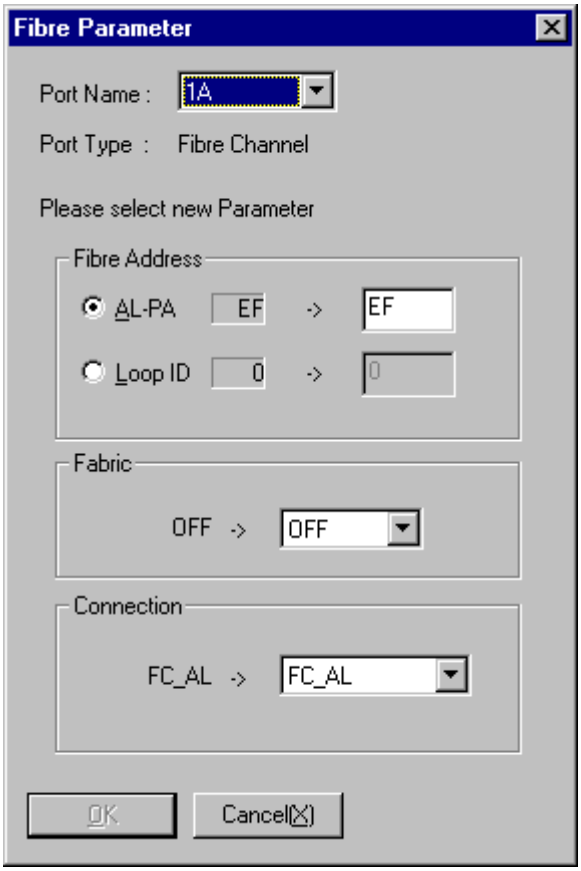

Figure 2.2 Fibre Parameter Panel (from the LUN Manager Software)

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 and Recording the Disk Numbers

After you have defined the fibre channel ports for the 9900 devices and configured the 9900 fibre-channel ports, you are ready to connect the 9900 subsystem to the Windows NT system. The 9900 subsystem comes with all hardware and cabling required for connection to the host system(s).

To connect the 9900 subsystem to the NT 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 channel ports and fibre device parameters to make sure that all 9900 devices are unique for each host system.
2. **Shut down and power off the NT system**. The user should perform this activity. You must shut down and power off the NT system before connecting the 9900.
   a) Shut down the NT system as usual.
   b) When shutdown is complete, power off the NT display.
   c) Power off all peripheral devices except for the 9900 subsystem.
   d) Power off the NT system. You are now ready to connect the 9900 subsystem.
3. **Connect the 9900 to the NT system**. The Hitachi Data Systems representative installs the fibre-channel cables between the 9900 and the NT system. ***Note:*** The Hitachi Data Systems representative must use the 9900 Maintenance Manual during all installation activities. Follow all precautions and procedures, and check all specifications to ensure proper installation and configuration.
4. **LUN security**. If the 9900 is connected to fabric, you must control access to the 9900 LUs before you power on the NT system to prevent failures on other systems connected to the same fabric. You can use any of several methods for ensuring that the NT system sees only the LUs it owns (e.g., LUN Security, set FC adapters not to automap LUNs, switch zoning). For further information on using the 9900 LUN Security remote console software, please refer to the *9900 LUN Manager User's Guide* (MK-90RD006).
5. **Power on and start booting the NT system** (if you powered off in step 2). The user should perform this activity. To power on the NT system after connecting the 9900:
   a) Power on the NT system display.
   b) Power on all peripheral devices. The 9900 should already be on. The fibre channel ports should already be defined. If not, the NT 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 NT system connected to the 9900.

6. **Record the disk numbers**. When the adapter connected to the 9900 starts displaying the new devices, pause the screen and record the disk number for each new device on your SCSI Device Information worksheet (refer to Table 2.5). If you did not have to shut down, scan the fibre-channel ports for new devices and record the disk numbers. You will need the disk numbers for the devices when you write signatures on the devices (see section 3.1).

   Note: You will also need the disk numbers for the HMDE devices when you create the HMDE volume definition file (**datasetmount.dat**). For example, if disk number 3 is a 3390-3B HMDE device, the entry for this volume in the HMDE volume definition file is:
   `\\.\PHYSICALDRIVE3 volser 3390-3B` ('volser' = mainframe volume serial number)

7. **Reboot**. After recording the disk numbers, reboot the NT system and get ready to access the fibre-channel adapter utilities while the system is booting up.

Note: The NT system assigns the disk numbers sequentially starting with the local disks and then by adapter and by TID/LUN. If the 9900 is attached to the first adapter (displayed first during system start-up), the disk numbers for the new devices will start at 1 (the local disk is 0). If the 9900 is not attached to the first adapter, the disk numbers for the new devices will start at the next available disk number. For example, if 40 disks are attached to the first adapter (disks 1-40) and the 9900 is attached to the second adapter, the disk numbers for the 9900 will start at 41.

Note: When disk devices are added to or removed from the NT system, the disk numbers are reassigned automatically. For the HMDE devices, make sure to update your HMDE volume definition file (**datasetmount.dat**) with the new disk numbers.

2.5 Configuring the Host Fibre-Channel Adapters

After connecting the 9900 subsystem and recording the disk numbers for the new devices, you are ready to configure the fibre-channel adapter(s) connected to the 9900. The HBA setup utilities allow you to configure the adapter settings while the system is booting up. The host bus adapters have many configuration options. This section provides the following minimum requirements for configuring FC adapters for operation with the 9900 subsystem.

- The disk I/O timeout value (TOV) requirement for the 9900 is 60 seconds (0x3c hex).
- The queue depth requirements for the 9900 devices are specified in Table 2.3.
- The BIOS may need to be disabled to prevent the system from trying to boot from the 9900. ***Note:*** If you want to configure the NT boot disk on the 9900 subsystem, see section 2.5.1 and Appendix C.
- In addition to the disk I/O TOV, queue depth, and BIOS, several other parameters (e.g., FC fabric) may also need to be set. Please refer to the user documentation, which came with your HBA to determine whether other options are required to meet your operational requirements.

Note: Make sure to use the same settings and device parameters for all 9900 devices.

The following sample instructions apply to the QLogic 2100F FC adapter. For instructions on configuring other adapters, refer to the user documentation for the adapter. Table 2.4 shows the FC adapter configuration requirements for the NT boot disk and HORC/HOMRCF volumes.

Note: If your HBA does not have a setup utility, or if your HBA setup utility does not provide access to the required parameters, you must use the Windows NT Registry Editor to set the required parameters. See section 2.6 for instructions on configuring the adapter settings using the Registry Editor.

Table 2.3 Queue Depth Requirements for the 9900 Devices

Parameter	Required Value
Queue depth per LU	queue-depth ≤ 32 per LU
Queue depth per port (MAXTAGS)	queue-depth ≤ 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.

Table 2.4 Fibre Adapter Configuration Requirements: Boot Disk and HORC/HOMRCF Volumes

Volumes under Host FC Adapter			BIOS Settings	
NT Boot Disk	HORC/HOMRCF	Availability	QLogic 2100F	Emulex LightPulse
No	No	Allowed	Host Adapter BIOS = Disabled	default
No	Yes	Allowed	Host Adapter BIOS = Disabled	default
Yes	No	Allowed*	Host Adapter BIOS = Disabled*	Host Adapter BIOS = Enabled
Yes	Yes	Not allowed	—	—

*The NT boot disk cannot be configured under the QLogic 2100 adapter.

To configure a QLogic 2100F fibre-channel adapter connected to the 9900:

1. While the NT system is booting up, launch the HBA setup utility as follows: when the message **Press <Alt-Q> to Run QLogic Fast! Utility** appears, press **Alt-Q**.
2. Select the QLogic adapter to configure.
3. Go to **<Configuration Settings>**, select **<Host Adapter Settings>**, and then verify the following settings: Host Adapter BIOS: Disabled
   Frame Size: 2048
   Execution Throttle: See Table 2.3 *(execution throttle = queue depth)*
4. Select **<Adapter Hard ID Settings>**, and then verify the following settings:
   Adapter Hard ID: Enabled
   Hard ID: Less than all 9900 fibre port addresses
5. Verify all other required settings for your operational environment. For example, the QLogic adapter defaults to eight LUNs per target, so you may need to change that setting. Refer to the user documentation for the adapter as needed.
6. Repeat steps (2)-(5) for each QLogic FC adapter connected to the 9900 subsystem. When you are finished configuring QLogic adapters, exit the HBA setup utility.

WARNING for Emulex FC adapter with Intel Pentium Pro PCI chipsets

On Windows NT 4.0 systems using the Intel Pentium Pro PCI chipsets, the Emulex adapter is not recognized by the secondary power-control interface (PCI) bus of a dual peer-bus system. This problem is caused by the Microsoft HAL improperly assigning resources to some PCI devices. Microsoft provides a workaround for this problem (for NT 4.0) which causes the HAL to use the BIOS-assigned defaults rather than reassign PCI resources. This workaround involves editing the NT system's **boot.ini** file.

Note: You may need to remove the read-only file attribute in order to edit the **boot.ini** file. You can use the Windows **File-Properties** panel or the DOS **attrib** command (e.g., attrib -r -h -s c:\boot.ini) to remove the read-only file attribute.

Edit the **boot.ini** file as follows to enable both PCI buses to recognize the Emulex FC adapter:

1. Use a text editor (e.g., Notepad) to open the **boot.ini** file.
2. Add the **/PCILOCK** option to the system boot entry, and then save your changes.
3. Close the **boot.ini** file. You must reboot the system for these changes to take effect.

2.5.1 Configuring the NT Boot Disk on the 9900 Subsystem

To configure the NT boot disk on the 9900 subsystem, your system must have the following:

- Emulex LightPulse HBA for the 9900 subsystem.
- Emulex LightPulse HBA firmware version 3.00 or later. The firmware must contain boot BIOS version 1.20 or later and POST version 3.00 or later.
- Fibre-channel topology must be FC-AL (private loop technology).

To download the firmware, see http://www.emulex.com. For further information on setting up the NT system to boot from the 9900, please contact your Hitachi Data Systems representative.

2.6 Verifying the Disk and Device Parameters

After you have configured the fibre-channel adapters during boot-up, you need to verify the required disk and device parameters using the Windows NT Registry. You must verify the disk I/O timeout value (TOV) and the queue depth, and you should also verify other required parameters such as FC fabric support, and link down timeout.

2.6.1 Verifying the Disk I/O Timeout Value (TOV)

The disk I/O TOV parameter, which applies to all SCSI disk devices attached to the NT system, must be set to 60 seconds. The default setting is hexadecimal 0x3c, which is decimal 60. ***Note:*** The driver installation for the Emulex FC adapter sets the disk I/O TOV to 60, if there was no previously assigned value.

WARNING: The following procedure utilizes the Windows NT Registry Editor and is intended for the system administrator and the Hitachi Data Systems representative. **Use the Registry Editor with extreme caution.** Always use **regedt32** (not **regedit**). Do not make any changes to the system registry other than those specified below. For instructions on editing the registry, please refer to the online help for the Registry Editor. When specifying multiple parameters, separate each parameter by a semi-colon and a space. If you have questions or concerns, please contact the Hitachi Data Systems Support Center before beginning this procedure.

Verify the disk I/O TOV using the Registry Editor as follows (see Figure 2.3):

1. Start the Windows NT Registry Editor: from the **Start** menu click on **Run** and enter **regedt32**, or double-click on **regedt32** in the Windows NT system directory.
2. Display the disk parameters as follows: go to **HKEY_LOCAL_MACHINE → SYSTEM → CurrentControlSet → Services → Disk** (see Figure 2.3).
3. Make sure that the **TimeOutValue** disk parameter is set to 60 seconds (0x3c).
   - If the **TimeOutValue** is not set to 60 seconds, reboot the NT system, and set the TOV to 60 seconds using one of the HBA setup utilities. If you do not want to reboot, ***carefully*** edit the **TimeOutValue** entry using the NT Registry Editor (see CAUTION above). For instructions on adding or modifying the **TimeOutValue**, refer to the online help for the Registry Editor.
4. Save your changes (if any), and exit the Registry Editor.

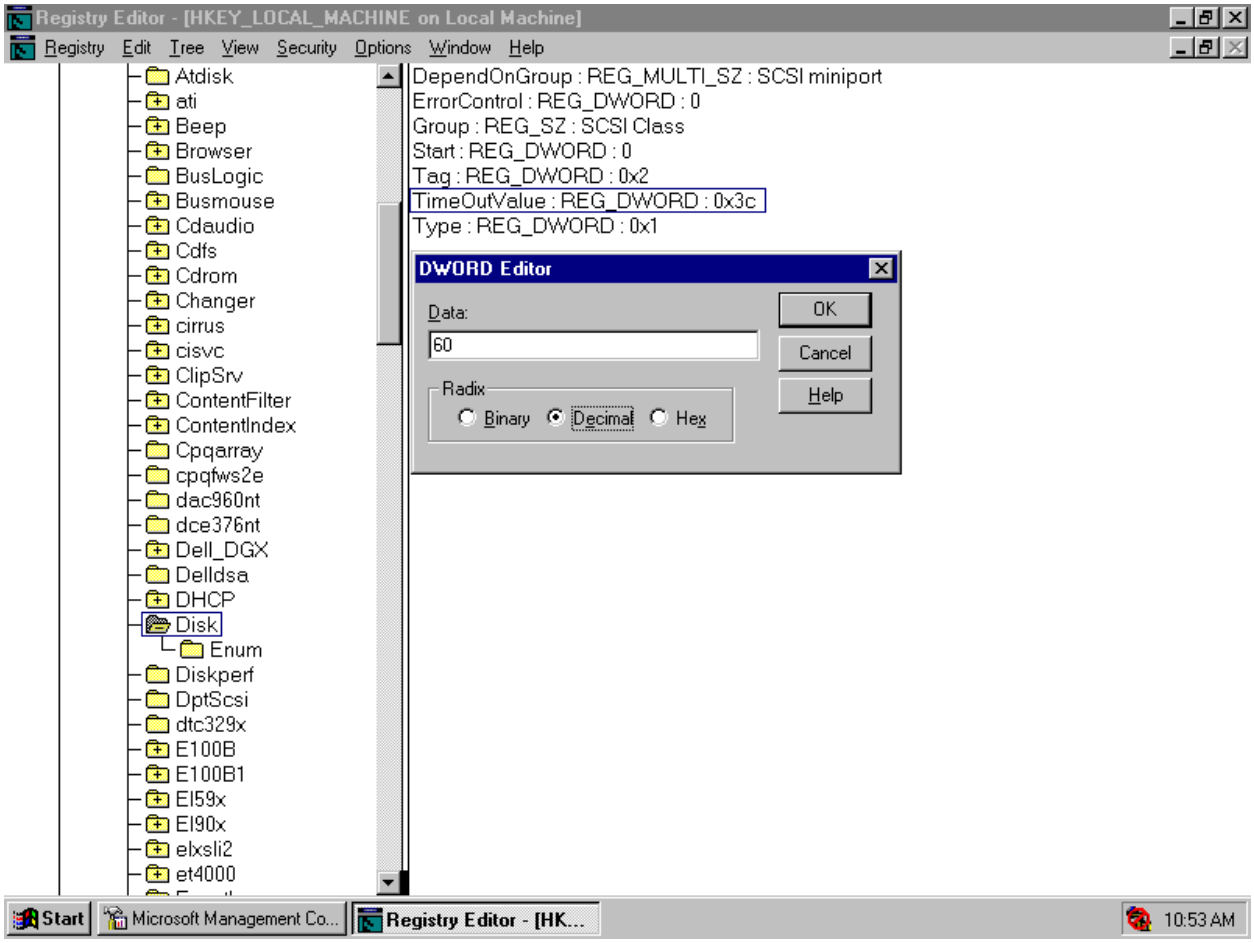

Figure 2.3 Verifying/Setting the Disk I/O TOV Using the Registry Editor

2.6.2 Verifying the Device Parameters

The queue depth parameter for the 9900 devices must be set as specified in Table 2.3 (refer to section 2.5). In addition to queue depth, you should also verify all other required settings for your operational environment (e.g., FC fabric support). You must also make sure that the device parameters are the same for all 9900 devices. This section provides sample instructions for the Emulex and QLogic FC adapters. For other adapters, please refer to the user documentation for the adapter.

CAUTION: The following procedure utilizes the Windows NT Registry Editor and is intended for the system administrator. **Use the Registry Editor with extreme caution.** Always use **regedt32** instead of **regedit**. Do not make any changes to the system registry other than those specified below. For instructions on editing the registry, please refer to the online help for the Registry Editor. When specifying multiple parameters, separate each parameter by a semi-colon and a space. If you have questions or concerns, please contact the Hitachi Data Systems Support Center before beginning this procedure.

Note: If the parameters do not exist, use the Registry Editor to add them.

Verify the queue depth and other device parameters using the Registry Editor as follows:

1. Start the Windows NT Registry Editor.
2. For each Emulex FC adapter (see Figure 2.4):
   e) For the SCSI Mini Port Driver, go to: **HKEY_LOCAL_MACHINE → SYSTEM → CurrentControlSet → Services → lp6nds35 → Parameters → Device**.
   For the SCSI Port Driver, go to: **HKEY_LOCAL_MACHINE → SYSTEM → CurrentControlSet → Services → exsli2**.
   f) Make sure that the **DriverParameter** device parameter has the following values:
   SCSI Mini Port Driver: **QueueDepth=X** (X meets the requirements in Table 2.3)
   SCSI Port Driver: **MaximumQueueDepth=X** (X meets the requirements in Table 2.3)
   g) If the Emulex adapter (SCSI Mini Port Driver) is connected to a fabric switch, verify that the **DriverParameter** device parameter **Topology=1**.
   h) Verify all other required settings for your operational environment. Refer to the user documentation for the adapter as needed.

3. For each QLogic FC adapter:

   i) Display the device parameters for the QLogic FC adapter as follows: go to **HKEY_LOCAL_MACHINE → SYSTEM → CurrentControlSet → Services → ql2100 → Parameters → Device**.

   j) Add the link down timeout parameters to the **DriverParameter**: **LipFFrecovery=1** and **LinkTimeOut=60**. These parameters assist in resolving "hung" loop conditions.

   k) If connected to a fabric switch, add **FabricSupported=1** to the **DriverParameter**.

   l) Verify all other required settings for your operational environment (e.g., support for more than eight LUNs per target ID). Refer to the user documentation for the adapter as needed.

4. If you need to change any adapter settings, reboot the NT system, and use the HBA setup utility. If you do not want to reboot, edit the registry ***carefully*** (see CAUTION above) using the Registry Editor.

5. Save your changes (if any), and exit the Registry Editor.

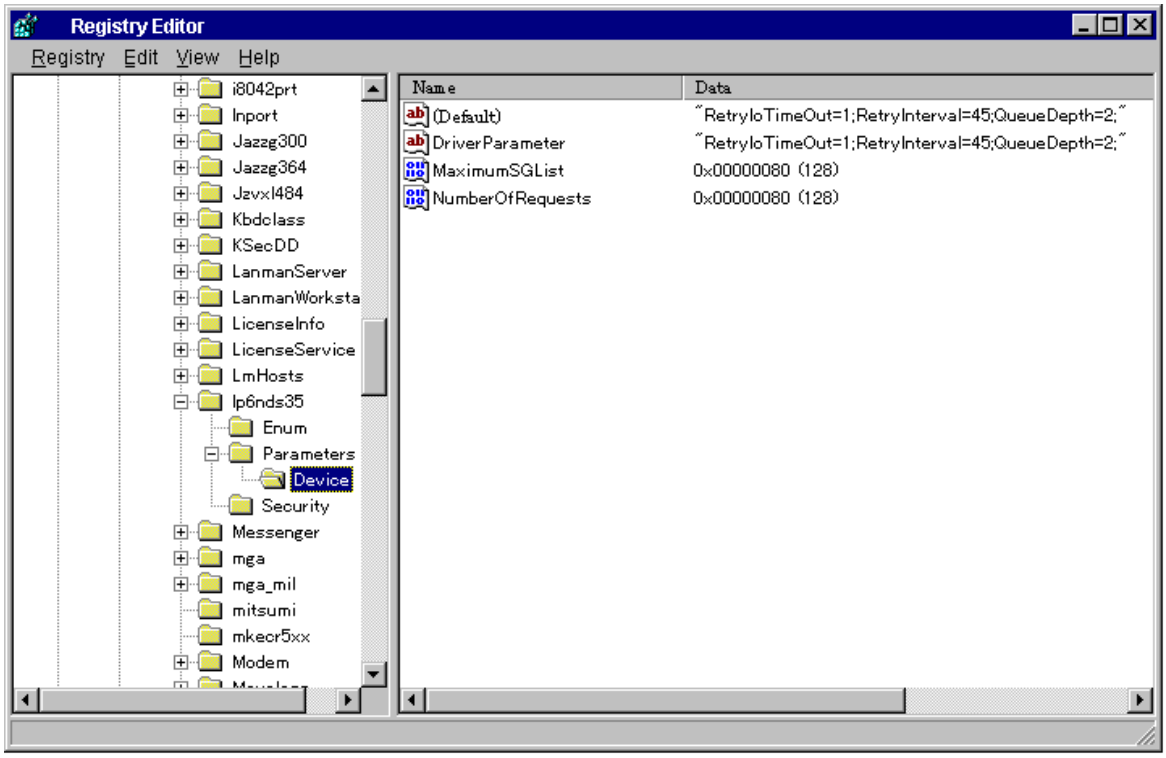

Figure 2.4 Verifying the Queue Depth (Emulex FC adapters shown)

2.7 Verifying New Device Recognition

The last step in preparing for new device configuration is to verify that the NT system recognizes the devices on the newly installed 9900 subsystem. You will display the SCSI device information using the SCSI Adapters control panel to verify that the NT system has recognized the new 9900 devices.

To verify that the NT system recognizes the new devices:

1. Open the Windows NT Control Panel (**Start-Settings-Control Panel**), and double-click on **SCSI Adapters** to open the SCSI Adapters control panel.
2. Select the adapter connected to the 9900, and display the devices connected to the adapter by expanding the view (click on ⊞ next to the adapter name).
3. Verify that all new 9900 devices are displayed (see Figure 2.5). All new devices (SCSI disk and HMDE devices) should be listed. To check the TID and LUN of any device, select the device, click on **Properties**, and then select the **Settings** tab.
4. Add the disk number for each new device to your SCSI Device Information Worksheet (refer to Table 2.5). You will need this information when you write the signatures. To view the disk number of a device, select the device, click on **Properties**

Note: You will use the disk number for each HMDE device in your HMDE volume definition file (**datasetmount.dat**). For example, if disk 3 is a 3390-3B HMDE device, then the entry for this volume in the HMDE volume definition file is:

```
\\.\PHYSICALDRIVE3 volser 3390-3B
```
('volser' = mainframe volume serial number)

5. If you used more than one adapter for the 9900, repeat steps (3) and (4) to verify the new devices on each adapter.

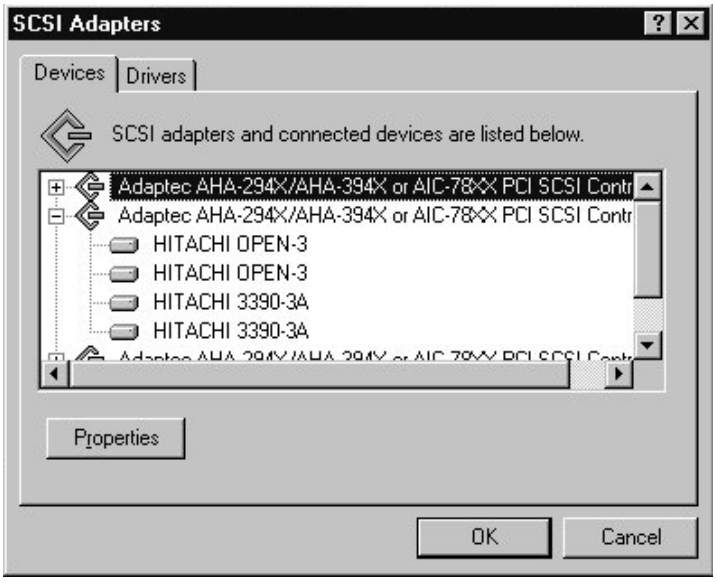

Figure 2.5 Verifying New Device Recognition – SCSI Adapters Control Panel

Table 2.5 SCSI Device Information Worksheet

LDEV (CU:LDEV) (CU = control unit)	Device Type	LUSE (✓)	SCSI Bus Number	Disk 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 installation has been completed and new device recognition has been verified, the new 9900 devices are ready to be configured for use. Configuration of the 9900 devices is performed by the user. The activities involved in configuring the 9900 devices are:

- Writing the Signatures (see section 3.2),
- Creating and Formatting the Partitions (see section 3.3),
- Verifying File System Operations (see section 3.4), and
- Verifying Auto-Mount (see section 3.5).

WARNING: Do not write signatures or create partitions on the HMDE devices, unless the HMDE devices will be used in the Microsoft Cluster Server (MSCS) environment. If so, you must write a signature on each HMDE device, but do not create partitions.

3.1 Writing the Signatures

The first step in configuring the new devices is to write a signature on each device using the NT Disk Administrator. You must write a signature on each SCSI disk device to enable the NT system to vary the device online. For MSCS environments, you must also write a signature on each HMDE device. The 32-bit signature identifies the disk to the NT system. If the disk's TID and/or LUN is changed, or even if the disk is moved to a different controller, the Disk Administrator and Windows NT fault-tolerant driver will continue to recognize it.

To write the signatures on the new disk devices (see Figure 3.1 and Figure 3.2):

1. From the **Start-Programs** menu, select **Administrative Tools (Common)**, and then select **Disk Administrator** to start the Disk Administrator. Initialization takes a few seconds.
2. When the Disk Administrator notifies you that one or more disks have been added, select **OK** to allow the system configuration to be updated. ***Note:*** If you removed any disks, the Disk Administrator will also notify you at this time.
3. The Disk Administrator now displays each new device by disk number and asks if you want to write a signature on the disk (see Figure 3.1). Refer to your completed SCSI Information Worksheet (see Table 2.5) to verify the device type for each disk number. For all SCSI disk devices, select **Yes** to write a signature. For HMDE devices without MSCS, select **No**. For HMDE devices with MSCS, select **Yes** and observe this warning:

   WARNING: After a signature has been written on an HMDE device, there is no way to distinguish the HMDE device from a SCSI disk device. The user must exercise extreme caution not to accidentally partition and format an HMDE device. This will overwrite any data on the HMDE device and also prevent the HMDE software from accessing the device.
4. After you have written a signature (or declined to write a signature) on each new device, the Disk Administrator main panel opens and displays the devices by disk number (see Figure 3.2). The total capacity and free space are displayed for each disk device with a signature. **Configuration information not available** indicates no signature. Do not exit the Disk Administrator yet. You will create partitions on the new SCSI disk devices next.

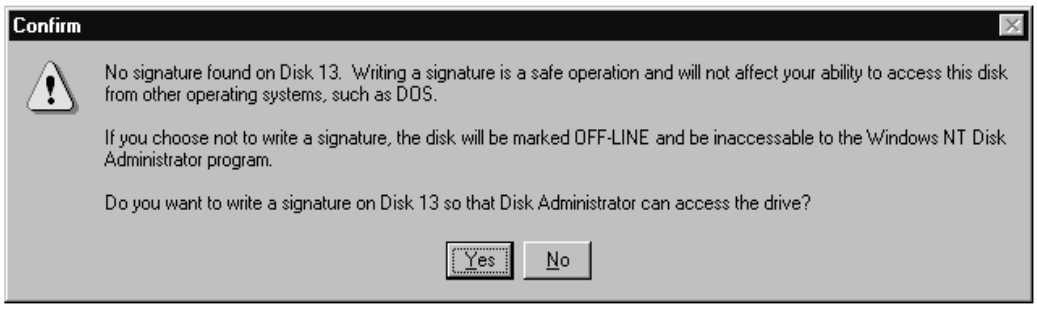

Figure 3.1 Writing the Signatures

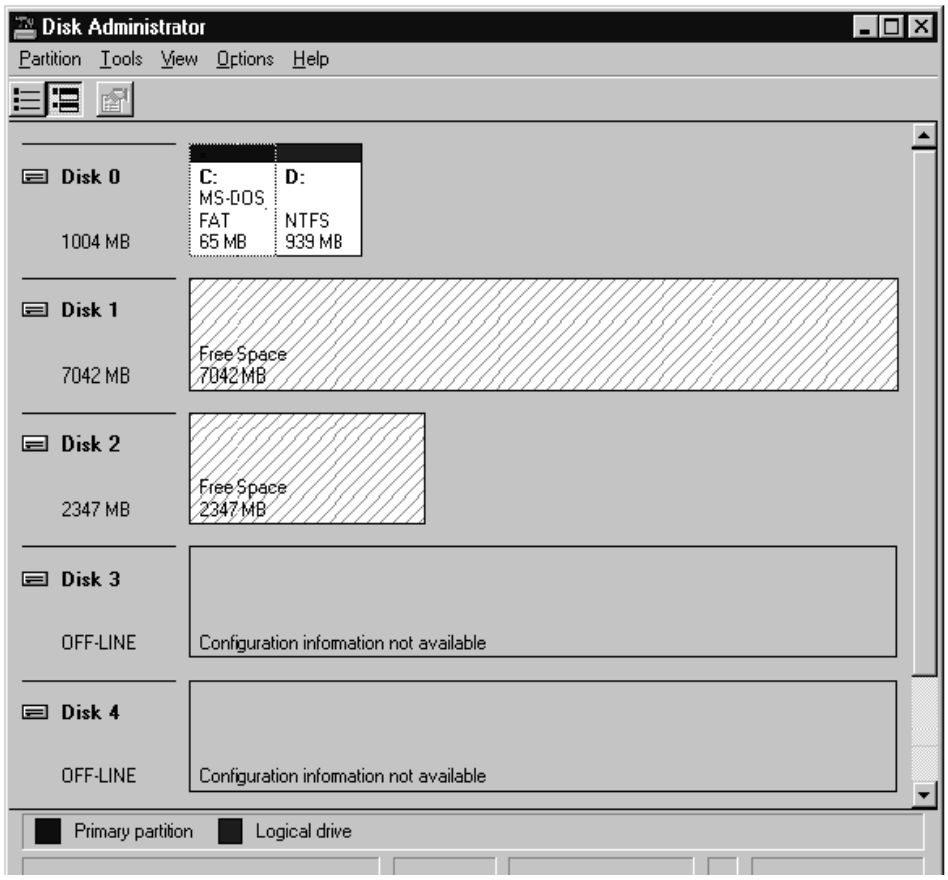

Note: In this example, disk 0 is the local disk disk 1 is an OPEN-9 device, disk 2 is an OPEN-3 device, disk 3 is a 3390-3B device, and disk 4
is a 3390-3A device. The entries in the HMDE volume definition file (**datasetmount.dat**) for these HMDE volumes are:

\\.\PHYSICALDRIVE3 'volser' 3390-3B
\\.\PHYSICALDRIVE4 'volser' 3390-3A ('volser' is the mainframe volume serial number)

Figure 3.2 Disk Administrator Panel Showing New Devices

3.2 Creating and Formatting the Partitions

After you have written the signatures on the new devices, you are ready to create and format the partitions on the new SCSI disk devices. Do not create partitions on the HMDE devices. Use your completed SCSI Device Information Worksheet (refer to Table 2.5) as needed to verify disk numbers and device types

To create and format partitions on the new SCSI disk devices (see Figure 3.3 through Figure 3.9):

1. On the Disk Administrator main panel, select the free space area for the SCSI disk you want to partition, select the **Partition** menu, and then select **Create…** to open the Create Primary Partition panel (see Figure 3.3).
2. On the Create Primary Partition panel, enter the desired partition size (see Figure 3.3), and select **OK**. If the specified partition size is greater than 1024 MB, the Disk Administrator will request confirmation to create the partition.
3. The Disk Administrator panel now shows the new unformatted partition for the selected device. Make sure that the correct partition size is displayed. If the partition size is not correct, repeat steps (1) through (4) to re-enter the correct partition size.
4. Select the **Partition** menu, and select **Commit Changes Now…** (see Figure 3.5). When the confirmation panel appears, select **Yes** to save the changes to your disk configuration.
5. When the disk update confirmation message appears (see Figure 3.6), select **OK**. On the Disk Administrator main panel, verify that the newly created partition changes from **Unformatted** to **Unknown**.
6. On the Disk Administrator main panel, select the newly created partition, select the **Tools** menu, and then select **Format…** (see Figure 3.7) to open the Format panel. The Format panel displays the partition name in its title bar (**G:** in Figure 3.8).
7. Enter the following information on the Format panel (see Figure 3.8):
   - **Capacity: Unknown capacity**. Do not change this entry.
   - **File System**: Select **NTFS** (enables the NT system to write to the disk).
   - **Allocation Unit Size: Default allocation size**. Do not change this entry.
   - **Volume Label**: Enter a volume label, or leave this field blank for no label.
   - **Format Options**: **Select Quick Format** to decrease the time required to format the partition; select **Enable Compression** only if you want to enable compression.
8. Select **Start** to format the partition as specified. When the format warning is displayed (this new format will erase all existing data on disk), select **OK** to continue. The Format panel displays the progress of the format partition operation.
9. When the format complete message is displayed, select **OK**, and then select **Close** to close the Format panel. Verify that the Disk Administrator main panel displays the correct file system (NTFS) for the formatted partition (see Figure 3.9).
10. Repeat steps (1) through (9) for each new SCSI disk device. When you are finished creating and formatting partitions, exit the Disk Administrator (select **Partition-Exit**). When the disk configuration change message comes up, select **Yes** to save your changes. ***Note:*** Make sure to make your new Emergency Repair Disk using RDISK.EXE.

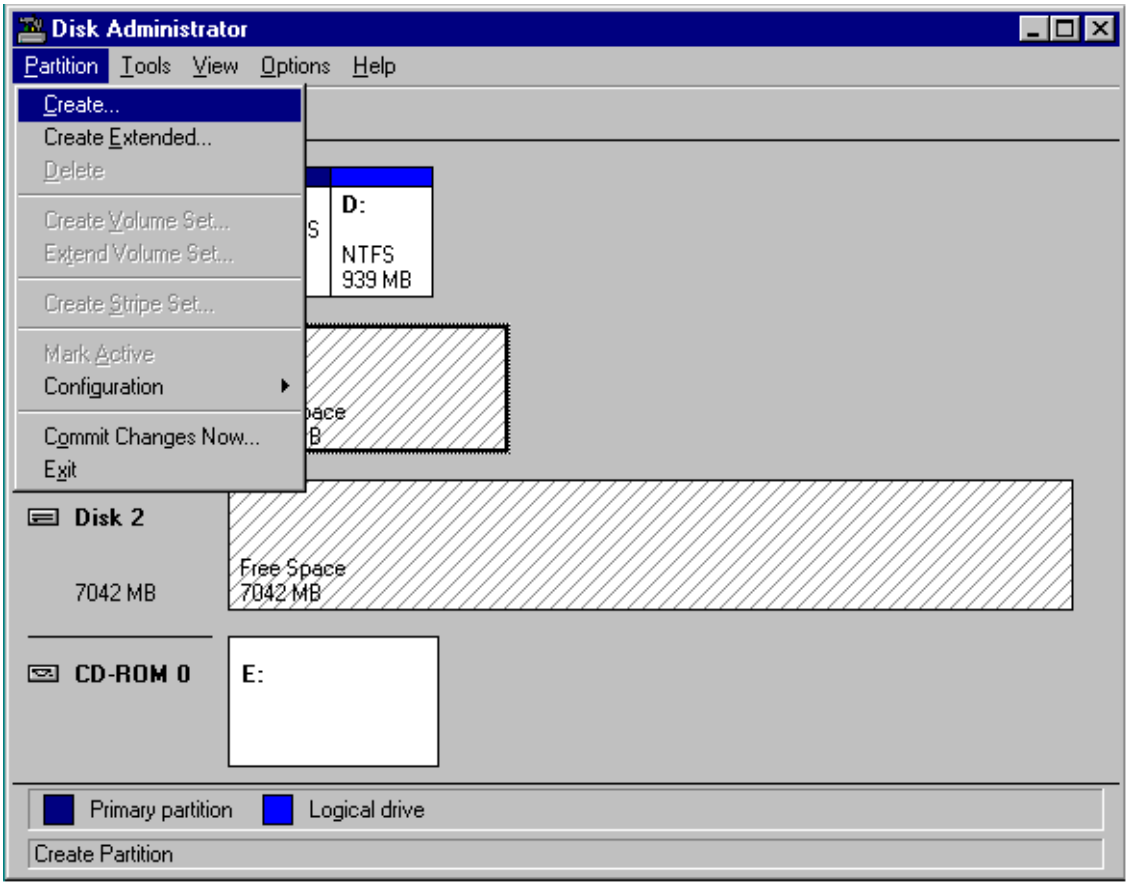

Figure 3.3 Opening the Create Primary Partition Panel

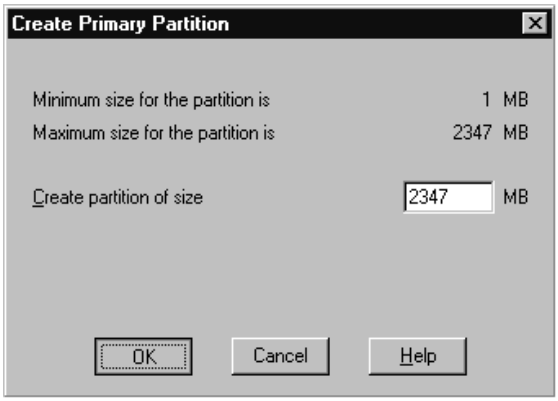

Figure 3.4 Entering the Partition Size

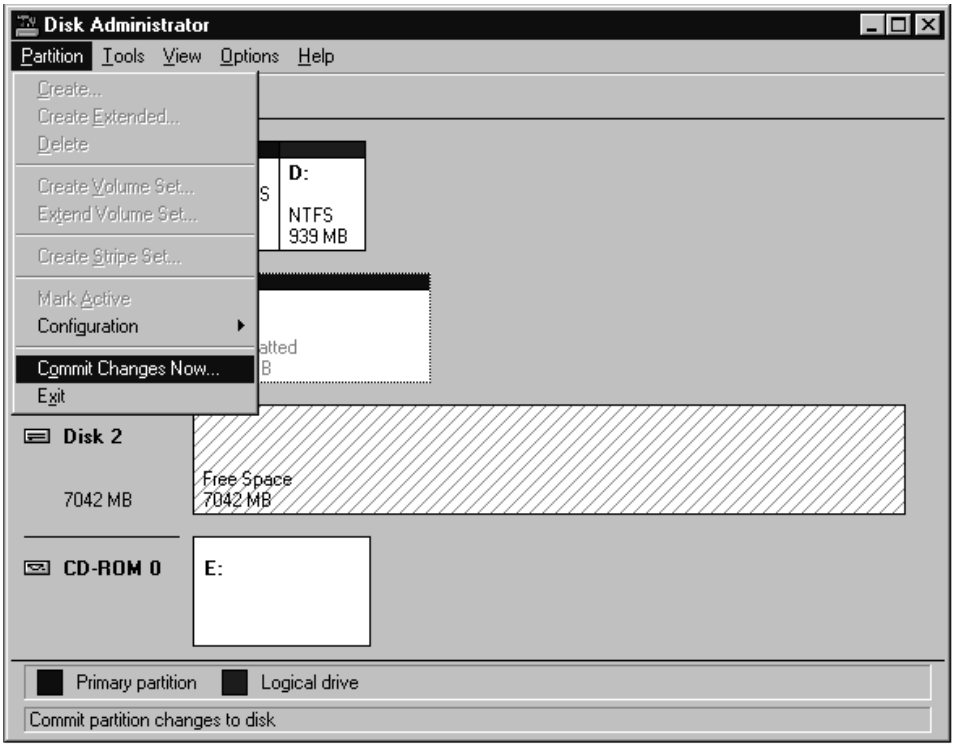

Figure 3.5 Saving the Changes to Disk Configuration

Figure 3.6 Confirming Disk Configuration Update

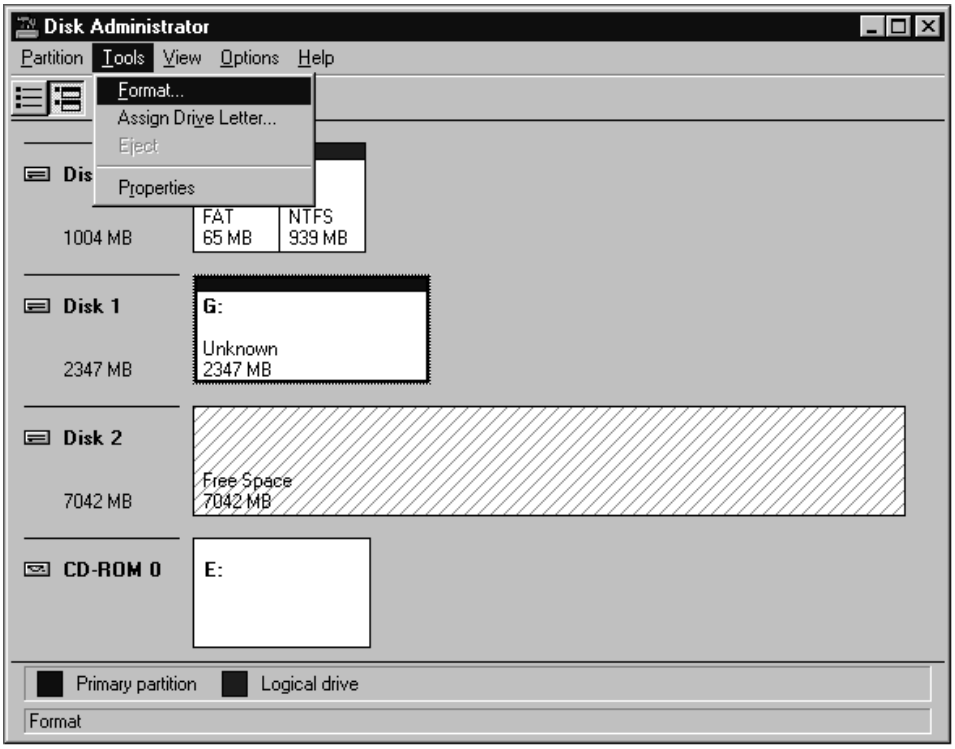

Note: After committing the changes, notice that the newly created partition changes from **Unformatted** to **Unknown**.

Figure 3.7 Opening the Format Panel

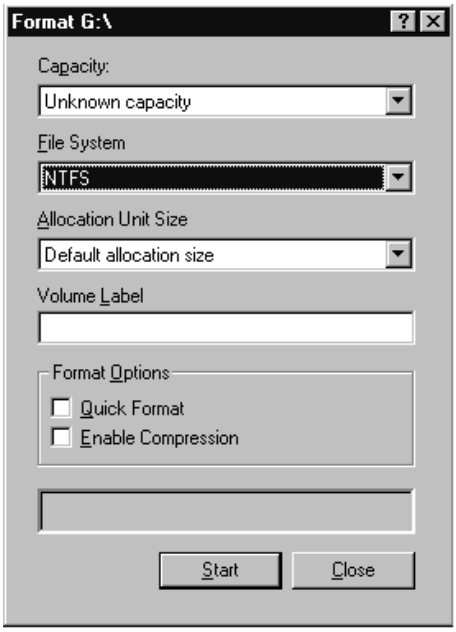

Note: In this example, the name of the partition being formatted is **G:**.

Figure 3.8 Formatting the Partition

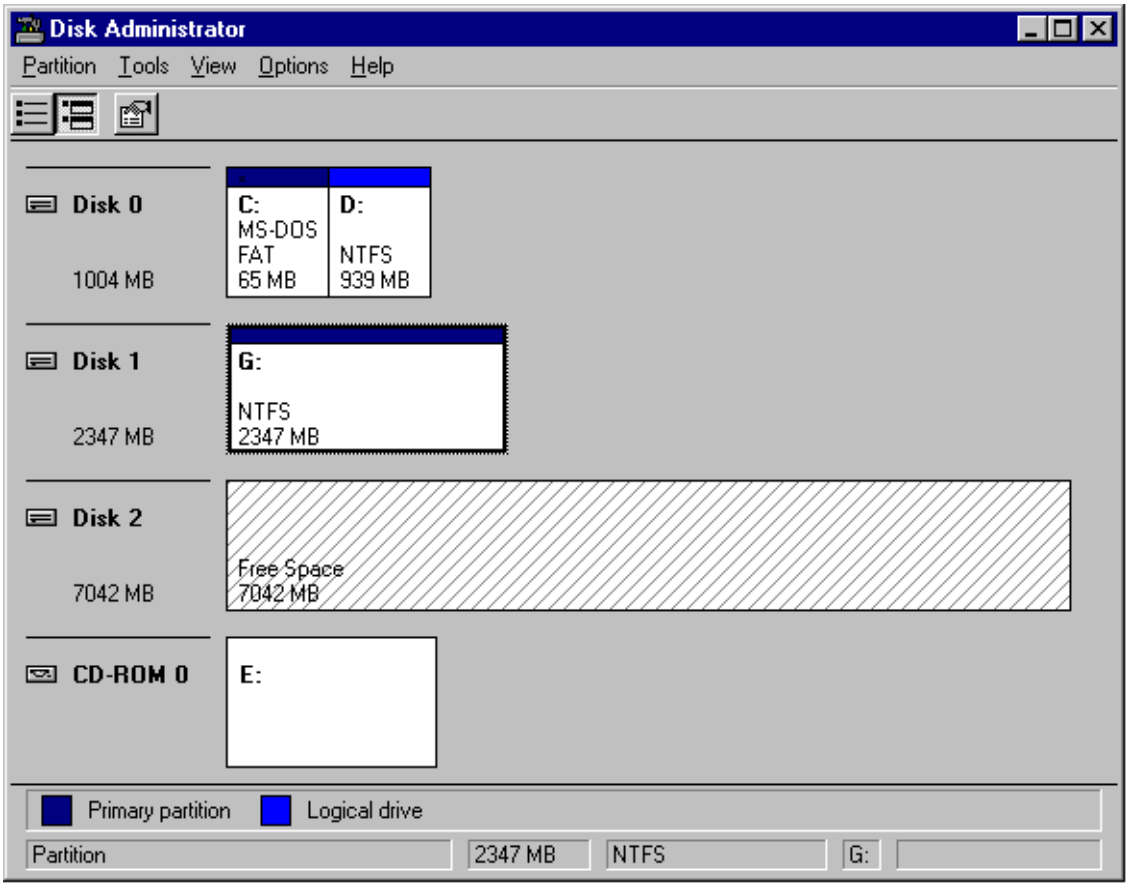

Figure 3.9 Verifying the Formatted Partition

3.3 Verifying File System Operations

After creating and formatting the partitions, you need to verify that the file system is operating properly on each new SCSI disk device (OPEN-x, CVS, and LUSE). The file system enables the NT system to access the devices. You can verify file system operation easily by copying a file onto each new device. If the file is copied successfully, this verifies that the file system is operating properly (i.e., the NT system can access the new device).

Note: Do not perform this procedure for HMDE devices. You must use the HMDE File Conversion Utility (FCU) or File Access Library (FAL) to access the HMDE devices.

To verify file system operations for the new SCSI disk devices:

From the NT desktop, double-click on **My Computer** to display all connected devices. All newly partitioned disks should appear in this window (see Figure 3.10).

1. Select the device you want to verify, and then display its Properties (select the **File** menu and then select **Properties**, or right-mouse-click and then select **Properties**).
2. On the Properties panel (see Figure 3.11), verify that the properties are correct: label (optional), type, capacity, and file system.
3. Copy a file to the new device. Any file will do, so choose a small one to speed things up.
4. Display the contents of the new device to make sure that the copy operation completed successfully (see Figure 3.12). The copied file should be displayed with the correct file size. If desired, compare the copied file with the original file to verify no differences.

5. Delete the copied file from the new device, and verify the file was deleted successfully.
6. Repeat steps (2) through (6) for each new SCSI disk device.

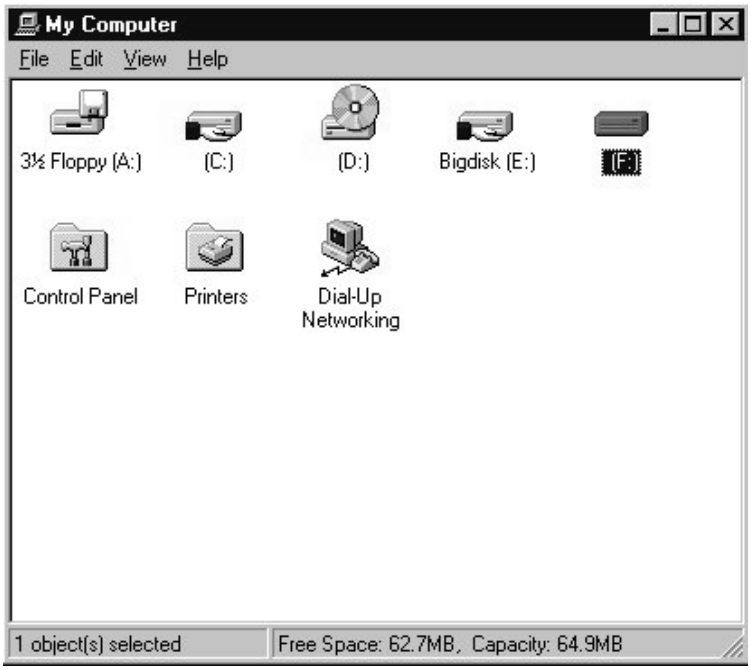

Note: In this example, [F:] is the only new device.

Figure 3.11 Displaying the Connected Devices

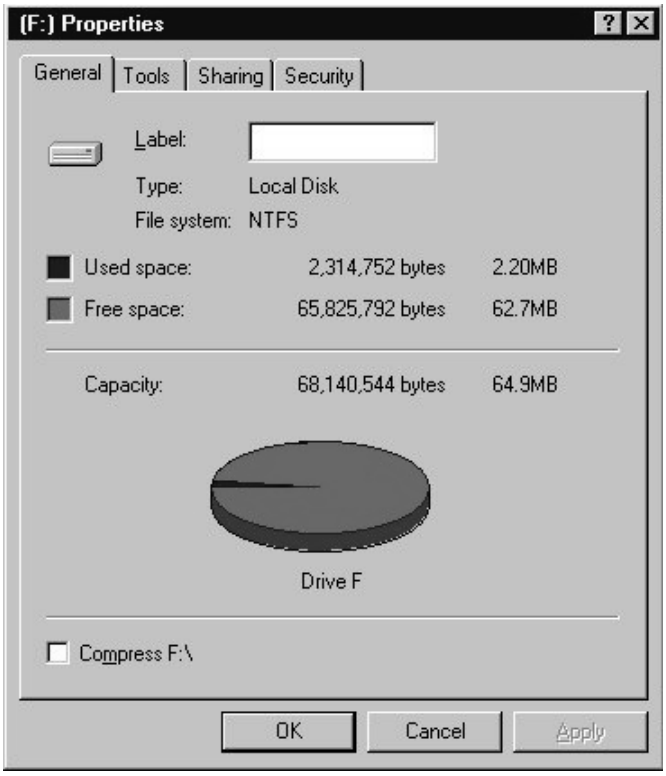

Figure 3.12 Verifying the New Device Properties

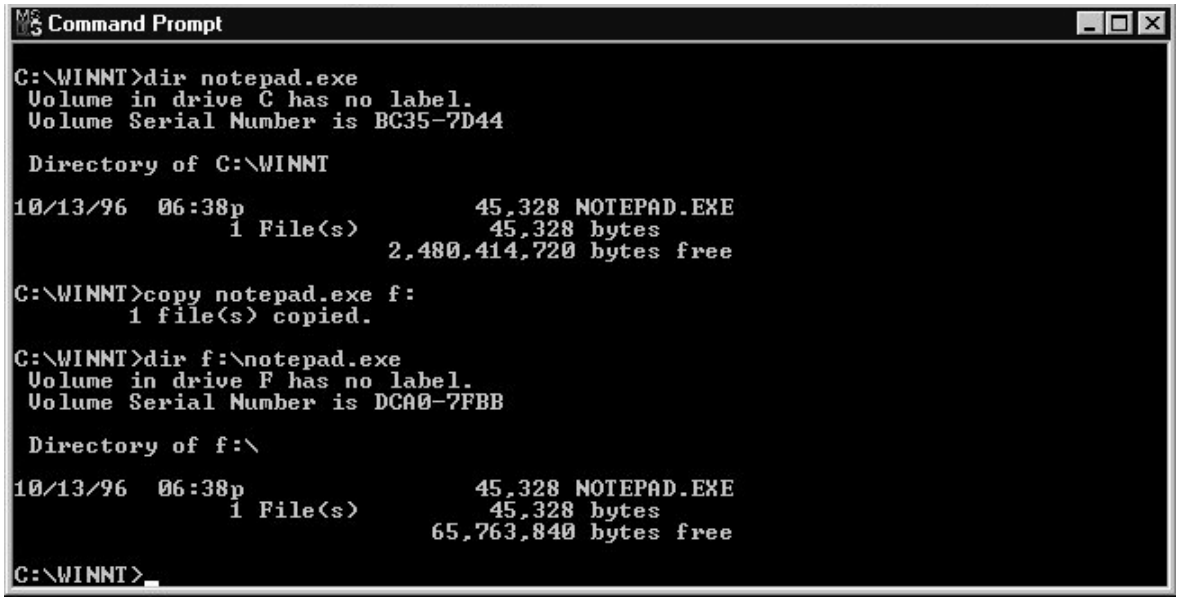

Figure 3.13 Verifying the File Copy Operation

3.4 Verifying Auto-Mount

The last step in configuring the new devices is to verify that all new devices are automatically mounted at system boot-up. To verify auto-mount of the new devices:

1. Shut down and then restart the Windows NT system.
2. Open **My Computer**, and verify that all new SCSI disk devices are displayed.
3. Verify that the NT system can access each new device by repeating the procedure in the previous section:
   a) Verify the device properties for all new devices (refer to Figure 3.11).
   b) Copy a file to each new device to make sure that the devices are functioning properly (refer to Figure 3.12).

3.5 Verifying Client Access

The last step in new device configuration is to verify that the NetWare® clients can access the new volumes. To verify access:

1. Copy an existing file onto each new volume.

   Note: This document does not provide instructions for copying a file to a mounted volume.
2. Verify that the file was copied successfully. If not, see Chapter 5 for troubleshooting instructions.

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 the Windows NT® operating system, the 9900 supports the following middleware products:

- Microsoft® Cluster Server (MSCS) for host fail-over (see section 4.1), and
- Hitachi Path Manager for I/O path failover (see section 4.2).

  Note: The logical volume management functions are included in the Windows NT® operating system (e.g., Disk Administrator).

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

Note: The user is responsible for configuring the middleware and SNMP management software on the UNIX®/PC server host. For assistance with host middleware and/or SNMP configuration, please refer to the user documentation, or contact the vendor's technical support.

4.1 Host Fail-Over

The 9900 subsystem supports the Microsoft® Cluster Server (MSCS) host fail-over feature of the Windows NT® OS. Please contact Microsoft® for the latest information on MSCS.

When the 9900 devices will be operating in an MSCS environment, you must perform the following additional configuration activity: writing **signatures** (see section 3.1). For MSCS operations, allow the NT Disk Administrator to write a signature on each HMDE device (e.g., 3390-3A/B/C, 3380-KA/B/C, OPEN-3 for HMDEoto).

After 9900 device configuration is complete, make sure to configure the MSCS software as needed to recognize the devices on the newly attached 9900 subsystem(s). For assistance with MSCS operations, please refer to the Microsoft® user documentation or contact Microsoft® customer support.

4.2 Alternate I/O Path

The 9900 subsystem supports the Hitachi Path Manager alternate I/O path middleware product for the Windows NT® OS. After you have completed 9900 device configuration as described in Chapter 3, make sure to configure Hitachi Path Manager as needed to recognize the devices on the newly attached 9900 subsystem(s). For assistance with Hitachi Path Manager operations, please refer to the *Hitachi Path Manager User's Guide* (MK-90RD018), or contact the Hitachi Data Systems Support Center (see section 5.2).

4.3 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 UNIX®/PC server 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 UNIX®/PC server host. Notification of 9900 error conditions is made in real time, providing UNIX® and PC server users with the same level of monitoring and support available to S/390® mainframe users. The SIM reporting via SNMP enables the user to monitor the 9900 subsystem from the UNIX®/PC server host without having to check the Remote Console PC for remote SIMs (R-SIMs).

When a SIM occurs, the 9900 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.

Note: The user is responsible for configuring the SNMP manager on the UNIX®/PC server host. For assistance with SNMP manager configuration on the UNIX®/PC server host, please refer to the user documentation, or contact the vendor's technical support.

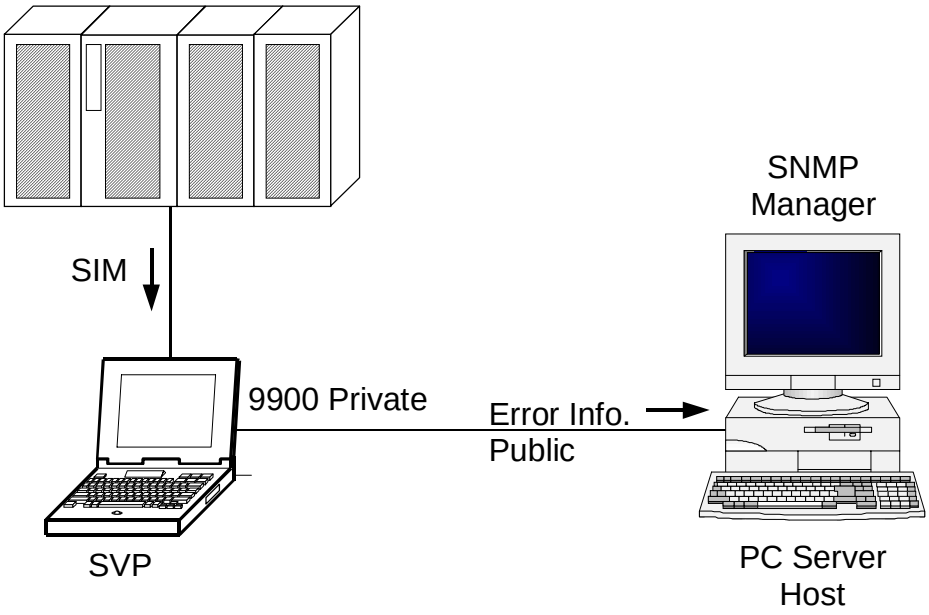

Figure 4.1 9900 SNMP Environment

Chapter 5 Troubleshooting

5.1 Troubleshooting

The Hitachi Lightning 9900™ array subsystem provides continuous data availability. For troubleshooting information for the 9900 subsystem, please refer to the *Hitachi Lightning 9900™ User and Reference Guide* (MK-90RD008).

Table 5.1 lists potential error conditions during 9900 device 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. See section 5.2 for instructions on calling the Hitachi Data Systems Support Center.

Table 5.1 Troubleshooting

Error Condition	Recommended Action
The devices are not recognized by the system.	Make sure that the READY indicator lights on the 9900 subsystem are ON. Make sure that the fibre cables are correctly installed and firmly connected. For some adapters, the LUNs must start at 0 and continue sequentially without skipping any numbers. Make sure that the fibre channel ports, LUN security, and/or switch zoning is defined properly.
The NT system does not reboot properly after hard shutdown.	If the NT system is powered off unexpectedly (without the normal shutdown process), wait three minutes before restarting the NT system. This allows the 9900's internal time-out process to purge all queued commands so that the 9900 is available (not busy) during system startup. If the NT system is restarted too soon, the 9900 will continue trying to process the queued commands, and the NT system will not reboot successfully.

5.2 Calling the Support Center

If you need to call the Hitachi Data Systems Support Center, make sure to provide as much information about the problem as possible, including the circumstances surrounding the error or failure and the exact content of any error messages displayed on the host system(s). Please check the remote service information messages (R-SIMs) logged on the Remote Console PC, and note the reference codes and severity levels of the recent R-SIMs.

The worldwide Hitachi Data Systems 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
blk	block
CVS	Custom Volume Size
ESCON	Enterprise System Connection (IBM trademark for optical channels)
FAL	File Access Library (HMDE software component)
FC	fibre channel
FCP	fibre-channel protocol
FCU	File Conversion Utility (HMDE software component)
HBA	host bus adapter
HDS	Hitachi Data Systems
HMDE	Hitachi Multiplatform Data Exchange
HMDEmto	HMDE mainframe-to-open
HMDEotm	HMDE open-to-mainframe
OPEN-x-HMDEoto	HMDE open-to-open
HOMRCF	Hitachi Open Multiple RAID Coupling Feature (also called ShadowImage)
HORC	Hitachi Open Remote Copy
I/O	input/output
IBM	International Business Machines Corporation
LDEV	logical device
LU	logical unit
LUN	logical unit number, logical unit
LUSE	LU Size Expansion
MSCS	Microsoft Cluster Server
NTFS	NT file system
OFC	open fibre control
OS	operating system
PC	personal computer system
PCI	power control interface
R-SIM	remote service information message
SCSI	small computer system interface
SIM	service information message
SNMP	simple network management protocol
SSB	sense byte
SVP	service processor
TID	target ID
TOV	timeout value

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, Windows NT "maps" a TID to each AL-PA.

Table B.1 and B.2 identify the fixed mappings between the bus/TID/LUN addresses assigned by Windows NT and the FC native addresses (AL_PA/SEL_ID) for FC adapters. There are two potential mappings depending on the value of the ScanDown registry parameter:

- For ScanDown = 0 (default) see Table B.1.
- For ScanDown = 1 see Table B.2.

Note: When 9900 devices and other types of devices are connected in the same arbitrated loop, the mappings defined in Tables B.1 and B.2 cannot be guaranteed.

Note: The Emulex driver emulates six SCSI busses per adapter to map all 126 possible AL-PAs to target IDs. The first bus (bus 0) is a dummy bus.

Table B.1 SCSI TID Map for Emulex FC Adapter (ScanDown=0)

Bus #	TID	LUN	AL_PA	SEL_ID
0	0-31	0-7	NONE	NONE
1	0	0-7	0x01	0x7D
	1	0-7	0x02	0x7C
	2	0-7	0x04	0x7B
	3	0-7	0x08	0x7A
	4	0-7	0x0F	0x79
	5	0-7	0x10	0x78
	6	0-7	0x17	0x77
	7	0-7	0x18	0x76
	8	0-7	0x1B	0x75
	9	0-7	0x1D	0x74
	10	0-7	0x1E	0x73
	11	0-7	0x1F	0x72
	12	0-7	0x23	0x71
	13	0-7	0x25	0x70
	14	0-7	0x26	0x6F
	15	0-7	0x27	0x6E
	16	0-7	0x29	0x6D
	17	0-7	0x2A	0x6C
	18	0-7	0x2B	0x6B
	19	0-7	0x2C	0x6A
	20	0-7	0x2D	0x69
	21	0-7	0x2E	0x68
	22	0-7	0x31	0x67
	23	0-7	0x32	0x66
	24	0-7	0x33	0x65
	25	0-7	0x34	0x64
	26	0-7	0x35	0x63
	27	0-7	0x36	0x62
	28	0-7	0x39	0x61
	29	0-7	0x3A	0x60
	30	0-7	0x3C	0x5F
	31	0-7	NONE	NONE

Bus #	TID	LUN	AL_PA	SEL_ID
2	0	0-7	0x43	0x5E
	1	0-7	0x45	0x5D
	2	0-7	0x46	0x5C
	3	0-7	0x47	0x5B
	4	0-7	0x49	0x5A
	5	0-7	0x4A	0x59
	6	0-7	0x4B	0x58
	7	0-7	0x4C	0x57
	8	0-7	0x4D	0x56
	9	0-7	0x4E	0x55
	10	0-7	0x51	0x54
	11	0-7	0x52	0x53
	12	0-7	0x53	0x52
	13	0-7	0x54	0x51
	14	0-7	0x55	0x50
	15	0-7	0x56	0x4F
	16	0-7	0x59	0x4E
	17	0-7	0x5A	0x4D
	18	0-7	0x5C	0x4C
	19	0-7	0x63	0x4B
	20	0-7	0x65	0x4A
	21	0-7	0x66	0x49
	22	0-7	0x67	0x48
	23	0-7	0x69	0x47
	24	0-7	0x6A	0x46
	25	0-7	0x6B	0x45
	26	0-7	0x6C	0x44
	27	0-7	0x6D	0x43
	28	0-7	0x6E	0x42
	29	0-7	0x71	0x41
	30	0-7	0x72	0x40
	31	0-7	NONE	NONE

Bus #	TID	LUN	AL_PA	SEL_ID
3	0	0-7	0x73	0x3F
	1	0-7	0x74	0x3E
	2	0-7	0x75	0x3D
	3	0-7	0x76	0x3C
	4	0-7	0x79	0x3B
	5	0-7	0x7A	0x3A
	6	0-7	0x7C	0x39
	7	0-7	0x80	0x38
	8	0-7	0x81	0x37
	9	0-7	0x82	0x36
	10	0-7	0x84	0x35
	11	0-7	0x88	0x34
	12	0-7	0x8F	0x33
	13	0-7	0x90	0x32
	14	0-7	0x97	0x31
	15	0-7	0x98	0x30
	16	0-7	0x9B	0x2F
	17	0-7	0x9D	0x2E
	18	0-7	0x9E	0x2D
	19	0-7	0x9F	0x2C
	20	0-7	0xA3	0x2B
	21	0-7	0xA5	0x2A
	22	0-7	0xA6	0x29
	23	0-7	0xA7	0x28
	24	0-7	0xA9	0x27
	25	0-7	0xAA	0x26
	26	0-7	0xAB	0x25
	27	0-7	0xAC	0x24
	28	0-7	0xAD	0x23
	29	0-7	0xAE	0x22
	30	0-7	0xB1	0x21
	31	0-7	NONE	NONE

Table B.1 SCSI TID Map for Emulex FC Adapter (ScanDown=0) (continued)

Bus #	TID	LUN	AL_PA	SEL_ID
4	0	0-7	0xB2	0x20
	1	0-7	0xB3	0x1F
	2	0-7	0xB4	0x1E
	3	0-7	0xB5	0x1D
	4	0-7	0xB6	0x1C
	5	0-7	0xB9	0x1B
	6	0-7	0xBA	0x1A
	7	0-7	0xBC	0x19
	8	0-7	0xC3	0x18
	9	0-7	0xC5	0x17
	10	0-7	0xC6	0x16
	11	0-7	0xC7	0x15
	12	0-7	0xC9	0x14
	13	0-7	0xCA	0x13
	14	0-7	0xCB	0x12
	15	0-7	0xCC	0x11
	16	0-7	0xCD	0x10
	17	0-7	0xCE	0x0F
	18	0-7	0xD1	0x0E
	19	0-7	0xD2	0x0D
	20	0-7	0xD3	0x0C
	21	0-7	0xD4	0x0B
	22	0-7	0xD5	0x0A
	23	0-7	0xD6	0x09
	24	0-7	0xD9	0x08
	25	0-7	0xDA	0x07
	26	0-7	0xDC	0x06
	27	0-7	0xE0	0x05
	28	0-7	0xE1	0x04
	29	0-7	0xE2	0x03
	30	0-7	0xE4	0x02
	31	0-7	NONE	NONE

Bus #	TID	LUN	AL_PA	SEL_ID
5	0	0-7	0xE8	0x01
	1	0-7	0xEF	0x00
	2	0-7	NONE	NONE
	3	0-7	NONE	NONE
	4	0-7	NONE	NONE
	5	0-7	NONE	NONE
	6	0-7	NONE	NONE
	7	0-7	NONE	NONE
	8	0-7	NONE	NONE
	9	0-7	NONE	NONE
	10	0-7	NONE	NONE
	11	0-7	NONE	NONE
	12	0-7	NONE	NONE
	13	0-7	NONE	NONE
	14	0-7	NONE	NONE
	15	0-7	NONE	NONE
	16	0-7	NONE	NONE
	17	0-7	NONE	NONE
	18	0-7	NONE	NONE
	19	0-7	NONE	NONE
	20	0-7	NONE	NONE
	21	0-7	NONE	NONE
	22	0-7	NONE	NONE
	23	0-7	NONE	NONE
	24	0-7	NONE	NONE
	25	0-7	NONE	NONE
	26	0-7	NONE	NONE
	27	0-7	NONE	NONE
	28	0-7	NONE	NONE
	29	0-7	NONE	NONE
	30	0-7	NONE	NONE
	31	0-7	NONE	NONE

Table B.2 SCSI TID Map for Emulex FC Adapter (ScanDown=1)

Bus #	TID	LUN	AL_PA	SEL_ID
0	0-31	0-7	NONE	NONE
1	0	0-7	0xEF	0x00
	1	0-7	0xE8	0x01
	2	0-7	0xE4	0x02
	3	0-7	0xE2	0x03
	4	0-7	0xE1	0x04
	5	0-7	0xE0	0x05
	6	0-7	0xDC	0x06
	7	0-7	0xDA	0x07
	8	0-7	0xD9	0x08
	9	0-7	0xD6	0x09
	10	0-7	0xD5	0x0A
	11	0-7	0xD4	0x0B
	12	0-7	0xD3	0x0C
	13	0-7	0xD2	0x0D
	14	0-7	0xD1	0x0E
	15	0-7	0xCE	0x0F
	16	0-7	0xCD	0x10
	17	0-7	0xCC	0x11
	18	0-7	0xCB	0x12
	19	0-7	0xCA	0x13
	20	0-7	0xC9	0x14
	21	0-7	0xC7	0x15
	22	0-7	0xC6	0x16
	23	0-7	0xC5	0x17
	24	0-7	0xC3	0x18
	25	0-7	0xBC	0x19
	26	0-7	0xBA	0x1A
	27	0-7	0xB9	0x1B
	28	0-7	0xB6	0x1C
	29	0-7	0xB5	0x1D
	30	0-7	0xB4	0x1E
	31	0-7	NONE	NONE

Bus #	TID	LUN	AL_PA	SEL_ID
2	0	0-7	0xB3	0x1F
	1	0-7	0xB2	0x20
	2	0-7	0xB1	0x21
	3	0-7	0xAE	0x22
	4	0-7	0xAD	0x23
	5	0-7	0xAC	0x24
	6	0-7	0xAB	0x25
	7	0-7	0xAA	0x26
	8	0-7	0xA9	0x27
	9	0-7	0xA7	0x28
	10	0-7	0xA6	0x29
	11	0-7	0xA5	0x2A
	12	0-7	0xA3	0x2B
	13	0-7	0x9F	0x2C
	14	0-7	0x9E	0x2D
	15	0-7	0x9D	0x2E
	16	0-7	0x9B	0x2F
	17	0-7	0x98	0x30
	18	0-7	0x97	0x31
	19	0-7	0x90	0x32
	20	0-7	0x8F	0x33
	21	0-7	0x88	0x34
	22	0-7	0x84	0x35
	23	0-7	0x82	0x36
	24	0-7	0x81	0x37
	25	0-7	0x80	0x38
	26	0-7	0x7C	0x39
	27	0-7	0x7A	0x3A
	28	0-7	0x79	0x3B
	29	0-7	0x76	0x3C
	30	0-7	0x75	0x3D
	31	0-7	NONE	NONE

Bus #	TID	LUN	AL_PA	SEL_ID
3	0	0-7	0x74	0x3E
	1	0-7	0x73	0x3F
	2	0-7	0x72	0x40
	3	0-7	0x71	0x41
	4	0-7	0x6E	0x42
	5	0-7	0x6D	0x43
	6	0-7	0x6C	0x44
	7	0-7	0x6B	0x45
	8	0-7	0x6A	0x46
	9	0-7	0x69	0x47
	10	0-7	0x67	0x48
	11	0-7	0x66	0x49
	12	0-7	0x65	0x4A
	13	0-7	0x63	0x4B
	14	0-7	0x5C	0x4C
	15	0-7	0x5A	0x4D
	16	0-7	0x59	0x4E
	17	0-7	0x56	0x4F
	18	0-7	0x55	0x50
	19	0-7	0x54	0x51
	20	0-7	0x53	0x52
	21	0-7	0x52	0x53
	22	0-7	0x51	0x54
	23	0-7	0x4E	0x55
	24	0-7	0x4D	0x56
	25	0-7	0x4C	0x57
	26	0-7	0x4B	0x58
	27	0-7	0x4A	0x59
	28	0-7	0x49	0x5A
	29	0-7	0x47	0x5B
	30	0-7	0x46	0x5C
	31	0-7	NONE	NONE

Table B.2 SCSI TID Map for Emulex FC Adapter (ScanDown=1) (continued)

Bus #	TID	LUN	AL_PA	SEL_ID
4	0	0-7	0x45	0x5D
	1	0-7	0x43	0x5E
	2	0-7	0x3C	0x5F
	3	0-7	0x3A	0x60
	4	0-7	0x39	0x61
	5	0-7	0x36	0x62
	6	0-7	0x35	0x63
	7	0-7	0x34	0x64
	8	0-7	0x33	0x65
	9	0-7	0x32	0x66
	10	0-7	0x31	0x67
	11	0-7	0x2E	0x68
	12	0-7	0x2D	0x69
	13	0-7	0x2C	0x6A
	14	0-7	0x2B	0x6B
	15	0-7	0x2A	0x6C
	16	0-7	0x29	0x6D
	17	0-7	0x27	0x6E
	18	0-7	0x26	0x6F
	19	0-7	0x25	0x70
	20	0-7	0x23	0x71
	21	0-7	0x1F	0x72
	22	0-7	0x1E	0x73
	23	0-7	0x1D	0x74
	24	0-7	0x1B	0x75
	25	0-7	0x18	0x76
	26	0-7	0x17	0x77
	27	0-7	0x10	0x78
	28	0-7	0x0F	0x79
	29	0-7	0x08	0x7A
	30	0-7	0x04	0x7B
	31	0-7	NONE	NONE

Bus #	TID	LUN	AL_PA	SEL_ID
5	0	0-7	0x02	0x7C
	1	0-7	0x01	0x7D
	2	0-7	NONE	NONE
	3	0-7	NONE	NONE
	4	0-7	NONE	NONE
	5	0-7	NONE	NONE
	6	0-7	NONE	NONE
	7	0-7	NONE	NONE
	8	0-7	NONE	NONE
	9	0-7	NONE	NONE
	10	0-7	NONE	NONE
	11	0-7	NONE	NONE
	12	0-7	NONE	NONE
	13	0-7	NONE	NONE
	14	0-7	NONE	NONE
	15	0-7	NONE	NONE
	16	0-7	NONE	NONE
	17	0-7	NONE	NONE
	18	0-7	NONE	NONE
	19	0-7	NONE	NONE
	20	0-7	NONE	NONE
	21	0-7	NONE	NONE
	22	0-7	NONE	NONE
	23	0-7	NONE	NONE
	24	0-7	NONE	NONE
	25	0-7	NONE	NONE
	26	0-7	NONE	NONE
	27	0-7	NONE	NONE
	28	0-7	NONE	NONE
	29	0-7	NONE	NONE
	30	0-7	NONE	NONE
	31	0-7	NONE	NONE

Appendix C Configuring the Boot Disk on the 9900

You can configure the Windows NT® boot disk on the 9900 subsystem under the following conditions:

- **OS version**: The Windows NT® OS version must be 4.0.
  Firmware version: The Windows NT® 9000 firmware version must be 3.00 (full Flash firmware image) or later and include the following: v1.20 Boot BIOS or later; v1.30 POST or later; and v3.00 Operational F/W SLI-1 and SL-2 or later
- The **host bus adapter** must be Emulex LP7000E/LP8000

To set up the Windows NT® boot disk on the 9900 subsystem:

1. Load and install the appropriate firmware. If your system does not have the necessary firmware, download and install it from http://www.emulex.com.

For further assistance, please contact your Hitachi Data Systems representative.