ptx®/HITACHI-RAIDMP V2.3.0
Licensed Program Specification

45P0322

Printed in the United States of America.

Contents

Chapter 1 Introduction

About This Manual 1-1
Document Organization 1-1
Audience 1-2
How to Use This Document 1-2
Related Documentation 1-2
Hitachi RAID Subsystem Support Overview 1-3
DF350F Support Summary 1-4
RAID200 Support Summary 1-5
RAID300 Support Summary 1-6
RAID400 Support Summary 1-8
DF500 Support Summary 1-9
FC Switch Support and Operation 1-10
FC Switches Supported 1-10
FC Switch Must Be Set to 1000 Compatibility Mode 1-10
New or Changed in This Release 1-11
Problems Fixed in This Release 1-11
Known Problems in This Release 1-12
ptx/HITACHI-RAIDMP Issues 1-12
PR#251951 - PANIC: cfg_coalesce: found recursing coalesce scenario 1-12
PR#251854 - "sar: corrupted disk response and active times" w/ HITACHI disks 1-12
PR#247568 - S/W Management Package Consistency Check Fails With Files Not Found 1-13
RAID400 Issues 1-13
FC Switch Ports May Log In Incorrectly 1-13
DF500 Issues 1-13
FC Switch Ports May Log In Incorrectly 1-13
RAID200/RAID300 Issues 1-13
SCSI Connections—Path Failure May Require SCSI Bus Reset 1-13
FCB SCSI Connections—FCB Microcode May Run Out of Memory at Boot 1-14
RAID300 Issues 1-15
Issues Relating to Microcode Downloads 1-15
Fibre Connections — Changing LDEV Configurations Causes Lost Connections 1-16
RAID200 Issues 1-17
Sequential I/O May Be Less Than Expected 1-17
DF350F Issues 1-17
PR #236699—Disks May Go Offline Under Heavy Load in Multi-Initiator Environments 1-17

Chapter 2 ptx/HITACHI-RAIDMP Overview

Requirements for Use 2-2
Restrictions 2-3
General Restrictions 2-3
DF350F Subsystem Restrictions 2-4
DF500 Subsystem Restrictions 2-4
RAID300 Subsystem Restrictions 2-4
RAID300/RAID400 Subsystem Fibre Channel Restrictions 2-4
Software Package Contents 2-6
Sample VTOCs 2-7
For DF350F Use 2-7
For RAID200, RAID300, or RAID400 Use 2-7
For RAID300 or RAID400 Use 2-7
For RAID400 Use ONLY 2-7
For Use With ptx/SVM 2-7
Installing ptx/HITACHI-RAIDMP 2-8
Installation Notes 2-8
First-Time Installation 2-9
ptx/HITACHI-RAIDMP Upgrade 2-11
Alt-Disk O/S Upgrade 2-13
Notes on Use and Administration 2-14
Using ptx/HITACHI-RAIDMP with ptx/SVM 2-14
Missing, Invalid or Expired License Key Will Cause Errors 2-14
Replace VTOCs With Caution 2-15

Chapter 3 DF350F Subsystem Integration

DF350F Requirements for Use 3-2
DF350F Restrictions 3-3
DF350F RAID Concepts 3-4
RAID Groups 3-4
Logical Units (LUs) 3-7
Multiporting on the DF350F Subsystem 3-8
DF350F Integration Overview 3-10
Step A: Install ptx/HITACHI-RAIDMP 3-11
Step B: Prepare for Integration 3-11
Step C: Rack Mount the Subsystem, If Required 3-12
Install Mounting Rails in the Expansion Cabinet 3-12
Install the 10-Drive Subsystem in the Expansion Cabinet 3-15
Install the 20-Drive Subsystem in the Expansion Cabinet 3-17
Step D: Check and Adjust Configuration Settings 3-19
Critical Settings 3-19
DF350F Subsystem SVP Function and EEPROM Settings 3-20
Checking or Adjusting Settings 3-20
Entering the CONTROLLER NAME/ENABLE Value 3-21
Summary of DF350F EEPROM Settings 3-21
Step E: Make Cable Connections 3-24

Step F: Verify the Integration 3-26
Notes on Use and Administration 3-29
Sample VTOCs Supplied With ptx/HITACHI-RAIDMP 3-29
The **make_vtoc**(1M) Utility 3-29
The **dfutil**(1M) Utility 3-29

Chapter 4 RAID200 Subsystem Integration

RAID200 Requirements for Use 4-2
DYNIX/ptx System Requirements 4-2
RAID200 Subsystem Requirements 4-2
SCSI Cabling and Termination Requirements 4-2
RAID200 Restrictions 4-3
General Subsystem Restrictions 4-3
SCSI Connectivity Restrictions 4-3
RAID200 RAID Concepts 4-4
Terminology 4-4
Multiported Devices 4-5
RAID200 Integration Overview 4-6
Step A: Install ptx/HITACHI-RAIDMP 4-7
Step B: Prepare for Integration 4-7
Step C: Configure the RAID200 Subsystem 4-8
Getting to the Configuration Screens 4-8
Setting the Host Mode 4-10
Adding, Changing, or Deleting a Path Definition 4-10
Renewing the Subsystem (Making New Settings Take Effect) 4-11
Step D: Make Cable Connections 4-12
Step E: Verify the Integration 4-14
Notes on Usagee and Administration 4-17
Obtaining Subsystem Status 4-17
Decoding CHS/Port Errors Reported in **ktlog** 4-19
Sample VTOCs Supplied with ptx/HITACHI-RAIDMP 4-19
The **make_vtoc**(1M) Utility 4-19

Chapter 5 RAID300 Subsystem Integration

RAID300 Requirements for Use 5-2
DYNIX/ptx System Requirements 5-2
IBM NUMA-Q Requirements 5-2
Symmetry S5000 Requirements 5-2
RAID300 Subsystem Requirements 5-2
RAID300 Host and System Modes 5-3
Fibre Cabling Requirements 5-4
SCSI Cabling and Termination Requirements 5-4
IBM NUMA-Q Connections 5-4
Symmetry S5000 Connections 5-4
SCSI Termination—FCB SCSI and QCIC-E SCSI 5-4
RAID300 Restrictions 5-5

General Restrictions 5-5
Fibre Channel Connectivity Restrictions 5-5
SCSI Connectivity Restrictions 5-5
NUMA-Q SCSI Restrictions 5-5
Symmetry S5000 SCSI Restrictions 5-6
RAID300 RAID Concepts 5-7
Terminology 5-7
Multiported Devices 5-8
RAID300 Integration Overview 5-9
Step A: Install ptx/HITACHI-RAIDMP 5-10
Step B: Prepare for Integration 5-10
Step C: Configure the RAID300 Subsystem 5-11
Configuration for Fibre Channel Connection 5-12
Changing Configuration Settings 5-12
Changing the Host Mode 5-14
Setting the Fibre Topology 5-15
Adding a Path Configuration 5-16
Configuring Volumes for LUSE 5-18
Deleting Path-LDEV Configurations 5-21
Completing the Configuration Process for Fibre Channel Connection 5-22
Configuration for SCSI Connection 5-26
Changing Configuration Settings 5-26
Changing the Host Mode 5-28
Deleting Path-LDEV Configurations 5-29
Adding a Path Configuration 5-30
Completing the Configuration Process for SCSI Connection 5-31
Step D: Make Cable Connections 5-32
Step E: Verify the Integration 5-35
Notes on Useage and Administration 5-43
Obtaining Subsystem Status 5-43
Typical **ktlog** Output Showing Route Failure and Recovery 5-45
Sample VTOCs Supplied with ptx/HITACHI-RAIDMP 5-45
The **make_vtoc**(1M) Utility 5-45
LUN Security 5-46
LUN Security Overview 5-46
Getting Ready to Use LUN Security 5-47

Chapter 6 RAID400 Subsystem Integration

RAID400 Requirements for Use 6-2
DYNIX/ptx System Requirements 6-2
RAID400 Subsystem Requirements 6-2
RAID400 Host and System Modes 6-2
Fibre Cabling Requirements 6-3
Related Documentation 6-4
RAID400 Restrictions 6-4
RAID400 Concepts 6-5

Terminology 6-5
Multiported Devices 6-6
RAID400 Integration Overview 6-7
StepA: Install ptx/HITACHI-RAIDMP 6-8
Step B: Prepare for Integration 6-8
Step C: Configure the RAID400 Subsystem 6-9
Configuration Overview 6-9
Configuration Notes 6-9
Step D: Make Cable Connections 6-10
Step E: Verify the Installation 6-12
Notes on Use and Administration 6-18
Obtaining Subsystem Status 6-18
Typical **ktlog** Output Showing Route Failure and Recovery 6-20
Sample VTOCs Supplied With ptx/HITACHI-RAIDMP 6-20
The **make_vtoc**(1M) Utility 6-20
LUN Security 6-21
LUN Security Overview 6-21
Getting Ready to Use LUN Security 6-22

Chapter 7 DF500 Subsystem Integration

DF500 Requirements for Use 7-2
Related Documentation 7-2
DF500 Restrictions 7-3
DF500 RAID Concepts and Architecture 7-4
Basic DF500 Architecture 7-4
Multiporting on the DF500 Subsystem 7-6
DF500 Integration Overview 7-8
Step A: Install ptx/HITACHI-RAIDMP 7-9
Step B: Prepare for Integration 7-9
Step C: Check and Adjust Configuration Settings 7-10
Required Settings 7-10
Controller 0 and 1, Common Parameters 7-10
Controller 0, Controller 1 Separately 7-10
Fibre Channel 7-10
Step D: Make Cable Connections 7-11
Step E: Verify the Integration 7-13
Notes on Use and Administration 7-16
The **make_vtoc**(1M) Utility 7-16
LUN Security 7-17
LUN Security Overview 7-17
Getting Ready to Use LUN Security 7-18

Appendix A Man Pages

List of Illustrations

Figure	Description	Page
3-1	Basic Subsystem Structure and Disk Layout	3-4
3-2	RAID-0 or RAID-5 Group	3-5
3-3	RAID-1 and RAID-0/1 (RAID-A) Groups	3-6
3-4	Larger RAID Group	3-6
3-5	Logical Units Within a RAID Group	3-7
3-6	Multiport Configuration	3-8
3-7	Route Failure Example	3-9
3-8	Installing the Guide-Rail Adapter (Left-Side Mounting Rail)	3-13
3-9	Installing the Mounting Rails for a 10-Drive Unit	3-14
3-10	Installing the Mounting Rails for a 20-Drive Unit	3-14
3-11	Rack Mounting the 10-Drive Subsystem	3-16
3-12	Rack Mounting the 20-Drive Subsystem	3-18
3-13	Single Controller Connection	3-24
3-14	Dual Controller Connection to One FC-SW Switch	3-25
3-15	Dual Controller Connection to Two FC-SW Switches	3-25
3-16	**dumpconf**, **diskid**, and **infodev** Output for Switched Fabric (FC-SW) Connection	3-27
4-1	Sample Multiported Disk Storage Unit	4-5
4-2	Multiport RAID200 Configuration With Two 4-Port Channel Adapters for SCSI (CHSs)	4-13
4-3	Sample **dumpconf** and **diskid** Output	4-15
4-4	CHS, Port Mappings and Typical Connection	4-16
5-1	Sample Multiported Disk Storage Unit	5-8
5-2	Switched Fabric (FC-SW) Connection to an IBM NUMA-Q System	5-33
5-3	Arbitrated Loop (FC-AL) Connection to an IBM NUMA-Q System	5-33
5-4	FCB SCSI Connection to an IBM NUMA-Q System	5-34
5-5	QCIC-E SCSI Connection to an IBM Symmetry S5000 System	5-34
5-6	Sample **dumpconf**, **infodev** and **diskid** Output for Switched Fabric (FC-SW) Connections	5-37
5-7	Sample **dumpconf**, **infodev** and **diskid** Output for Arbitrated Loop (FC-AL) Connections	5-38
5-8	Sample **dumpconf**, **infodev** and **diskid** Output for SCSI Connection	5-39
5-9	CHS, Port Mappings for SCSI Connection	5-37
5-10	Example Connection Scheme with LUN Security Feature	5-46
6-1	Sample Multiported Disk Storage Unit	6-6
6-2	Switched Fabric (FC-SW) Connection to an IBM NUMA-Q System	6-11
6-3	Arbitrated Loop (FC-AL) Connection to an IBM NUMA-Q System	6-11

6-4 Sample **dumpconf**, **infodev**, and **diskid** Output for Fibre Channel (FC-SW or FC-AL) Connections 6-14
6-5 Sample **hrutil** Output 6-15
6-6 Example Connection Scheme with LUN Security Feature 6-21

7-1 DF500 System Architecture 7-5
7-2 Multiport Configuration 7-6
7-3 Single Primary Path Failure 7-7
7-4 Dual Primary Path Failure 7-7
7-5 DF500 Back Panel Connections 7-11
7-6 Typical Switched Fabric (FC-SW) Fibre Channel Connection 7-12
7-7 Typical Direct Connect Arbitrated Loop (FC-AL) Fibre Channel Connection 7-12
7-8 Sample **dumpconf**, **diskid**, and **infodev** Output for FC-SW Connection 7-14
7-9 Sample **dumpconf**, **diskid**, and **infodev** Output for FC-AL Connection 7-15
7-10 Example Connection Scheme with LUN Security Feature 7-17

List of Tables

Table	Description	Page
1-1	DF350F Models Supported	1-4
1-2	DF350F Feature Support Summary	1-4
1-3	RAID200 Models Supported	1-5
1-4	RAID200 Feature Support Summary	1-5
1-5	RAID300 Models Supported	1-6
1-6	RAID300 Feature Support Summary	1-6
1-7	RAID400 Models Supported	1-8
1-8	RAID400 Feature Support Summary	1-8
1-9	DF500 Models Supported	1-9
1-10	DF500 Feature Support Summary	1-9
1-11	FC Switches Supported	1-10
3-1	DF350F EEPROM Settings—Microcode Revision 4356/M	3-22
4-1	RAID200 Terminology	4-4
4-2	CHS/Port Mappings	4-15
4-3	Status Icons and Contents	4-18
5-1	Host and System Modes	5-3
5-2	RAID300 Terminology	5-7
5-3	CHS/CHF-to-Port Mappings	5-41
5-4	Status Icons and Contents	5-44
6-1	Host and System Modes	6-3
6-2	RAID400 Terminology	6-5
6-3	CHF-to-Port Mappings	6-16
6-4	Status Icons and Contents	6-19

Chapter 1
Introduction

About This Manual

This integration guide gives an overview of the ptx® / HITACHI-RAIDMP V2.3.0 layered software package, and tells how to install the software and how to integrate selected Hitachi RAID disk array subsystems with DYNIX / ptx® host systems running ptx / HITACHI-RAIDMP.

Document Organization

This guide is organized as follows:

Chapter 1: Introduction
Overview of this document, a summary of the RAID subsystems supported, and release-specific information, including new or changed functionality or support, problems fixed, and problems known to exist in this release.

Chapter 2: ptx/HITACHI-RAIDMP Overview
Overview of the ptx / HITACHI-RAIDMP software package, requirements for use, restrictions, installation instructions, and notes on the use and administration of ptx / HITACHI-RAIDMP.

Chapter 3: DF350F Subsystem Integration
Chapter 4: RAID200 Subsystem Integration
Chapter 5: RAID300 Subsystem Integration
Chapter 6: RAID400 Subsystem Integration
Chapter 7: DF500 Subsystem Integration
Each of these chapters summarizes the information necessary to integrate a particular Hitachi RAID subsystem with a DYNIX / ptx host system. Topics include: subsystem features supported, requirements for use, restrictions, instructions for configuring and connecting the subsystem to the host system, and notes on the use and administration of the subsystem with ptx / HITACHI-RAIDMP.

Appendix A: Man Pages
Hard copies of the on-line man pages for this release of ptx / HITACHI-RAIDMP.

Audience

This document assumes that the reader is familiar with:

- Hardware and software installation procedures for the appropriate host system.
- DYNIX/ptx system administration tasks
- General RAID concepts and terminology

How to Use This Document

To integrate one of the RAID subsystems discussed in this manual with a host system running DYNIX/ptx, complete the following steps.

1. Read Chapter 1 (this chapter) carefully.

2. Read Chapter 2 carefully, making sure that you understand the requirements for use and the restrictions of the ptx/HITACHI-RAIDMP software.

3. Refer to the integration chapter for the particular RAID subsystem and do the following:

   a. Make sure that you understand the RAID concepts as implemented in the subsystem.

   b. Make sure that you understand the requirements for integrating the subsystem with the particular host system and DYNIX/ptx release, as well as any restrictions that apply to the integration and use of this subsystem.

   c. Using the instructions provided in that chapter, along with the software installation instructions in Chapter 2, complete the procedures required to integrate the subsystem with the host system.

When integration has been completed, refer to this document as necessary for the ongoing administration and use of the Hitachi RAID subsystem.

Related Documentation

The following documentation contains essential or useful additional information for this product:

Appropriate host system hardware and software documentation
DYNIX/ptx V4.4 software installation documentation

In addition, you should have access to all documentation supplied with the RAID subsystem. If you are using the LUN Security feature of the RAID300, RAID400, or DF500 subsystem, be sure to consult the LUN security documentation for that subsystem.

Hitachi RAID Subsystem Support Overview

NOTES

Some earlier versions of ptx/HITACHI-RAIDMP included support for the Hitachi DF400. Because DF400 support is included in the DYNIX/ptx V4.4.8 base operating system, ***the DF400 is not supported*** *as part of this release of ptx/HITACHI-RAIDMP.*

*Also note that with DYNIX/ptx V4.4.8, support for SCSI connection of DF350 subsystems is provided via ptx/HITACHI-RAID (**hrd** device driver). The **hrd** driver and ptx/HITACHI-RAIDMP driver, along with their associated RAID subsystems, can co-exist on the same DYNIX/ptx host system.*

The ptx / HITACHI-RAIDMP software package enables host systems running selected versions of DYNIX / ptx to connect to and control selected Hitachi RAID disk array subsystems. This release of ptx / HITACHI-RAIDMP supports the following subsystems and connectivities:

- DF350F
  IBM NUMA-Q switched fabric (FC-SW) Fibre Channel connection
- RAID200
  IBM NUMA-Q Fibre Channel Bridge (FCB) SCSI connection
- RAID300
  IBM NUMA-Q switched fabric (FC-SW) and direct connect arbitrated loop (FC-AL) Fibre Channel connections
  IBM NUMA-Q Fibre Channel Bridge (FCB) SCSI connection
  Symmetry S5000 QCIC-E SCSI connection
- RAID400
  IBM NUMA-Q switched fabric (FC-SW) and direct connect arbitrated loop (FC-AL) Fibre Channel connections
- DF500
  IBM NUMA-Q switched fabric (FC-SW) and direct connect arbitrated loop (FC-AL) Fibre Channel connections

The discussions on the following pages summarize the supported models and features.

NOTES

Hitachi distributors vary with geography. The following support summaries include model designations for selected Hitachi distributors. However, throughout the remainder of this guide, Hitachi model designations are used exclusively.

On certain IBM NUMA-Q platforms, other connectivities may be possible via the ptx/HITACHI-RAID software package. Contact your IBM NUMA-Q sales representative for more information.

Minimum microcode levels for each supported Hitachi RAID subsystem are given in the following support summary tables. Contact the IBM NUMA-Q Customer Support Center for the latest supported microcode releases.

DF350F Support Summary

The following tables summarize support for Hitachi DF350F RAID disk array subsystems in this release of ptx/HITACHI-RAIDMP.

Table 1-1
DF350F Models Supported

Hitachi Model	HDS Model[a]
DF350F-RKH	5742E
DF350F-RKWH	5744E
DF350F-CK	5750E

a. HDS chassis models with nomenclature 57*XX*-A1001, where *XX* = 42, 44, or 50, are configured for fibre connection.

Table 1-2
DF350F Feature Support Summary

RAID Features	**RAID Levels Supported**	0, 1, 5, and “0+1” (“A”)
	Configurability	Multiple Logical Units per RAID group Mixed RAID groups
Connectivity	**IBM NUMA-Q Connectivity**	Fibre Channel (FC-SW) *only* Levels I, II, and III supported
	Network Connectivity	SNMP LAN (optional) with management software
Redundancy Support	**Failover Support**	Yes
	Component Redundancy	Redundant “N+1” AC power supplies, cooling fans, and AC power inputs
Hardware Support	**Microcode Levels**	4356/M or above; contact the IBM NUMA-Q Customer Support Center for information on the latest supported microcode release(s).
	RAID Controllers	1 or 2
	Cache	Up to 1GB of mirrored, battery-backed cache
	Disk Types	9GB, 18GB 7200RPM
	Total Disk Storage per Subsystem	356GB (268GB, RAID-5)
	Spare Disk Support	One (1) global hot spare

RAID200 Support Summary

The following tables summarize support for Hitachi RAID200 RAID disk array subsystems in this release of ptx/HITACHI-RAIDMP.

Table 1-3
RAID200 Models Supported

Hitachi Model	HDS Model
RAID200	7700

Table 1-4
RAID200 Feature Support Summary

RAID Features	**RAID Levels Supported**	5
	Configurability	6D + 1P
Connectivity	**IBM NUMA-Q Connectivity**	FCB SCSI *only* Up to 16 SCSI connections are supported
	Network Connectivity	Hi-Track only
Redundancy Support	**Failover Support**	DYNIX/ptx multiport capability Single-port and multiport configurations are supported Up to 2 paths are supported
	Component Redundancy	Full hardware redundancy
Hardware Support	**Microcode Levels**	03-16-60-00/00, minimum; contact the IBM NUMA-Q Customer Support Center for information on the latest supported microcode release(s).
	RAID Controllers	2 to 4
	Cache	4GB
	Disk Types	9GB disk modules
	Total Disk Storage per Subsystem	1.2TB (RAID-5)
	Spare Disk Support	Up to 10 global spares
	LU Formats	OPEN-3, OPEN-9, LUSE

RAID300 Support Summary

The following tables summarize support for Hitachi RAID300 RAID disk array subsystems in this release of ptx/HITACHI-RAIDMP.

Table 1-5
RAID300 Models Supported

Hitachi Model	HDS Model
RAID300	7700E

Table 1-6
RAID300 Feature Support Summary

RAID Features	**RAID Levels Supported**	1, 5
	RAID-1 Configurability	6D + 1P (9GB disk modules) 3D + 1P (6GB, 15GB, 18GB, 36GB, 47GB disk modules) LUN Security supported
	RAID-5 Configurability	FCB SCSI: Up to 32 SCSI connections Fibre Channel connection FC-SW or direct connect FC-AL, Levels I, II, and III Up to 16 Fibre connections
Connectivity	**IBM NUMA-Q**	FCB SCSI: Up to 32 SCSI connections Fibre Channel: FC-SW and direct connect FC-AL, Levels I, II, and III Up to 16 Fibre connections
	IBM Symmetry S5000	QCIC-E SCSI
	Network Connectivity	Hi-Track, Remote Console, Graph Track
Redundancy Support	**Failover Support**	DYNIX/ptx multiport capability. Single-port and multiport configurations are supported. Up to 2 paths are supported for SCSI connections. Up to 8 paths are supported for FC connections.
	Component Redundancy	Full hardware redundancy
Hardware Support	**Microcode Levels**	52-47-71-00/00, minimum; contact the IBM NUMA-Q Customer Support Center for information on the latest supported microcode release(s).
	RAID Controllers	2 to 8 (1 to 4 pairs)
	Cache	Up to 16GB
	Disk Types	6K RPM: 9GB, 18GB, 36GB disk modules 12K RPM: 6GB, 15GB, 47GB disk modules
	Total Disk Storage per Subsystem	Up to 2.9TB (RAID-5, with 18GB disk modules) Up to 5.8TB (RAID-5, with 36GB disk modules) Up to 10.4TB (RAID-5, with 47GB disk modules)
	Spare Disk Support	Up to 8 global spares

Table 1-6
RAID300 Feature Support Summary

	LU Formats	OPEN-3, OPEN-8, OPEN-9, LUSE, 3380-X, 3390-X, 6586-X, 6588-X, OP-C-X

RAID400 Support Summary

The following tables summarize support for Hitachi RAID400 RAID disk array subsystems in this release of ptx/HITACHI-RAIDMP.

Table 1-7
RAID400 Models Supported

Hitachi Model	HDS Model
RAID400	9900

Table 1-8
RAID400 Feature Support Summary

RAID Features	**RAID Levels Supported**	1, 5
	RAID-1 Configurability	1D + 1M (18GB, 47GB, 72GB disk modules)
	RAID-5 Configurability	3D + 1P (18GB, 47GB, 72GB disk modules) LUN Security supported
Connectivity	**IBM NUMA-Q Connectivity**	Fibre Channel: FC-SW or direct connect FC-AL Levels I, II, and III Up to 32 Fibre connections Up to 256 LUNs per port; max of 4096 LUNs per system
	Network Connectivity	Hi-Track, Remote Console, Graph Track
Redundancy Support	**Failover Support**	DYNIX/ptx multiport capability Single-port and multiport configurations are supported Up to 8 paths are supported for fibre connections
	Component Redundancy	Full hardware redundancy
Hardware Support	**Tested Microcode Levels**	01-11-23-00/04, minimum; contact the IBM NUMA-Q Customer Support Center for information on the latest supported microcode release(s).
	RAID Controllers	2 to 8 (1 to 4 pairs)
	Cache	Up to 32GB
	Disk Types	10K RPM: 18GB, 47GB, 72GB disk modules
	Total Disk Storage per Subsystem	Up to 4.3TB (RAID-5, with 18GB disk modules) Up to 11.2TB (RAID-5, with 47GB disk modules) Up to 17.3TB (RAID-5, with 72GB disk modules
	Spare Disk Support	Up to 16 global spares
	LU Formats	OPEN-3, OPEN-8, OPEN-9, OPEN-E, 3380-X, 3390-X, LUSE, 6586-X, 6588-X, OP-C-X

DF500 Support Summary

The following tables summarize support for Hitachi DF500 RAID disk array subsystems in this release of ptx/HITACHI-RAIDMP.

Table 1-9
DF500 Models Supported

Hitachi Model	HDS Model
DF500	9200

Table 1-10
DF500 Feature Support Summary

RAID Features	**RAID Levels Supported**	0, 1, 5, 0+1
	RAID-0	2D to 16D
	RAID-1 Configurability	1D + 1P
	RAID-5 Configurability	2D + 1P to 15D + 1P LUN Security supported
	RAID 0+1 Configurability	2D + 2P to 8D + 8P
Connectivity	**IBM NUMA-Q Connectivity**	Fibre Channel: FC-SW or FC-AL, Levels I, II, and III Up to 4 connections
	Network Connectivity	Disk Array Management Program (DAMP); Ethernet or RS232-to-PC connection
Redundancy Support	**Failover Support**	DYNIX/ptx multiport capability Single-port and multiport configurations are supported Up to 4 paths are supported for fibre connections
	Component Redundancy	Full hardware redundancy
Hardware Support	**Tested Microcode Levels**	0555/C, minimum; contact the IBM NUMA-Q Customer Support Center for information on the latest supported microcode release(s).
	RAID Controllers	Two (2) controllers; four (4) FC interfaces
	Cache	Up to 2GB
	Disk Types	36GB, 72GB disk modules
	Total Disk Storage per Subsystem	Up to 3.6TB with 36GB disk modules Up to 7.2TB with 72GB disk modules
	Spare Disk Support	Up to 5 global spares

FC Switch Support and Operation

FC Switches Supported

Connection of Hitachi RAID subsystems to IBM NUMA-Q host systems via an FC switch must be via one of the following switches. The switch must be running the minimum firmware code revision shown below.

Table 1-11
FC Switches Supported

FC Switch	Minimum Firmware Revision
IBM Model 2109	2.1.3
Brocade® Silkworm® 2000 (equivalent to IBM Model 2109)	2.1.3
Brocade® Silkworm® 1000	1.6.3

FC Switch Must Be Set to 1000 Compatibility Mode

For correct operation with IBM NUMA-Q systems running DYNIX/ptx V4.4.x, the IBM Model 2109 (Silkworm 2000) switch *must be* in Silkworm 1X00 compatibility mode. Some switches were shipped from the factory with 1X00 compatibility mode already set; other models will need to be set at the installation site.

To set or verify that the Model 2109 (Silkworm 2000) is set to 1X00 compatibility mode:

1. Log onto the switch as **admin**.
2. Enter **switchDisable**.
3. Enter the following commands in the order shown:
   a. **configDefault** (to begin with the factory default settings)
   b. **configure** and make sure that the fabric parameters are set as follows:

      ```
      VC Encoded Address = 1
      Disable Translative Mode = 1
      ```

   c. **iodSet** (sets up the switch for “in order delivery”)
   d. **dlsReset**
4. Reboot the switch.
5. Enter **configShow** (to verify the settings).
6. Log off of the switch.

New or Changed in This Release

The following items are new or changed in this release of ptx/HITACHI-RAIDMP.

- This release of ptx/HITACHI-RAIDMP runs on DYNIX/ptx V4.4.7 or V4.4.8 *only*.
- This release supports Fibre Channel switched fabric (FC-SW) and direct connect arbitrated loop (FC-AL) connections of the Hitachi RAID400 and DF500 subsystems.
- This release of ptx/HITACHI-RAIDMP has been tested with RAID400 microcode release 01-11-23-00/04. Contact the IBM NUMA-Q Customer Support Center for information on the latest supported microcode release(s).
- This release of ptx/HITACHI-RAIDMP has been tested with RAID300 microcode release 52-47-71-00/00. Contact the IBM NUMA-Q Customer Support Center for information on the latest supported microcode release(s).
- A new utility, **hrutil**, reports device information for RAID400 subsystems.
- The **dfutil** utility has been enhanced to report device information for DF500 subsystems.
- OPEN-E volumes are supported for the RAID400 subsystem.
- This release supports the RAID300, RAID400, and DF500 LUN Security features.
- This release supports the DF500 TurboLU feature.

Problems Fixed in This Release

No user-detectable problems have been fixed in this release of ptx/HITACHI-RAIDMP.

Known Problems in This Release

ptx/HITACHI-RAIDMP Issues

PR#251591 - PANIC: cfg_coalesce: found recursing coalesce scenario

On IBM NUMA-Q host systems running DYNIX/ptx V4.4.8, if the system is brought up with a Hitachi RAID subsystem connected but without ptx/HITACHI-RAIDMP installed, the DYNIX/ptx host system may panic or exhibit other undesirable behavior.

Workaround. Be sure that ptx/HITACHI-RAIDMP is installed before connecting the Hitachi RAID subsystem to the DYNIX/ptx system.

PR#251854 - "sar: corrupted disk response and active times" w/ HITACHI disks

> **NOTE**
> *The following is a problem in the accuracy of **sar** information reports, and **does not indicate actual data corruption**.*

Under certain conditions, the **sar** utility may report data corruption. This **does not** indicate actual data corruption, but rather a problem with **sar** itself.

This problem has been seen in test scenarios running disk tests against an EFS filesystem comprising Hitachi RAID disk devices. These errors have not been seen on raw disk I/O.

The following is a sample **sar** log that illustrates the reporting error:

```
# sar -d 6 10

Wed Mar  1 13:22:35 PST 2000

DYNIX/ptx elm9a41 4.0 V4.4.8 i386    03/01/00

13:22:35  device   %busy   avque   r+w/s  blks/s  avwait  avserv

13:22:41      sd0       0     1.0       0       3     0.0    11.0
                 sd20      94   128.7      22    5604  5481.2    42.9

              sd30     103    88.1      17    4473  5146.7    59.1
              sd31     102   242.6      20    4984 12510.3    51.8
              sd32     101    95.2      21    5263  4448.8    47.2
              sd34   sar: corrupted disk response and active times
  100   525.9      24    6007 22164.2    42.2
              sd35     101   150.5      19    4984  7791.9    52.1
              sd37      94   502.7      18    4600 26118.8    52.1

sar: corrupted disk response and active times
sar: corrupted disk response and active times
sar: corrupted disk response and active times
```

PR#247568 - S/W Management Package Consistency Check Fails With Files Not Found

When you use the ptx/ADMIN menu system to install ptx/HITACHI-RAIDMP, the "Software Management" option under "Check Package Consistency" fails because several files are not found for **raidmp**. The files are identified by a number of messages similar to the following:

```
/usr/lib/raidmp/sd_r200.o: file not found or can't lstat!
```

Files listed in these messages are placed in */usr/lib/raidmp/.save* for safekeeping.

Workaround. None. You can ignore these messages. *Do not* attempt to delete the files listed in these messages from */usr/lib/raidmp/.save.*

RAID400 Issues

FC Switch Ports May Log In Incorrectly

When a microcode load occurs on the RAID400, the FC switch ports may re-login as lports, instead of as fports. This problem has been seen when adding or changing LDEVs, when the switch reboots, or when switch ports are disabled, then re-enabled.

Workarounds. To correct this problem, unplug and re-connect the fiber cabling, then use the **switchShow** command to ensure that the ports have logged in correctly. If the ports are still not logged in correctly, you may need to repeat this procedure.

DF500 Issues

FC Switch Ports May Log In Incorrectly

If a connection between the DF500 subsystem and the FC switch is broken, then re-connected, the FC switch ports may re-login as lports, instead of as fports.

Workarounds. To correct this problem, you can (1) power cycle the FC switch, (2) unplug and re-connect the fiber cabling, or (3) use a switch **telnet** session with the **portDisable** and **portEnable** commands to disable, then re-enable the ports. Then, use the **switchShow** command to ensure that the ports have logged in correctly. If the ports are still not logged in correctly, you may need to repeat this procedure.

RAID200/RAID300 Issues

SCSI Connections—Path Failure May Require SCSI Bus Reset

In a multiport configuration where more than one SCSI path exists for the Logical Devices (LDEVs), if a path failure occurs in the middle of an I/O, the SCSI bus may be rendered unusable until a bus reset occurs. (This has occurred when a SCSI cable was pulled, but might also occur with path failures due to other causes.) This problem is due to an interaction between the DYNIX/ptx host system's SCSI controllers and the subsystem's RAID SCSI controllers.

FCB SCSI Problem Indications

With FCB SCSI connections, if this problem occurs, messages similar to the following will appear in **ktlog**:

```
36545989 09:46:49 tolog/warn q1/e6 fcbr0 (FC port): Hardware Error on SCSI 0x0:
Test Unit Ready   code=0x70<0x70> Current Error seg=0x0 key=0x4<0x4>
36545989 09:46:49 tolog/cont q1/e6 info=0x0 addlen=0x20 cmdinfo=0x0 asc=0x44 asc
q=0x0   internal target failure   fru=0x0 keyinfo=0x0
```

Re-routing I/O from the failed path to other paths will work successfully. This problem only affects the recovery of the failed path, once the cable is re-installed.

Workarounds. A workaround for this problem would be to do one of the following:

- Schedule down time to restore the failed path and power cycle the FCB.
- Offline the FCB and cycle power to the subsystem, then use **devctl -c** to bring the FCB and SCSI buses back online.

QCIC-E SCSI Problem Indications

On subsystems attached to Symmetry S5000 systems via QCIC-E SCSI, if this problem occurs, messages similar to the following will appear in **ktlog**:

```
371e63af 16:47:59 tolog/warn e4/p4656 sd5(port 82): I/O timeout, hung bus?
371e63b0 16:48:00 tolog/warn e9 qcic2 (at slic 26): internal or bus protocol
error on SCSI channel 1
```

Workaround. Reboot the Symmetry host system.

FCB SCSI Connections—FCB Microcode May Run Out of Memory at Boot

Older Fibre Channel Bridges (FCBs) have less RAM than new models. On the older model FCBs, the microcode can run out of memory during boot as it tries to allocate software structures for devices. The algorithm used for allocation is based on SCSI Target IDs (TIDs) and Logical Unit Numbers (LUNs) that the operating system is detecting. For a large RAID configuration that uses many TIDs, the FCB allocates memory for all of the available LUNs automatically, and can quickly run out of memory. The error in **ktlog** would look something like this:

```
35ae09ed 07:10:53 tolog/warn q1/e7 +fcbr3 (FC port): No Sense on SCSI
      0x12: Inquiry   code=0x70<0x70> Current Error seg=0x0 key=0x0<>
35ae09ed 07:10:53 tolog/cont q1/e7 info=0x0 addlen=0x20 cmdinfo=0x0
      asc=0x25 ascq=0x0   logical unit not supported   fru=0x0
      keyinfo=0x0
35ae09ed 07:10:53 warn q1/e7 +scsibus32: probe of target ID 13, lun 1:
      adapter status 0xc (controller error)
```

Workarounds. Either of the following workarounds can be used to alleviate this problem:

- Use LUNs more than TIDs. For example, for seven LDEV assignments, instead of using TIDs 8 - 14, all with LUN0, use TID 8, LUNs 0 - 6.
- Use Fibre Channel Bridge Model 45P0540.

RAID300 Issues

Issues Relating to Microcode Downloads

The following issues are associated with a number of microcode releases for RAID300 subsystems, including the release(s) tested with this release of ptx/HITACHI-RAIDMP.

FC Ports Do Not Log In Properly After Microcode Installation

As a microcode installation is completed, Tachyon Fibre Channel ports are reset and do not log back into the switches as expected. (Refer to Hitachi's ECN DKC310I-M042, #11.)

From the perspective of the ptx/HITACHI-RAIDMP driver, this is considered a non-critical problem because the vendor does not support microcode downloads with I/O present (see below). Thus, all I/O should already have been stopped to enable the microcode download. However, this issue and the following one are under investigation.

Workarounds. If this problem is encountered, do one of the following.

- Pull and re-attach the cable between the RAID300 FC port and the Silkworm 2X00 switch.

  or

- Log onto the switch, disable the ports, then re-enable them. This causes the switch to re-login.

Vendor Warns Against Downloading Microcode With I/O Present

Currently, Hitachi states that a microcode download cannot be performed with I/O present. This issue is under investigation.

Workaround. The following *Online Microcode Installation* procedure serves as a workaround for this problem.

Online Microcode Installation Procedure with Manual Intervention

Downloading microcode can be done without stopping I/O operations *as long as* (1) there exist at least two (2) paths to each LDEV, and (2) there is at least one path into each subsystem cluster (Cluster 1 and Cluster 2). *This procedure requires manual intervention, as well as an Ethernet connection to the switch.* To download the microcode using this scenario, do the following:

1. Download the microcode as usual, selecting Alternate SCSI Path.

2. You will be prompted, "Please switch SCSI Path which is under Cluster 1". At this prompt, press RETURN to start the actual upgrade to the ports located on Cluster 1.

3. You will then be prompted, "Please switch SCSI Path, which is under Cluster 2". ***Before you press RETURN***, do the following:

   a. Log into the switch or switches to which the Cluster 1 ports are connected.

   b. Issue a **portDisable** to all ports that are connected to Cluster 1 ports. (Use **switchShow** to view these ports; they will appear as L-ports in the output of **switchShow**.)

   c. Issue a **portEnable** to all relevant ports.

d. Verify that these ports are all properly logged into the switch.

4. Now press RETURN on the SVP to continue with the download.

5. After the relevant ports on Cluster 2 are upgraded, repeat Step 3 above for Cluster 2 ports.

6. Verify that all routes have been restored.

This procedure should enable a successful microcode download.

Fibre Connections — Changing LDEV Configurations Causes Lost Connections

Changing LDEV configurations on the RAID300 subsystem causes each affected port to do a Link Reset and to re-login to the fabric. While the RAID300 is establishing its login sequence with the switch, the first host adapter attempts to establish a connection with the RAID300 port, and this attempt is aborted with a P_RJT. The current host adapters do not re-probe a port that has sent a P_RJT. This situation is under investigation.

Workaround. *If* there are *at least two (2)* paths to all LUNs, the following will allow you to add or remove LDEVs.

Adding LDEVs. To add an LDEV, do the following:

1. Log into the fabric (via **telnet**).

2. Issue a **portDisable *port#***, where ***port#*** is the fabric port connected to the RAID300 path that is being changed.

3. Add LDEV(s) to that port as described in the RAID300 integration chapter of this manual.

4. Issue a **portEnable *port#*** for the fabric port disabled in Step 2.

5. Verify that the RAID300 ports are logged in correctly.

6. Issue a **devctl -c fabric<*#*>** on the host system to reprobe the fabric for the new LDEV(s).

7. Verify that all previously existing routes have been restored.

Removing LDEVs. To remove an LDEV, do the following:

1. On the host system, issue a **devctl -d sd<*x*>** to all LDEVs that are being removed.

2. Log into the fabric (via **telnet**).

3. Issue a **portDisable** for the ports that are being removed.

4. Remove the LUNs on the disabled path(s).

5. Issue a **portEnable** for the relevant ports.

6. Verify that these ports are logged in correctly.

RAID 200 Issues

Sequential I/O May Be Less Than Expected

Due to the internal caching algorithm of the RAID200 subsystem's microcode, sequential I/O performance on multiport configurations may be less than expected.

Workaround. There is no workaround. However, to minimize the effects of this problem, until the problem is resolved, it is recommended that only single-port connections be used. If redundancy is required and if performance from single logical devices is essential, consider implementing software mirroring using ptx/SVM.

DF350F Issues

PR #236699—Disks May Go Offline Under Heavy Load in Multi-Initiator Environments

In multi-initiator environments, individual disk drives in the DF350F may go offline under heavy usage loads. Generally, two or more disks on the same port but in different rows will fail. The subsystem will run with the spare disk or in regressed mode. After removing and re-inserting the drives, they will recover successfully.

Workaround. None.

Chapter 2

ptx/HITACHI-RAIDMP Overview

This chapter discusses the ptx/HITACHI-RAIDMP layered software package, and includes discussions of the following topics:

- Requirements for use
- Restrictions
- Overview of the ptx/HITACHI-RAIDMP software
- List of sample VTOCs provided with the software
- Software installation instructions
- Notes on the ongoing administration and use of ptx/HITACHI-RAIDMP

Requirements for Use

CAUTION

*This version of ptx/HITACHI-RAIDMP runs **only** under the DYNIX/ptx versions listed below.*

NOTE

The particular platforms and connectivities supported depend on the DYNIX/ptx release and on the particular RAID subsystem. Refer to the integration chapter for a particular subsystem for details.

The following are required to run this version of ptx/HITACHI-RAIDMP:

- Host system running DYNIX/ptx V4.4.7, V4.4.8 or V4.4.9.
- A valid software license key is required for the ptx/HITACHI-RAIDMP software to execute. Refer to the software installation instructions, later in this chapter, for details.
- For SCSI connections, special SCSI cables and terminators are required. Refer to the integration chapter for the particular subsystem for detailed cable and terminator requirements.
- If RAID400 or DF500 devices are to be connected, fastpatch FP# 254922 *must be installed.*

Restrictions

The following restrictions apply to this release of ptx/HITACHI-RAIDMP.

General Restrictions

- No data throughput speeds are expressed or implied. Throughput is dependent on data and system loading and on other site-specific conditions that IBM does not control.
- Any restrictions on this solution does not affect any other standard product restrictions, prerequisites, or configuration guidelines for IBM NUMA-Q or IBM Symmetry systems.
- This software product supports only the Hitach/HDS RAID subsystems documented in this manual. No other Hitachi OEM storage array subsystems are supported as part of this ptx/HITACHI-RAIDMP solution.
- Microcode for the Hitachi/HDS subsystem must be at a level supported by IBM. Contact the IBM NUMA-Q Customer Support Center for information about the correct firmware level(s).
- The Hitachi devices cannot be used as boot or dump devices.
- Installation of the ptx/HITACHI-RAIDMP software must be completed by IBM-approved personnel for IBM to offer further service offerings on this software product.
- For revision level compatibility, customer personnel must notify IBM of any hardware or microcode changes to their Hitachi storage subsystem, so that IBM can ensure that the software package is fully compatible with the customer's configuration.
- Support for newer revisions of Hitachi storage subsystem hardware and microcode may require additional testing by IBM. Such additional testing is *not* included as part of this software solution or its purchase by the customer.

  Refer to the support summaries in Chapter 1 of this document for minimum microcode levels for particular storage subsystems supported in this software release.
- The customer is responsible for upgrading their Hitachi storage subsystem hardware and microcode to meet the minimum revision levels required by IBM.
- This release of this software product may not support some vendor-specific features of the Hitachi storage subsystems. Specifically, IBM does not qualify or directly support the following Hitachi Storage Management features of the storage subsystems:
  - Hitachi Remote Copy (HRC)
  - Hitachi Extended Remote Copy (HXRC)
  - Hitachi Multiplatform Resource Sharing (HMRS)
  - Hitachi On-line Data Migration (HODM)
  - Hitachi Multiplatform Backup/Restore (HMBR)
  - Hitachi Multiplatform Data Exchange (HMDE)

  - Hitachi ShadowImage
  - Hitachi Flash Access

- This software program runs only under the DYNIX/ptx releases listed under *Requirements for Use*, earlier in this chapter, and when running with a specific Hitachi storage subsystem, only under the additional requirements for use listed in the integration chapter for the appropriate subsystem, later in this document. Contact the IBM NUMA-Q Customer Support Center for information on the latest DYNIX/ptx fastpatches required for the installation and use of this software product.
- Connection of any supported storage subsystem must be via the cables and terminator(s), if applicable, listed under *Requirements for Use*, earlier in this document.
- The standard DYNIX/ptx VTOC driver will work with ptx/HITACHI-RAIDMP devices. However, stand-alone diagnostics and format utilities *will not work*. Do not attempt to use these, as unpredictable results may occur.
- The Hitachi DF400 is not supported in this release of ptx/HITACHI-RAIDMP; support for the DF400 is included as part of the DYNIX/ptx V4.4.7 and higher base operating system.
- If you upgrade the operating system using the "Alternate Disk" (Alt-Disk) Installation, available through ptx/ADMIN, go to the discussion titled *Alt-Disk O/S Upgrade*, under *Installing ptx/HITACHI-RAIDMP*, later in this chapter. That discussion contains information critical to a successful Alt-Disk upgrade with ptx/HITACHI-RAIDMP.

DF350F Subsystem Restrictions

- This software package does not support SCSI connection of the DF350 subsystem. Devices controlled by ptx/HITACHI-RAIDMP and *hrd* devices controlled by the earlier Hitachi RAID device driver, ptx/HITACHI-RAID, can co-exist on the same host system. ***Note*** that DF350 devices can be SCSI-connected to the host system via the ptx/HITACHI-RAID ***hrd*** driver.
- A Fibre Channel (FC) switch is required to attach the DF350F to the IBM NUMA-Q host system. Refer to the discussion titled *FC Switch Support and Operation* in Chapter 1 of this document for approved FC switches.

DF500 Subsystem Restrictions

- This release of this software package supports direct-connect (FC-AL) and switched fabric (FC-SW) connections to the DF500 subsystem. No other types of FC connections are supported.
- This software package will not work in clustered server configurations with FC Arbitrated Loop configurations. Only switched fabric (FC-SW) clustered server configurations are supported.
- DF500 storage subsystems cannot be installed in IBM cabinets, only in cabinets supplied by Hitachi or HDS.

RAID200 Subsystem Restrictions

- Connection of RAID200 storage subsystems to DYNIX/ptx host systems can only be made through the IBM FC Bridge, Model 45P0540. Connection through any other SCSI adapter or SCSI bridge product is not supported.
- For the RAID200 subsystem, a maximum of two (2) SCSI paths per logical volume are supported.
- External SCSI cable length is limited to 20m.
- Multi-initiator SCSI configurations are not supported. Daisy-chaining is not allowed.
- If an FC bridge is connected to a RAID200 subsystem, no other third-party SCSI peripherals can be connected to that bridge.

RAID300 Subsystem Restrictions

- All connections between an IBM NUMA-Q host system and a RAID300 storage subsystem must be via the IBM FC Bridge, Model 45P0540. Attachments through other devices, SCSI adapters, or FC bridges are not allowed or supported.
- Connection of IBM Symmetry S5000 host systems to the RAID300 subsystem must be via the host system's QCIC-E SCSI ports.
- For the RAID300 subsystem, a maximum of two (2) SCSI paths per logical volume are supported.
- External SCSI cable length is limited to 20m.
- Multi-initiator SCSI configurations are not supported. Daisy-chaining is not allowed.
- If an FC bridge is connected to a RAID300 subsystem, no other third-party SCSI peripherals can be connected to that bridge.

RAID300/RAID400 Subsystem Fibre Channel Restrictions

- Both direct-connect (FC-AL) and switched fabric (FC-SW) connections are allowed. No other types of FC connections are supported.
- This software will not work in clustered server environments with Arbitrated Loop (FC-AL) configurations.
- Only FC switches listed under *FC Switch Support and Operation* in Chapter 1 of this manual are supported.
- FC connections of the RAID300/RAID400 subsystems are limited to eight (8) SCSI paths per logical volume.

Software Package Contents

CAUTION

The standard DYNIX/ptx VTOC driver will work with ptx/HITACHI-RAIDMP devices. However, stand-alone diagnostics and format utilities (such as /stand/CCSformat) ***will not work****. Do not attempt to use these, as unpredictable results will occur.*

The ptx/HITACHI-RAIDMP software uses the DYNIX/ptx **sd** device driver to control Hitachi RAID disk devices. The ptx/HITACHI-RAIDMP package includes the following:

- Several modules that are linked into the ***sd*** device driver:
  - A module *sd_df350* to provide control of Hitachi DF350F disk devices.
  - A module *sd_df500* to provide control of Hitachi DF500 disk devices.
  - A module *sd_r200* to provide control of Hitachi RAID200, RAID300, and RAID400 disk devices.
- A host utility, **dfutil**, allows the display and/or logging of subsystem status information for DF350F and DF500 subsystems.
- A host utility, **hrutil**, allows the display and/or logging of subsystem status information for RAID400 subsystems.
- A software utility, **make_vtoc**, that allows the creation of custom VTOCs to support variable-sized devices.
- Sample VTOCs.

Sample VTOCs

The ptx/HITACHI-RAIDMP distribution media includes several sample VTOC files. These files will be installed in */etc/vtoc*.

For DF350F Use

hd0_5x9gb	For RAID-0 disk arrays using 9GB disk modules
hd1_4x9gb	For RAID-1, RAID 0+1 disk arrays using 9GB disk modules
hd5_5x4gb	For RAID-5 disk arrays using 4GB disk modules
hd5_5x9gb	For RAID-5 disk arrays using 9GB disk modules

For RAID200, RAID300, or RAID400 Use

open-3	For OPEN-3 volumes
open-9	For OPEN-9 volumes

For RAID300 or RAID400 Use

open-8	For OPEN-8 volumes

For RAID400 Use ONLY

open-e	For OPEN-E volumes

For Use with ptx/SVM

In addition, there are sample VTOCs to support ptx/SVM *sliced disks*:

One-slice disk, 9GB disk modules

hd0_5x9gb.vx	RAID-0
hd1_4x9gb.vx	RAID-1 or RAID 0+1
hd5_5x9gb.vx	RAID-5

Many-slice disk, 9GB disk modules

hd0_5x9gb.vxsl	RAID-0
hd1_4x9gb.vxsl	RAID-1 or RAID 0+1
hd5_5x9gb.vxsl	RAID-5

RAID200, RAID300, or RAID400, one-slice disk

open-3.vxsl	OPEN-3 disk modules
open-9.vxsl	OPEN-9 disk modules
open-e.vxsl	OPEN-E disk modules

RAID300 only, one-slice disk

open-8.vxsl	OPEN-8 disk modules

Refer to the *ptx/SVM Administration Guide* for detailed information on ptx/SVM.

> **NOTE**
> *The standard VTOC, **hd5_5x18gb**, is available as part of the DYNIX/ptx base and can be used with Hitachi DF350F devices.*

Installing ptx/HITACHI-RAIDMP

Installation Notes

- Installing this release of ptx/HITACHI-RAIDMP overlays the following DYNIX/ptx base operating system files, to support Hitachi RAID devices:

  /usr/conf/uts/io/sd/sd_space.c
  /usr/conf/uts/symmetry/unix_std.a

- If RAID400 or DF500 devices are to be connected, fastpatch FP# 254922 *must be installed.*

- A valid software license key is required for the ptx/HITACHI-RAIDMP software to execute. To obtain a license key, contact the IBM NUMA-Q Customer Service Center and provide them with the **nodeid** of the target system.

  CAUTION
  If you install any upgrade or fastpatch to DYNIX/ptx that affects any of these files, you should first de-install ptx/HITACHI-RAIDMP, then install the upgrade or fastpatch, and then re-install ptx/HITACHI-RAIDMP to ensure that the proper versions of the above files are present.

  CAUTION
  If the DYNIX/ptx host system is brought up with Hitachi RAID hardware attached but without the ptx/HITACHI-RAIDMP software installed, the host system may panic or exhibit other undesirable behavior. ***Be sure*** *to install ptx/HITACHI-RAIDMP* ***before*** *bringing up the DYNIX/ptx host system with Hitachi RAID hardware attached.*

First-Time Installation

If this is a first-time installation of ptx/HITACHI-RAIDMP, do the following:

1. Log in as **root** and make sure that the correct version of DYNIX/ptx is installed.

2. If RAID400 or DF500 devices are to be connected, install fastpatch FP# 254922 at this time. You must install this fastpatch *before* installing the ptx/HITACHI-RAIDMP software.

3. Install the ptx/HITACHI-RAIDMP software. To do this:

a. Insert the ptx/HITACHI-RAIDMP CD-ROM into the system's CD drive.

b. Refer to the DYNIX/ptx V4.4 software installation documentation and install the software. This should be a standard CD-ROM installation, with these specifications:

Installation source type = **CD-ROM**
Installation type = **root**
Software Package Part Number = **45P0321**

There should be no file conflicts. If you abort the installation for any reason, then begin it again, the menu system asks if you want to continue the current installation; answer **yes**.

CAUTIONS

Do not *recompile the kernel at this time. The recompile should take place* ***only*** *after the license key has been entered, as described below.*

Be sure to save a hard copy of the ptx/HITACHI-RAIDMP license key. If you ever need to re-install or upgrade ptx/HITACHI-RAIDMP, you will need to re-enter the key.

4. You must enter a software license key for the ptx/HITACHI-RAIDMP software to run. If you do not have a key, contact the IBM NUMA-Q Customer Support Center and provide them with the *nodeid* of the target system. Once you have the key, enter it as follows:

a. Change to the directory */usr/conf/uts/io/sd*, and open the file *sd_space.c* for editing.

b. In the file *sd_space.c*, find the line that reads:

```
char *sd_raidmp_license_key = "No LicenseKey";
```

c. Edit this line, replacing the string `No LicenseKey` with the ptx/HITACHI-RAIDMP license key. *Be sure that the key is enclosed in quotes.*

d. Save and exit the file *sd_space.c*.

e. Save a hard copy of the license key, in case you ever need to re-install the ptx/HITACHI-RAIDMP software.

5. Compile the new kernel.

NOTE

To reduce system downtime, it is recommended that you do not reboot the system until the RAID hardware has been installed and connected.

6. If you have not already installed the RAID hardware, refer to the appropriate integration chapter in this manual and do so at this time. To reduce system downtime, it is recommended that you do not reboot the system until the RAID hardware is installed and connected.

7. When the ptx/HITACHI-RAIDMP software and the RAID hardware have been installed, reboot the host system.

8. Refer to the appropriate integration chapter to verify and complete the integration.

ptx/HITACHI-RAIDMP Upgrade

If you are upgrading ptx/HITACHI-RAIDMP from an earlier version, do the following.

1. Before installing or de-installing any software, go to the file */usr/conf/uts/io/sd/sd_space.c* and find the line that reads:

```
char *sd_raidmp_license_key = "<license key string>";
```

The string *<license key string>* between quote marks is the software license key. Copy down this character string. You will need to re-enter it when you install the new version of ptx/HITACHI-RAIDMP.

2. De-install ptx/HITACHI-RAIDMP.

3. Refer to the discussion titled *Requirements for Use*, earlier in this chapter, and install any required fastpatches.
NOTE: If RAID400 or DF500 devices are to be connected, install fastpatch FP# 254922 at this time. You must install this fastpatch *before* installing the ptx/HITACHI-RAIDMP software.

4. Refer to the DYNIX/ptx V4.4 software installation documentation and install the software. This should be a standard CD-ROM installation, with the specifications listed under *First-Time Installation*, above.

CAUTIONS

***Do not** recompile the kernel at this time. The recompile should take place **only** after the license key has been entered, as described below.*

Be sure to save a hard copy of the ptx/HITACHI-RAIDMP license key. If you ever need to re-install or upgrade ptx/HITACHI-RAIDMP, you will need to re-enter the key.

5. Using the license key that you copied from the old version of software, do the following:

a. Change to the directory */usr/conf/uts/io/sd*, and open the file *sd_space.c* for editing.

b. In the file *sd_space.c*, find the line that reads:

```
char *sd_raidmp_license_key = "No LicenseKey";
```

c. Edit this line, replacing the string `No LicenseKey` with the ptx/HITACHI-RAIDMP license key. *Be sure that the key is enclosed in quotes.*

d. Save and exit the file *sd_space.c*.

e. Save the hard-copy form of the license key in a known location, in case you should need to re-install or upgrade the software.

6. Compile the new kernel.

7. If this upgrade involves any hardware changes (e.g., additional hardware, replacements, etc.), and if these changes have not yet been made, refer to the appropriate integration chapter in this manual and make the hardware changes at this time. To reduce system downtime, we recommend that you do not reboot until the RAID hardware changes have been made.

8. When the ptx/HITACHI-RAIDMP software and RAID hardware have been installed, reboot the host system.

9. Refer to the appropriate integration chapter to verify and complete the integration.

Alt-Disk O/S Upgrade

If you upgrade the DYNIX/ptx operating system using the "Alternate Disk" (or "Alt-Disk") Installation option in ptx/ADMIN, be sure to do the following steps in the order shown here. ***Performing these steps, in the proper order, is critical for a successful installation.***

1. Go to the file */usr/conf/uts/io/sd/sd_space.c* and find the line that reads:

```
char *sd_raidmp_license_key = "<license key string>";
```

Copy down the license key and save it. You will need to re-enter the key when you re-install ptx/HITACHI-RAIDMP.

2. De-install ptx/HITACHI-RAIDMP.

3. Perform the Alt-Disk operating system upgrade, then reboot the new operating system.

4. Refer to *Requirements for Use*, earlier in this chapter, and install any required fastpatches.

5. Re-install ptx/HITACHI-RAIDMP, using the *<license key string>* string you copied in Step 1.

6. Compile the new kernel.

7. If this upgrade includes any RAID hardware changes, make those changes at this time.

> **NOTE**
> *To reduce system downtime, it is recommended that you do not reboot the system until the RAID hardware is installed and connected.*

8. Refer to the appropriate integration chapter to verify and complete the integration.

Notes on Use and Administration

Using ptx/HITACHI-RAIDMP with ptx/SVM

If you are using the ptx/HITACHI-RAIDMP software with ptx/SVM, you should *not* use **newfs** on volumes that are under SVM control. Instead, use **mkfs** to reduce the number of copies of superblocks, thereby increasing the usable disk space and improving performance. On RAID-5 partitions *not* under SVM control, you can use **newfs**, as usual.

Missing, Invalid or Expired License Key Will Cause Errors

If the ptx/HITACHI-RAIDMP software license key has expired, you will see warning messages in **ktlog** similar to the following:

```
38c4a75a 22:53:14 warn q0/e0/p590 raidmp:2c:: Error <0xbad2> RAIDMP Driver
disabled

38c4a75a 22:53:14 tolog/warn q0/e0/p590 mpt_set_device_num: +sd11: set device
num callout failed
```

In addition, on the DYNIX/ptx system console, you will see **devctl** failures; e.g.:

```
devctl: name: couldn't find sd12

devctl: Assigned +sd11 --> sd12
```

as the kernel attempts to assign permanent names to the Hitachi RAID devices. This will cause **dumpconf** to report devices with temporary device names of the form:

```
+sd<device_number>
```

The above temporary devices are rendered useless.

If the ptx/HITACHI-RAIDMP software license key is incorrect or invalid, you will see warning messages similar to the following on the system console or in */usr/adm/messages*:

```
WARNING: raidmp:2b:: Error <0xbad1> RAIDMP Driver disabled

WARNING: sd(unit 16695296 on +fabric0): unable to get unit serial number
(custom)

WARNING: raidmp:2a:: Error <0xbad0> RAIDMP Driver disabled

WARNING: sd(unit 16695296 on +fabric0): unable to get unit serial number
(custom)
```

If any of the above symptoms occur, check that the software license key has been entered correctly and is still valid. If you do not have a license key or if it has expired, contact the IBM NUMA-Q Customer Support Center.

Replace VTOCs With Caution

It is possible to replace an existing VTOC with a different one (e.g., one built with **make_vtoc**) without rebooting the system or using **devctl -M** (**devctl -c** will leave DF350F routing set in a degraded mode). However, if you do this, the old "partition / slice" device nodes will still be present and *no new device nodes will be built*. Therefore, it is possible that some of the nodes required by the new VTOC will be missing. If this occurs, writes attempted to that partition will occur on the main hard drive as a regular file write rather than to the drive in the RAID array.

If you attempt to replace VTOCs without rebooting or using **devctl**, you will need to check that all necessary device nodes are present; if they are not present, you will need to build them manually with **mknod**(1M).

The safest method is to always reboot or use **devctl** when changing VTOCs.

Chapter 3

DF350F Subsystem Integration

This chapter describes the steps needed to integrate a Hitachi DF350F RAID disk array subsystem with a DYNIX/ptx host system running ptx/HITACHI-RAIDMP. The following topics are included:

- **Requirements for use.** Gives hardware requirements for the use of this subsystem with ptx/HITACHI-RAIDMP.
- **Restrictions.** Includes restrictions specific to the functioning and use of this subsystem with ptx/HITACHI-RAIDMP.
- **DF350F RAID concepts.** Gives a brief summary of the RAID terminology and concepts used in association with this subsystem.
- **DF350F subsystem integration overview.** Gives a high-level overview of the steps required to integrate this subsystem with the DYNIX/ptx host system.
- **Integration procedures.** Includes detailed procedures for ensuring that the subsystem has the required configuration settings and for installing and connecting the subsystem to the DYNIX/ptx host system.
- **Notes on Use and Administration.** Includes any notes, helpful hints, or other information that may be useful in the ongoing administration and use of the subsystem.

DF350F Requirements for Use

The following are required to integrate a DF350F subsystem with a DYNIX/ptx host system and this release of ptx/HITACHI-RAIDMP:

- The host system must conform to the requirements listed in the software chapter of this manual.
- The Hitachi DF350F RAID subsystem must meet the following requirements:
  - Fibre Channel interface cards installed
  - Configured for Fibre Channel connectivity
  - Microcode level 4356/M, minimum; contact the IBM NUMA-Q Customer Support Center for information on the latest supported microcode release(s).

DF350F Restrictions

The following restrictions apply to the integration and use of the DF350F subsystem with a DYNIX/ptx host system and ptx/HITACHI-RAIDMP. These restrictions are in addition to those listed in Chapter 2 of this document.

- Only switched fabric Fibre Channel connection of the DF350F subsystem is supported. SCSI connections are *not supported* by ptx/HITACHI-RAIDMP. SCSI connection of the *DF350* (same hardware chassis, same functionality as the DF350F, but configured for SCSI connection) is supported by the *ptx/HITACHI-RAID* software package.
- Only those Fibre Channel switches discussed in Chapter 1 of this manual are supported.
- DF350F disks cannot be used as boot or dump devices.
- The DF350F-MK Mini-Tower subsystem is not supported.
- The subsystem's RS-232 port is *not* an ASCII terminal port. It should be used *only* to access the subsystem using the Hitachi SOMMET software, which runs on a PC with MicroSoft Windows™.
- This release of ptx/HITACHI-RAIDMP supports only the "single controller" or "dual controller/Dual-Active mode" configurations.
- Hot standby mode is not supported.
- When using the subsystem with ptx/Clusters, "Data Share" mode *must be* ON.

DF350F RAID Concepts

This section discusses basic RAID concepts as implemented in the Hitachi DF350F subsystem. An understanding of Hitachi's implementation of these concepts is important to understanding subsystem configuration and status, and error messages reported by the subsystem.

> **NOTE**
> *This discussion assumes that the reader is familiar with general RAID concepts and terminology. It covers **only** how RAID is implemented in the Hitachi subsystem.*

RAID Groups

The DF350F subsystem uses the following terms to determine and identify configuration and status and to report error messages. Figure 3-1 shows a representation of these concepts:

Port	A *port* can essentially be identified with an internal SCSI bus.
Row	A *row* consists of five disks, each on a different port (internal SCSI bus), but each having the same SCSI ID.

In this basic configuration, a *port* can be equated with one of the subsystem's internal SCSI buses; there is one port for each internal SCSI bus. A *row* consists of five disks with the same SCSI IDs, with each disk on a different bus. Thus, Row 0 is composed of the disks on Ports 0—4 with SCSI ID 0. **Note** that the "hot spare" disk is on Port 4, Row 4.

The DF350F has two controllers, each with an external fibre channel interface for connection to a fibre channel switch or to the host system. Each of these controllers has five *internal* SCSI controllers that handle I/O on the subsystem's five internal SCSI buses.

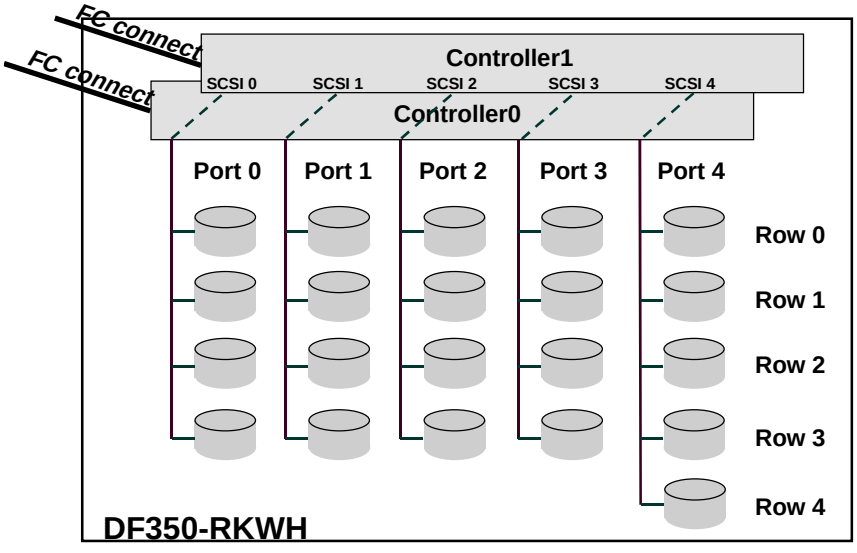

Figure 3-1.
Basic Subsystem Structure and Disk Layout.

A *RAID group* consists of one or more rows of disks, grouped together with the same RAID level. Figure 3-2 shows a typical RAID-0 or RAID-5 group. In this example, the entire row of five disks is used, although other configurations are possible. ***Note*** that different RAID groups within the same DF350F chassis can have different RAID levels.

RAID-5. For RAID-5, parity data is distributed across all five disks, each of which resides on its own internal SCSI bus. Also, for RAID-5, the spare disk on Port 4, Row 4 can be used to automatically assume the functions of a failed disk in any of the RAID-5 groups.

RAID-0. RAID-0 on the DF350F is not really striped, but is concatenated.

Figure 3-3 shows typical RAID-1 and RAID 0/1 (RAID-A) groups.

RAID-1. In this scheme, within each row:

The drive on Port 0 is mirrored with the drive on Port 1.
The drive on Port 2 is mirrored with the drive on Port 3.

Note the following:

- RAID-1 is not striped.
- This implementation of RAID-1 leaves one *remainder* disk in each row, the disk on Port 4, unused.

RAID-0/1. RAID-0/1 (also known as RAID-A) is mirrored in the same fashion as RAID-1 (Figure 3-3). Ports 0 and 2 (D and D') are a striped pair, and Ports 1 and 3 (M and M') are a striped pair.

Multiple rows, or all rows, can be used to make up a larger RAID group (Figure 3-4).

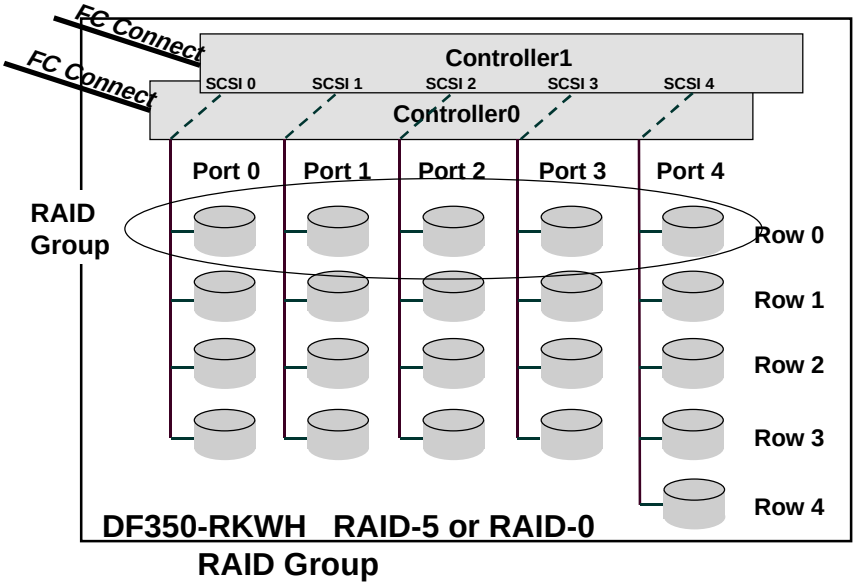

Figure 3-2.
RAID-0 or RAID-5 Group.

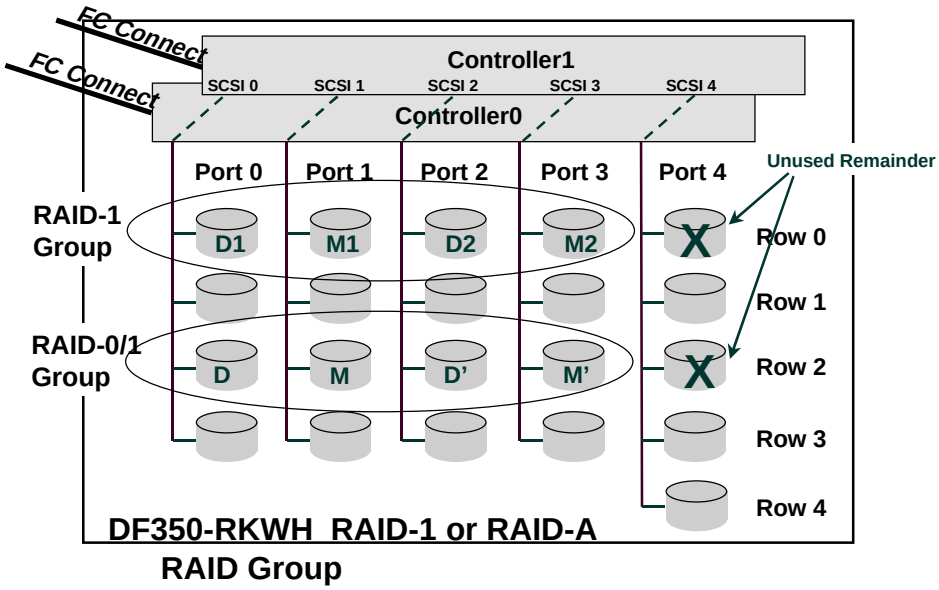

Figure 3-3.
RAID-1 and RAID-0/1 (RAID-A) Groups.

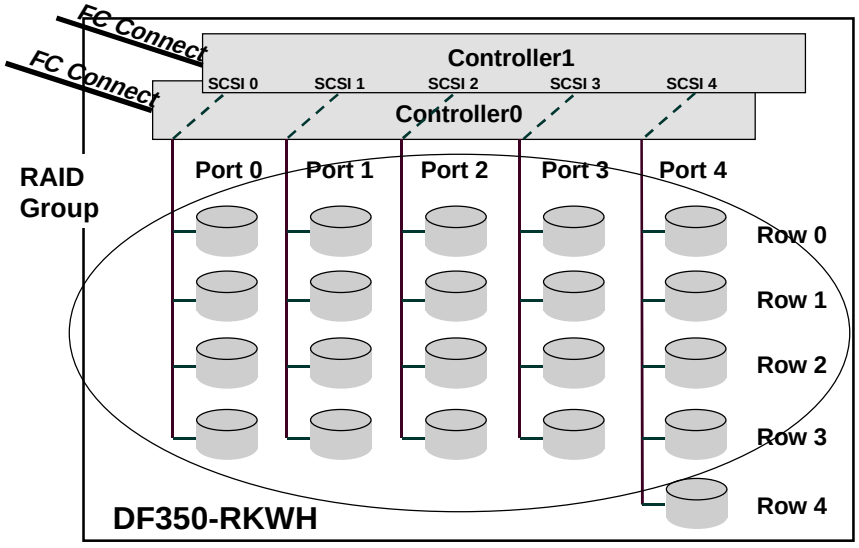

Figure 3-4.
Larger RAID Group.

Logical Units (LUs)

NOTE

Defining multiple LUs for a single RAID group requires knowledge of low-level technical details about disk geometry (e.g., block sizes, first and last block numbers, etc.).

With ptx/SVM, you can accomplish the same logical groupings. ptx/SVM offers custom VTOCs, and does not require you to know technical details of disk geometry. Therefore, if you are not familiar with this level of technical detail, you may find ptx/SVM easier and faster to use than defining multiple LUs for a RAID group.

A RAID group can be divided into more than one *logical unit (LU)*. Each logical unit is addressed by its *logical unit number (LUN)*. This is illustrated in Figure 3-5. A maximum of 8 LUs per DF350F subsystem are supported.

Note that all LUs within a RAID group will be configured at the same RAID level. E.g., in a RAID-5 group, all LUs within the group will be configured as RAID-5.

For RAID-5, LU0 data is striped across all five disks, using the specified stripe width of the DF350F. The stripe width must be set *before* defining the RAID groups. The default width is 64K.

For RAID-0/1 (RAID "A"), LU0 data is striped across two disks using the specified stripe width.

For RAID-0 and RAID-1, data is *not striped*.

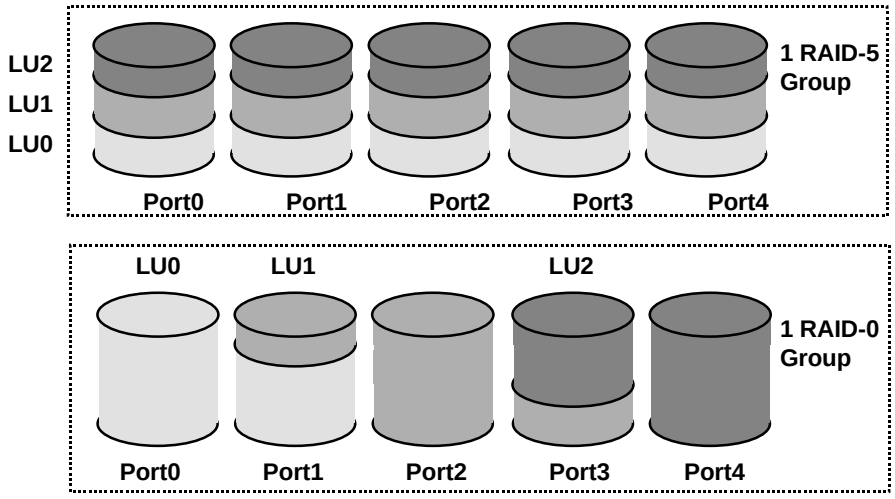

Figure 3-5.
Logical Units Within a RAID Group.

Multiporting on the DF350F Subsystem

Figure 3-6 shows a typical DF350F multiport configuration. In this configuration, each LU is "owned" by a single controller. A common configuration is for two LUs to be owned by each controller, with only one I/O route "active" for each LU at a given time. Ownership of each LU is defined at the time the LUs are configured, during hardware setup. Ownership of an LU can be changed via the front panel SVP or with Hitachi's SOMMET software.

The DYNIX/ptx multiport (**mpt**) driver is able to identify "Primary" and "Secondary" routes from DF350F Inquiry data. It is critical that the DF350F have the correct configuration settings (see the configuration settings discussion, later in this chapter). If these settings are not correct, either the devices will not be identified and configured correctly by DYNIX/ptx, or performance may be degraded.

- **Configuration for DF350F Multiporting**
  - **All LUs are "owned" by only one controller**
  - **The most common configuration for a DF350-RKHW is 2 LUs owned by each controller**
  - **Only one I/O route will be "Active" for an LU at a time**

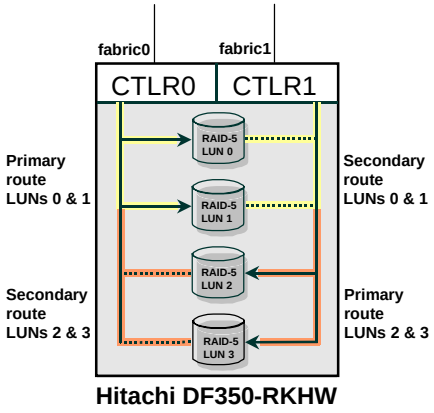

Figure 3-6.
Multiport Configuration.

In the event of a route failure (e.g., controller failure, bad cable, bad switch, etc.), the driver will automatically reroute I/O to the "Secondary" route (as in Figure 3-7). Because the secondary route's controller is not the "owning" controller, performance will be degraded but will at least allow I/O to continue until the Primary route is restored and/or the cause of the failure is resolved. Once the failure is resolved, the driver will automatically switch back to the Primary route.

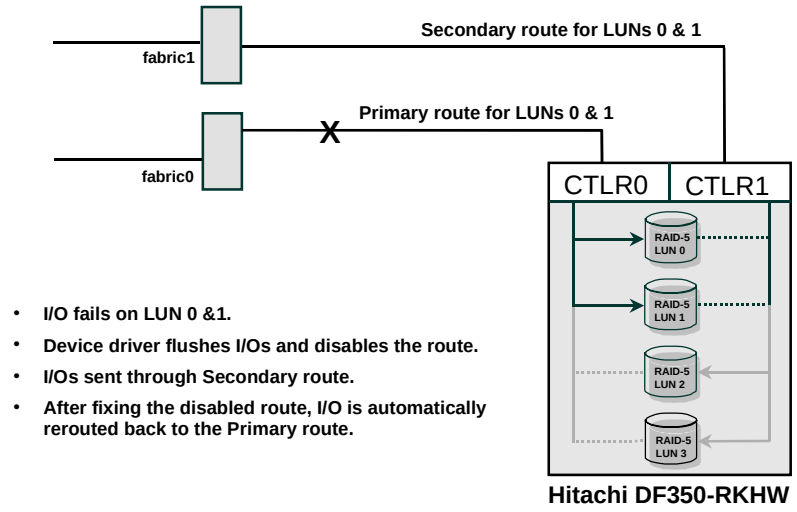

Figure 3-7.
Route Failure Example.

DF350F Integration Overview

To integrate the DF350F subsystem with a DYNIX/ptx host system, you must complete the following major steps:

- Step A: Install the ptx/HITACHI-RAIDMP software.
- Step B: Prepare the RAID hardware for integration.
- Step C: For rack-mount models to be mounted in an IBM NUMA-Q expansion cabinet, install mounting rails and rack-mount the subsystem.
- Step D: Check or adjust the subsystem's configuration settings.
- Step E: Make cable connections between the DYNIX/ptx host system and the subsystem.
- Step F: Verify the integration.

The following discussions cover these steps in detail or refer you to other documentation for details.

Step A: Install ptx/HITACHI-RAIDMP

At this time, refer to Chapter 2 of this manual and install the ptx/HITACHI-RAIDMP software.

Step B: Prepare for Integration

> **CAUTION**
> *Each of the supported rack-mount subsystems is too heavy for one person to move or lift. Obtain mechanized lifting equipment and/or additional persons before attempting to move or lift the subsystem.*
>
> *The DF350F 10-drive subsystem requires **at least two (2) people** to lift safely.*
> *The DF350F 20-drive subsystem requires **at least four (4) people** to lift safely.*

To unpack and locate the subsystem, refer to the subsystem documentation and do the following:

1. Unpack the unit.
2. Immediately after the unit is unpacked, inspect the unit for signs of shipping damage. If the unit appears damaged, contact the carrier in charge of shipping the unit and the vendor from which the unit was shipped.
3. Make sure that all parts and subassemblies were received.
4. Set aside and save the packing carton and all packing materials, in case the unit should ever need to be repackaged.
5. If the unit appears undamaged, locate the unit where it will be used.
6. Collect all cables and terminators, etc. that will be required.
7. Collect all required documentation.

Step C: Rack Mount the Subsystem, If Required

NOTE

This discussion tells how to rack mount the DF350F in an IBM NUMA-Q expansion cabinet. If you are installing the subsystem in an IBM NUMA-Q cabinet, complete the relevant procedures given in this section. If you are installing the subsystem in a vendor cabinet, consult the vendor documentation shipped with the subsystem and cabinet.

Most of the procedures required to rack-mount the subsystem are generic to the rack-mount models, and unless otherwise marked, a given procedure applies to *all* rack-mount models. However, certain of the following procedures are model-specific; these are clearly marked.

To rack-mount a subsystem in an IBM NUMA-Q expansion cabinet, you must:

- Install mounting rails into the expansion cabinet. The DF350F rack-mount subsystems require special rail adapters for mounting.
- Install the subsystem in the expansion cabinet.

Install Mounting Rails in the Expansion Cabinet

The rack-mount subsystems use the standard IBM NUMA-Q mounting rails, but require a special rail adapter to be installed on each rail. This discussion tells how to install the rail adapters and how to install the modified mounting rails.

1. Unpack the Hitachi Rackmount Kit (P / N 1003-75181). This kit should contain the following items:

- Standard IBM NUMA-Q rack mounting rails (qty 2)
- Assorted nuts, bolts, and screws for mounting, including four size-10 flat-head screws
- Hitachi guide-rail adapters (qty 2)

2. Modify each of the standard mounting rails by installing one guide-rail adapter onto the mounting rail. Figure 3-8 shows to install the guide-rail adapter on the left-side rail; installation on the right-side rail is similar. To do this:

a. Note that the left- and right-side mounting rails are *not* symmetrical.

b. Using Figure 3-8 as a guide, install the correct guide-rail adapter on the left-side mounting rail. Position the adapter along the *inside bottom* flange of the mounting rail, so that the two mounting holes in the adapter line up with the two screw holes in the lower portion of the side of the mounting rail.

c. Use two (2) flat-head screws per rail to attach the adapter to the mounting rail.

d. Repeat Steps **b** and **c** for the right-side mounting rail.

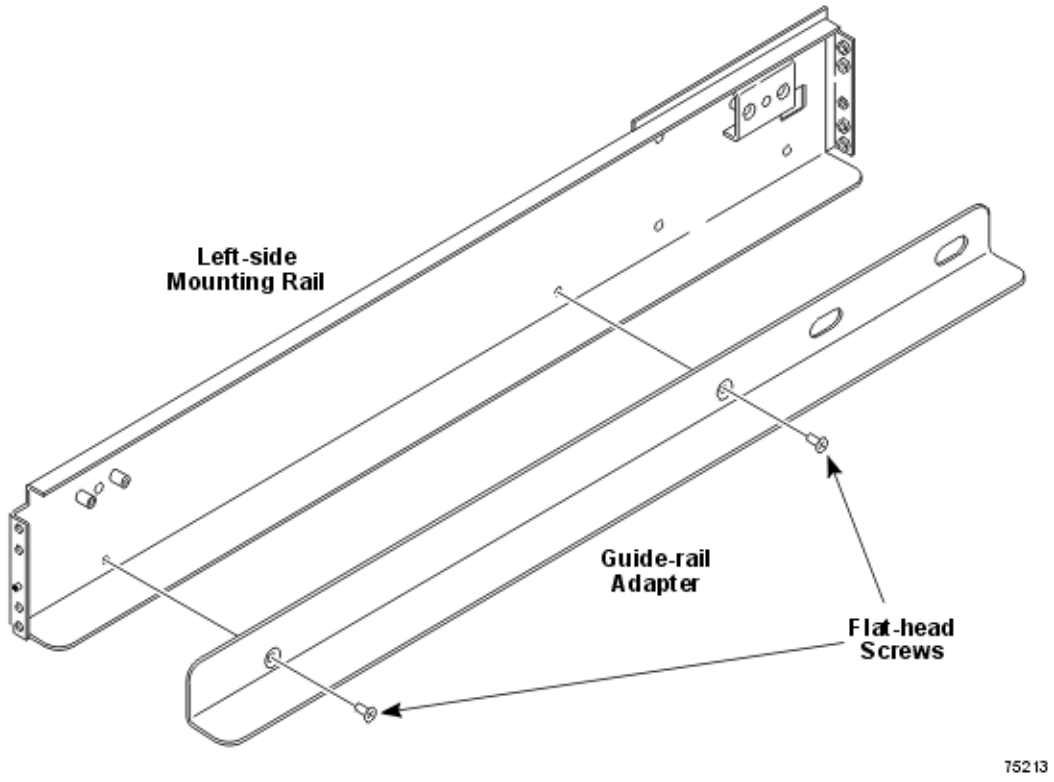

Figure 3-8.
Installing the Guide-Rail Adapter (Left-Side Mounting Rail).

3. Remove the front and rear doors from the expansion cabinet in which the rails will be installed.

4. Refer to Figure 3-9 (for 10-drive units) or Figure 3-10 (for 20-drive units) and install the left-side and right-side mounting rails at the appropriate location in the expansion cabinet. **Note** that each mounting rail is spring-loaded, with tabs to hold the rail in place until it can be secured to the cabinet's vertical rails with screws. Each end of the rail must be attached to one of the cabinet's vertical rails by *two (2)* screws. These screws should be installed as follows.

NOTE
On some models, the hole for the 2nd screw on the front of the rail may not be present.

- Front end of rail Hole locations 1, 2
- Back end of rail Hole locations 1, 5

Note that the fifth (top) hole location in the front of the mounting rail will be used to attach the subsystem to the cabinet's vertical rails.

5. Install a single clip nut onto the chassis. The clip nut location is model-dependent; hole locations are counted up, starting with the lowest mounting rail hole as location 1:

DF350F-RKH	Hole location 11
DF350F-RKWH	Hole location 23

This completes the procedure to install the mounting rails in the expansion cabinet. You are now ready to install the subsystem in the expansion cabinet.

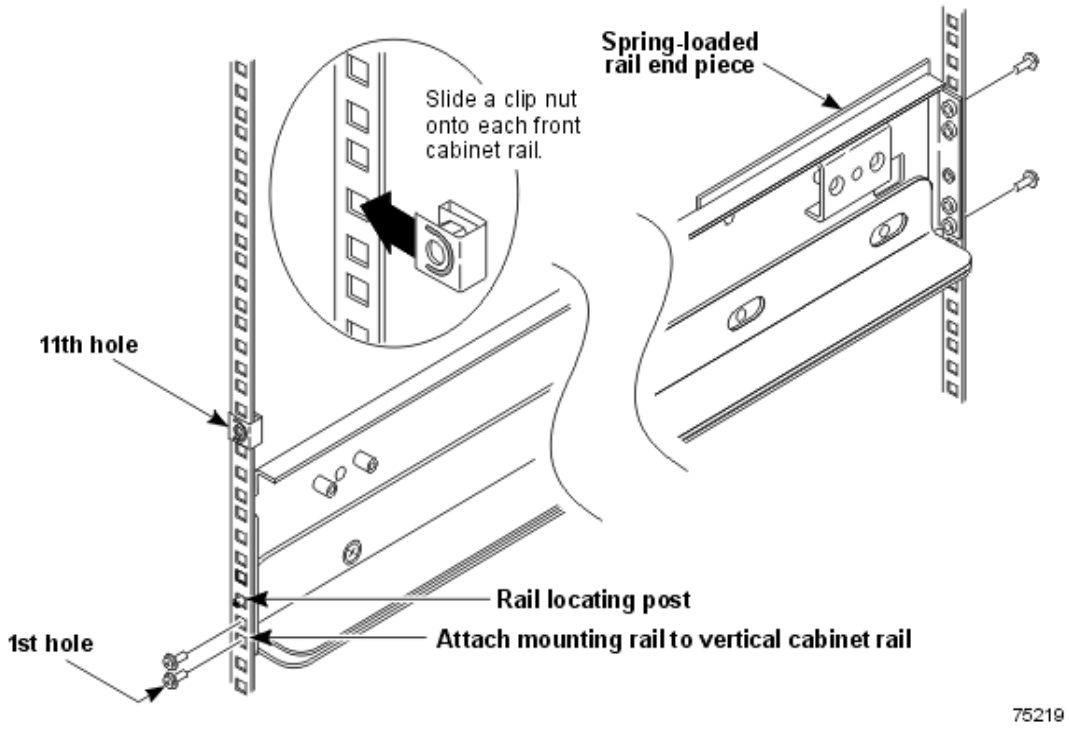

Figure 3-9.
Installing the Mounting Rails for a 10-Drive Unit.

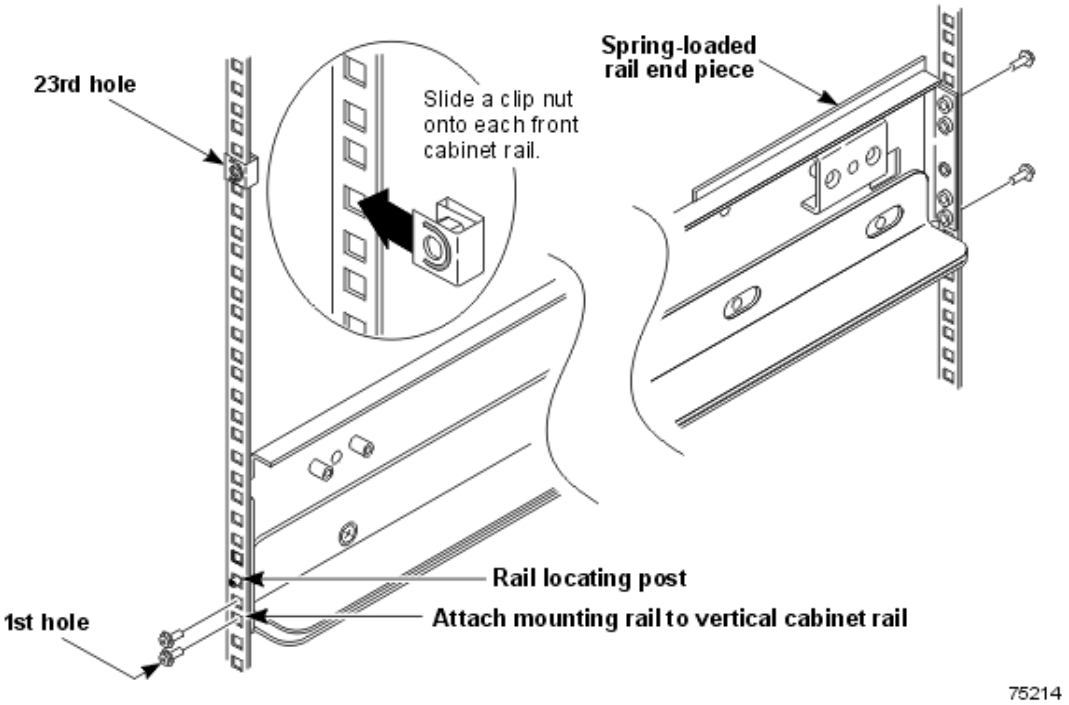

Figure 3-10.
Installing the Mounting Rails for a 20-Drive Unit.

Install the 10-Drive Subsystem in the Expansion Cabinet

CAUTION

The subsystem is too heavy for one person to move or lift. Attempting to move or lift the subsystem without adequate help could result in serious personal injury. Make sure that you have mechanized lifting equipment and/or enough persons to lift and maneuver the subsystem safely.

For the 10-drive subsystem, you should have at least two (2) persons to lift the subsystem. You should have one additional person to guide the subsystem into the cabinet.

To install the 10-drive subsystem in the expansion cabinet, do the following:

1. Make sure that the subsystem has been unpacked and that all packing material, parts containers, etc. have been removed from around the subsystem.

2. The subsystem's front bezel is held in place by metal friction tabs. Remove the front bezel by pulling it gently but firmly away from the subsystem frame until the friction tabs release their hold. It may be best to pull one side of the bezel away, then pull the other side away. Set the bezel aside for later re-installation.

3. Have ***at least*** two (2) persons lift the subsystem and slide it into the expansion cabinet and onto the mounting rails installed earlier. An additional person should be present to guide the subsystem into the expansion cabinet and onto the mounting rails.

4. Push the subsystem completely into the expansion cabinet, until the mounting flanges on the front of the subsystem are flush against the expansion cabinet's vertical rails.

5. Refer to Figure 3-11 and use the rail mounting screws supplied with the rack-mount kit to attach the subsystem to the expansion cabinet's vertical rails.

6. Re-attach the subsystem's front bezel to the subsystem by placing it against the front of the subsystem and pressing in toward the subsystem until *all four corners* of the bezel lock into place. *Do not* strike or shove the bezel into place, as this could damage sensitive internal components of the subsystem.

7. This completes the rack-mount installation procedure for the 10-drive subsystem.

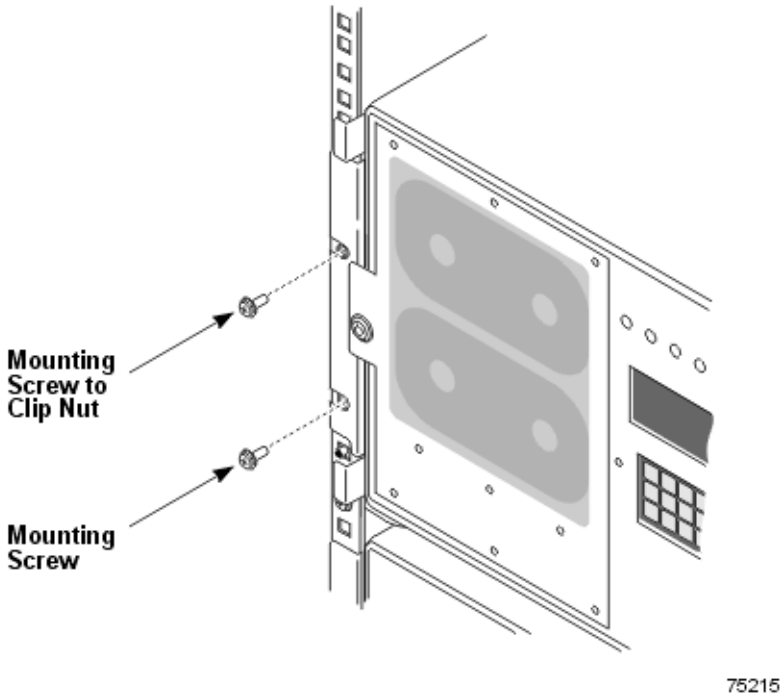

Figure 3-11.
Installing the 10-Drive Subsystem.

Install the 20-Drive Subsystem in the Expansion Cabinet

CAUTION

The subsystem is too heavy for one person to move or lift. Attempting to move or lift the subsystem without adequate help could result in serious personal injury. Make sure that you have mechanized lifting equipment and/or enough persons to lift and maneuver the subsystem safely.

For the 20-drive subsystem, you should have at least four (4) persons to lift the subsystem. You should have one additional person to guide the subsystem into the cabinet.

To install the 20-drive subsystem in the expansion cabinet, do the following:

1. Make sure that the subsystem has been unpacked and that all packing material, parts containers, etc. have been removed from around the subsystem.

2. The subsystem's front bezel is held in place by metal friction tabs. Remove the front bezel by pulling it gently but firmly away from the subsystem frame until the friction tabs release their hold. It may be best to pull one side of the bezel away, then pull the other side away. Set the bezel aside for later re-installation.

3. Detach the nut plate hardware from each of the subsystem's mounting flanges. Retain the nut plates, *but discard the nuts and bolts used to attach the nut plates to the subsystem*; you will use the screws supplied with the rack-mount kit to attach the nut plates and subsystem to the cabinet.

4. Have ***at least*** four (4) persons lift the subsystem and slide it into the expansion cabinet and onto the mounting rails installed earlier. An additional person should be present to guide the subsystem into the expansion cabinet and onto the mounting rails.

5. Push the subsystem completely into the expansion cabinet, until the mounting flanges on the front of the subsystem are flush against the expansion cabinet's vertical rails.

6. Refer to Figure 3-12 and note the position of the nut plate, subsystem mounting flange, and expansion cabinet's vertical rail. Note the following:

- Each nut plate is mounted with the cutouts for the friction tabs *toward the interior of the expansion cabinet*; the friction tabs face *away from the interior of the expansion cabinet*.
- Although there are four holes in the nut plate, the nut plate and subsystem chassis are attached using *two (2)* rail mounting screws:

  The lower screw goes into the *fifth (5th) hole*, counting up from the bottom of the mounting rail.

  The upper screw goes into the hole location of the nut clip you attached earlier. Make sure that the nut clip is still in its proper position; this will be hole 23.

7. Using Figure 3-12 as a guide, attach the nut plates and subsystem to the expansion cabinet's vertical rails using the screws included in the rack-mount kit.

8. Re-attach the subsystem's front bezel to the subsystem by placing it against the front of the subsystem and pressing in toward the subsystem until *all four corners* of the bezel lock into place. *Do not* strike or shove the bezel into place, as this could damage sensitive internal components of the subsystem.

This completes the rack-mount installation procedure for the 20-drive subsystem.

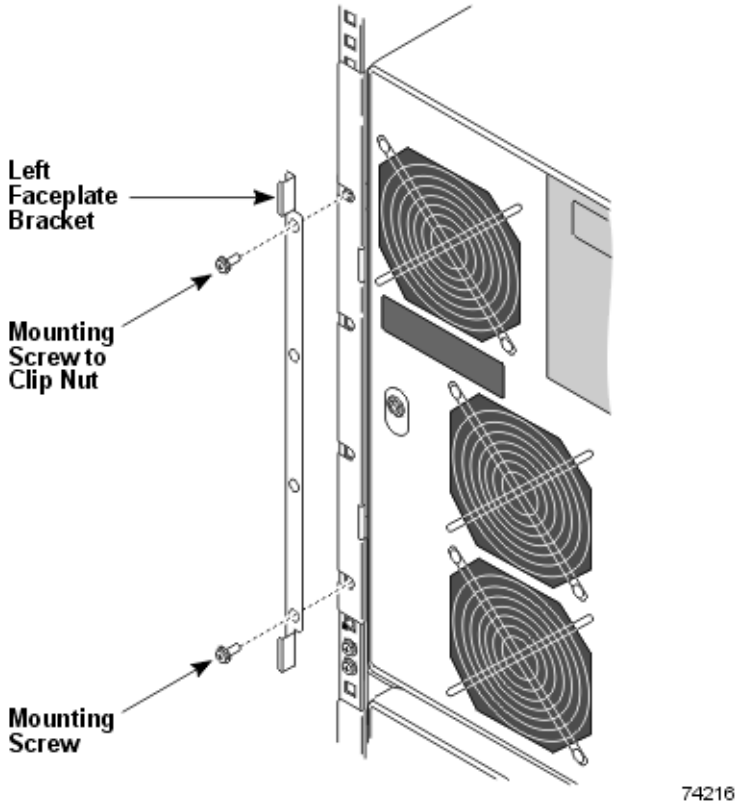

Figure 3-12.
Rack-Mounting the 20-Drive Subsystem.

Step D: Check and Adjust Configuration Settings

CAUTIONS

The subsystem has two power switches: a DC switch on the front panel, and an AC switch on the rear panel. ***Always turn off the front panel DC switch first!*** *Turning off the rear panel AC power switch first may remove power from the subsystem's backup battery.*

Use caution when setting the subsystem's DIP switches. One of these switches (switch #3) will erase the subsystem's NVRAM!

Use extreme caution when changing the subsystem's RAID configuration. Do not attempt to do this unless you are thoroughly familiar with RAID concepts as implemented in this subsystem. Errors in configuration can cause data corruption or loss. For example, removing a logical unit from the configuration causes all data on that logical unit to be lost!

Critical Settings

CAUTION

The following critical parameters ***must be set exactly as shown*** *for the DF350F subsystem and ptx/HITACHI-RAIDMP to work properly together.*

ROM Mode Parameters—Single Controller Operation

CONTROLLER NAME	ENABLE	Serial #-Controller combination
See *Entering the CONTROLLER NAME/ENABLE Value* discussion, below.		
SPECIAL CONNECT	ALTERNATE PATH3	
INQUIRY INF	VENDOR TYPE	HITACHI
	PRODUCT TYPE	DF350F
DUAL CONFIG	SINGLE MODE	

ROM Mode Parameters—Dual Controller Operation

CONTROLLER NAME	ENABLE	Serial #-Controller combination
See *Entering the CONTROLLER NAME/ENABLE Value* discussion, below.		
SPECIAL CONNECT	ALTERNATE PATH3	
INQUIRY INF	VENDOR TYPE	HITACHI
	PRODUCT TYPE	DF350F
DUAL CONFIG	DUAL MODE	DUAL-ACTIVE

System Settings for FC Multiport Operation (both single and dual controller)

SYSTEM OPT (SVD)	TRESPASS MODE	AUTO TRSPAS
TOPOLOGY (CTL0)	FABRIC P-TO-P	
TOPOLOGY (CTL1)	FABRIC P-TO-P	

NOTES

*ROM parameters can **only** be changed during initial boot of the DAS. DIP switch settings must be made on **both** controllers to boot into ROM mode.*

SPECIAL CONNECT - ALTERNATE PATH3 is required for both single and dual controllers.

DUAL CONFIG - DUAL MODE - DUAL-ACTIVE is required only for dual controller configurations.

DF350F Subsystem SVP Function and EEPROM Settings

The DYNIX/ptx system will expect the DF350F disk array subsystem to have the configuration settings listed in this section. Before attempting to connect the subsystem to the host system, check the subsystem's configuration settings to make sure that they match the values given here.

A single disk array subsystem may have either one or two RAID controllers installed. These controllers are identified as follows:

Controller 0 (CTL0)	First controller in a subsystem. If there is only one controller installed in a subsystem, it is identified as CTL0.
Controller 1 (CTL1)	Second controller in a subsystem.

Checking or Adjusting Settings

To check or adjust these settings, do the following:

1. Power off the disk array subsystem.
2. From the rear of the subsystem, locate the row of DIP switches on each controller board.
3. Set DIP switch 4 to **ON**. This will cause the subsystem to boot up in ROM mode.
4. Power up the subsystem.

NOTE

*If dual controllers are installed, you must check/adjust settings for each controller **separately**. On the display panel, the first character in the second line of the display indicates which controller is currently selected:*

*An **underscore** in the first field indicates Controller 0.*
*A **flashing block** in the first field indicates Controller 1.*

CAUTION

When checking or adjusting the settings on an existing DAS that contains data, be aware that changing the STRIPE SIZE and/or ROW/LAST LBA settings can cause loss of data.

5. Using the subsystem's keypad and LED display panel, check or adjust the settings ***for each controller*** to match those required for your microcode revision level, as shown in the appropriate table, below.
6. Power off the subsystem, then reset DIP switch 4 to OFF.

NOTE

*You will probably want to leave the subsystem powered off at this time. You **do not need** to power the subsystem back on until cable connections have been completed in Step E.*

This completes the procedure to check and adjust the subsystem's configuration settings.

Entering the CONTROLLER NAME/ENABLE Value

The CONTROLLER NAME/ENABLE menu item takes an *eight-character* value. **Note** that the field is 10 characters in length, but the last two characters (which will read C0 or C1, depending on the controller selected) cannot be changed by the user.

You must use the DAS's 5-digit "Mfg. No." (serial number, found on the subsystem's rear panel) and the controller number (00 or 01) to construct an 8-digit controller name. To do this, append "-00" or "-01", depending on the controller, to the Mfg. No.:

`<Mfg. No.>-00 or <Mfg. No.>-01`

For example, a unit with Mfg. No. 10298, the CONTROLLER NAME/ENABLE field for Controller 0 must read:

```
10298-00C0
```

The "Mfg. No." is the same for both controllers. Thus, if this is a dual controller subsystem, the CONTROLLER NAME/ENABLE field for Controller 1—the *second* controller—must read:

```
10298-01C1
```

NOTE

On the subsystem's keypad, numeric values must be entered digit by digit, as four-digit ASCII decimal equivalents (ADEs). The values you will need are:

Digit	***ADE***
0	*0300*
1	*0301*
...	*...*
9	*0309*
- (hyphen)	*0213*

Summary of DF350F EEPROM Settings

The following table summarizes the ROM mode settings. The various message levels refer to items on the DF350F's menu display. Note the following:

- A * indicates the factory default setting.
- A ** indicates a setting required for operation with ptx/HITACHI-RAIDMP that is *different than* the factory default setting.
- Recommended or required settings are shown in **bold**. Settings shown are for both single and dual controller operation, where these settings differ.

Table 3-1
DF350F EEPROM Settings—Microcode Revision 4356/M

Setting	Message 1	Message 2	Notes
WRITE & VERIFY	**ON*** OFF		
STRIPE SIZE	16KB 32KB **64KB***		
LAN CONST	CONNECT LAN	CONNECT **NOT CONNECT***	
	IP ADDRESS SUB NET MASK DEFAULT GATEWAY ETHER ADDRESS		
CONNECT VMS	CONNECT VMS **NOT CONNECT***		
SPECIAL CONNECT	OPEN SYSTEM **ALTERNATE PATH3**** (single or dual controller) TGET SHARE MODE		
HOST BLK SIZE	**512 BYTES*** 520 BYTES		
ERROR INF	ON (Normal) ON (HiTrack) **OFF***		
CONTROLLER	MINI TOWER RACK MOUNT HIGH RACKMOUNT **CABINET, WRACK***		
SPARE DISK	**EXISTENCE*** NOTHING		
SERIAL NO	**0000***		
ROM V/R	**0000***		
RAM V/R	**0000***		
CONTROLLER NAME	DISABLE* **ENABLE****	CTL NAME	Enter <serial#>-<CTLR#>. **Note** that last two characters in this field cannot be changed by the user.
ROW LAST LBA	ROW 0	**01046C97*** (8.8GB disks) **0215D9B3** (17.8GB disks)	
	ROW 1	**01046C97*** (8.8GB disks) **0215D9B3** (17.8GB disks)	
	ROW 2	**01046C97*** (8.8GB disks) **0215D9B3** (17.8GB disks)	
	ROW 3	**01046C97*** (8.8GB disks) **0215D9B3** (17.8GB disks)	

Table 3-1
DF350F EEPROM Settings—Microcode Revision 4356/M

Setting	Message 1	Message 2	Notes
BUZZER	**ENABLE*** DISABLE		
SYSTEM ERROR	SYSTEM DOWN **AUTO RESET***		
GENERATE SYS	**NEW SYSTEM*** OLD SYSTEM		
INQUIRY INF	VENDOR TYPE	**HITACHI***	
	PRODUCT TYPE	**DF350F***	
	COMMAND QUEUING	**ON***	
DUAL CONFIG	**DUAL MODE** (dual controller)	**DUAL TYPE***	**DUAL ACTIVE***
		DATA SHARE MODE	ON **OFF***
	SINGLE MODE (single controller)		
SYS REV	***x*356/M; *x*=0 or 4**		Value depends on the revision.
OPTION	**0000***		
BOOT PROG V/R	**FIXED:PE02*** **CHGABLE:PE29***		
RTC SET			Real-Time Clock
QUIT			

Step E: Make Cable Connections

1. Before connecting the subsystem to the host system, refer to the cabling diagrams on the following pages. Select the correct diagram for your installation.

2. Refer to the correct cabling diagram and connect the disk array subsystem(s) to the host system. Be sure to use the correct cables. See *Requirements for Use*, earlier in this chapter.

CAUTION

*The configurations shown in the following figures are the **only configurations supported in this release of ptx/HITACHI-RAIDMP**. Do not attempt to connect or use the subsystems in any other configuration.*

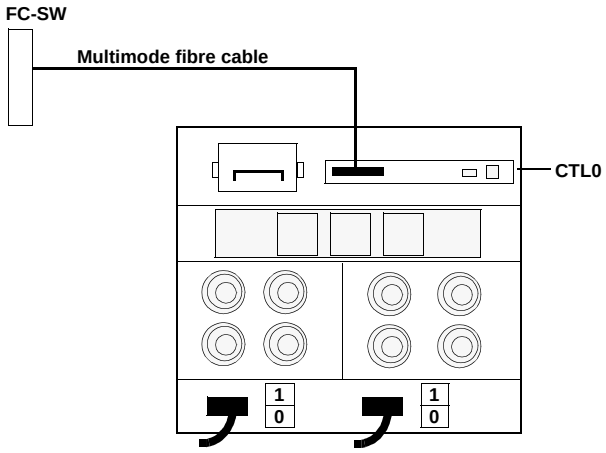

Figure 3-13.
Single Controller Connection.

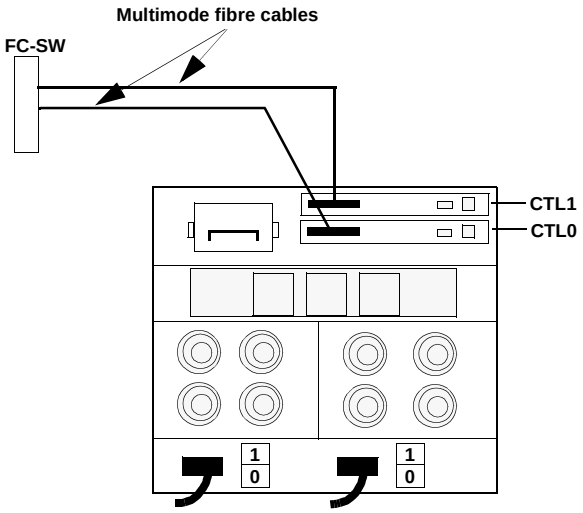

Figure 3-14.
Dual Controller Connection to One FC-SW Switch.

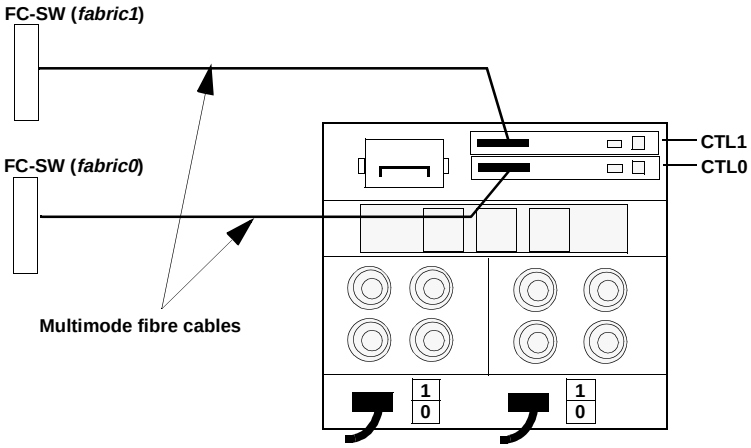

Figure 3-15.
Dual Controller Connection to Two FC-SW Switches.

Step F: Verify the Integration

To verify the installation and integration of the RAID subsystem:

1. Bring the RAID subsystem up. To do this:

a. Connect AC power to the subsystem, if this has not already been done.

b. Turn the subsystem's AC power switch on. This switch is located on the subsystem's rear panel.

c. Turn the subsystem's DC power switch on. This switch is located on the front panel.

2. Warn users that the system will be coming down.

3. Reboot the host system with the new kernel that you compiled when you installed the ptx/HITACHI-RAIDMP software.

4. When the system has booted, enter the following at the ptx command line:

/etc/dumpconf -d | grep sd | grep fabric

and verify that the appropriate **sd** devices were found at system boot. For single controller subsystems you should see one entry for each **sd** device; for dual controller subsystems you should see two entries for each **sd** device.

5. For each **sd** device, use the **diskid** and **infodev** commands to determine which **sd** devices are Hitachi RAID devices and fabric and path information. Figure 3-16 shows sample command outputs.

NOTE
*The **dumpconf** and **infodev** commands report unit information in hex; the **diskid** command reports the **cport** information in decimal. You will need to translate between the output of these commands.*

Refer to that Figure 3-16 and note the following.

dumpconf Output. In this example, there are four (4) **sd** devices, each with two (2) routes through fabric0 and fabric1.

In the fourth field of the **dumpconf** output, the first two hex digits show the Controller and Port IDs. For example, in the line for **sd4**:

```
sd4          sd            4  0x01fe1802  SM     fcp       fabric1
```

in the string 0x01f31802, the digits "0x01" indicate Controller 1, Port 0.

The last two digits of this hex string identify the LUN. In this example, the LUN is 02.

diskid Output. From the **diskid** output, you can determine the following:

- At least one of the fabrics to which this device is connected (in Figure 3-16, it is *fabric0*).
- The digit following "cport" identifies the controller port (here, port 0).

- In the fc id hex string, the third and fourth digits identify the Switch port ID (in the example in Figure 3-16, it is 03).
- In the fcp lun list, the first four digits identify the LUN (here, 03).
- The vendor information line identifies the unit serial number (here, 10213) and, again, the LUN (here, 3).

The **cport** number can be decoded to find the DF350F controller number and port number on that controller, as follows:

DF350F cport	CTL	CTL Port
0	0	0
1	1	0

NOTE: To find both fabrics, use **/etc/diskid** with no arguments. All **sd** devices will be displayed, with all routes configured.

infodev Output. You can also use **infodev** to determine the Switch port ID and the fabric.

```
/etc/dumpconf -d |grep sd|grep fabric

sd2         sd           2  0x00fe1800  SM     fcp      fabric0
sd3         sd           3  0x00fe1801  SM     fcp      fabric0
sd4         sd           4  0x00fe1802  SM     fcp      fabric0
sd5         sd           5  0x00fe1803  SM     fcp      fabric0
sd2         sd           2  0x01fe1800  SM     fcp      fabric1
sd3         sd           3  0x01fe1801  SM     fcp      fabric1
sd4         sd           4  0x01fe1802  SM     fcp      fabric1
sd5         sd           5  0x01fe1803  SM     fcp      fabric1
```

```
/etc/diskid sd5
sd5: fabric0 cport 0 fc id 0x200313, fcp lun 00:03:00:00:00:00:00:00
  No Pbay on fabric0
  Vendor HITACHI, product DF350F, revision 0000, serial # 10213-3
  capacity 68235 Mbytes
```

```
/etc/infodev -n sd5

Fibre Channel Address          : 0x200313 on fabric0 (unit 0xfe1803)
Fibre Channel Address          : 0x200313 on fabric1 (unit 0xfe1803)
```

Figure 3-16.
dumpconf, diskid and infodev Output for FCB SCSI Connection.

6. Perform I/O to the subsystem, and verify its operation.

7. Complete the preparation of the RAID subsystem for use. This may include building VTOCs and filesystems or configuring logical volumes using ptx/SVM. If you are using VTOCs, refer to the discussion titled *Notes on Usage and Administration*, at the end of this chapter, for details on the VTOCs supplied for DF350F devices.

8. Bring the host system back to multi-user mode and notify users that the system is now available for use.

Notes on Use and Administration

Sample VTOCs Supplied With ptx/HITACHI-RAIDMP

The ptx/HITACHI-RAIDMP distribution media includes several sample VTOC files. The characters “9gb” in the VTOC filename indicate the disk capacity.

hd0_5x9gb	For RAID-0 disk arrays
hd1_4x9gb	For RAID-1, RAID 0+1 disk arrays
h5_5x4gb	For RAID-5 disk arrays
hd5_5x9gb	For RAID-5 disk arrays

In addition, there are sample VTOCs to support ptx/SVM *sliced disks*:

hd0_5x9gb.vx	RAID-0, one-slice disk
hd1_4x9gb.vx	RAID-1 or RAID 0+1, one-slice disk
hd5_5x9gb.vx	RAID-5, one-slice disk
hd0_5x9gb.vxsl	RAID-0, many-slice disk
hd1_4x9gb.vxsl	RAID-1 or RAID 0+1, many-slice disk
hd5_5x9gb.vxsl	RAID-5, many-slice disk

Refer to the *ptx/SVM Administration Guide* for detailed information on ptx/SVM.

The make_vtoc(1M) Utility

The **make_vtoc**(1M) utility allows you to define custom VTOCs. Refer to the on-line man page or to the hard copy of this man page in Appendix A of this manual, for details. Refer to the DYNIX/ptx **mkvtoc** man page for additional details on VTOC creation.

The dfutil(1M) Utility

The **dfutil**(1M) utility allows the display and/or logging of subsystem status information for the DF350F subsystem. Refer to the online man page, or to the hard copy of the man page in Appendix A of this manual, for details.

Chapter 4

RAID200 Subsystem Integration

This chapter describes the steps needed to integrate a Hitachi RAID200 RAID disk array subsystem with a DYNIX/ptx host system running ptx/HITACHI-RAIDMP. The following topics are included:

- **Requirements for use.** Gives hardware requirements for the use of this subsystem with ptx/HITACHI-RAIDMP.
- **Restrictions.** Includes restrictions specific to the functioning and use of this subsystem with ptx/HITACHI-RAIDMP.
- **RAID200 RAID concepts.** Gives a brief summary of the RAID terminology and concepts used in association with this subsystem.
- **RAID200 subsystem integration overview.** Gives a high-level overview of the steps required to integrate this subsystem with the host system.
- **Integration procedures.** Includes detailed procedures for ensuring that the subsystem has the required configuration settings and for installing and connecting the subsystem to the host system.
- **Notes on Use and Administration.** Includes any notes, helpful hints, or other information that may be useful in the ongoing administration and use of the subsystem.

RAID200 Requirements for Use

The following are required to integrate a RAID200 subsystem with a DYNIX/ptx host system and this release of ptx/HITACHI-RAIDMP:

> **CAUTION**
> *This release of ptx/HITACHI-RAIDMP supports **only differential, FCB SCSI connection** of RAID200 subsystems. Switched fabric fibre channel connections are not supported.*

DYNIX/ptx System Requirements

The following are the DYNIX/ptx system requirements for integration of the RAID200 subsystem:

- The host system must conform to the requirements listed in the software chapter of this manual.
- The host system must have enough FCB SCSI channels available for connection of the RAID200 devices.

RAID200 Subsystem Requirements

Hitachi RAID200 subsystem (or distributor equivalent) must meet the following requirements:

- Must be configured for differential SCSI connectivity.
- Must be running microcode revisions 03-16-60-00/00, minimum; contact the IBM NUMA-Q Customer Support Center for information on the latest supported microcode releases.

SCSI Cabling and Termination Requirements

> **NOTE**
> *The Hitachi disk array subsystem's connectors use a smaller screw thread size (2-56) than the standard connectors (4-40) on DYNIX/ptx host systems. The cables and terminator listed below use the smaller, 2-56 thread size.*

- SCSI cables to connect the DYNIX/ptx system with the subsystem must be chosen from the following:

09P7837	15.0m cable
09P7839	7.5m cable
09P7799	3.0m cable

- For each SCSI termination on a disk array subsystem, the following SCSI terminator is required:

  09P7552

RAID200 Restrictions

The following restrictions apply to the integration and use of the RAID200 subsystem with this release of ptx/HITACHI-RAIDMP. These are in addition to any general or subsystem-specific caveats listed in Chapter 2 of this document.

General Subsystem Restrictions

- This release of ptx/HITACHI-RAIDMP supports ***only*** Fibre Channel Bridge (FCB) SCSI connections to the RAID200 devices. Switched fabric fibre channel connections *are not supported*.
- Single-port and multiport configurations are supported. Up to two (2) ports are supported.
- RAID200 disks cannot be used as boot or dump devices.

SCSI Connectivity Restriction

- Only differential SCSI connections are supported. Single-ended SCSI connections are not supported.
- External SCSI bus cable length is limited to 20m.
- Although multi-initiator SCSI is allowed with the RAID200, it is not recommended for use. For path redundancy, it is recommended that the multiport capabilities of the RAID200 be used, along with FC (fibre channel) multi-pathing with fibre channel switches and bridges.
- Fibre channel bridges with Hitachi RAID devices connected to them must be *dedicated* to the Hitachi RAID devices. No other types of devices can be connected to that bridge.
- The IBM NUMA-Q Fibre Channel Bridge (FCB) does not support Ultra-Wide SCSI protocol. However, it negotiates SCSI-2 protocol.

RAID200 RAID Concepts

The following discussions give an overview of RAID200 terminology and concepts.

Terminology

The following terms are used in the remainder of this document to describe the Hitachi RAID200 subsystem and software.

Table 4-1
RAID200 Terminology

CHA	Channel Adapter. The MVS (ESCON) connection.
CHS	Channel Adapter for SCSI. Each CHS supports four (4) SCSI buses.
DKA	Disk Adapter. The interface between DKC cache and the internal SCSI disk drives.
DKC	Disk Controller. The main chassis that houses the CHA, CHS, cache, and SVP laptop.
DKU	Disk Unit. A frame that houses the SCSI disk drives.
frame	Physical chassis that houses one or more hardware components.
LDEV	Logical Device. This is either a 3390-3 or 3390-9 compatible volume, in terms of capacity. Each LDEV is assigned and addressable as a SCSI target ID/LUN.
SCP	SCSI Control Processor. A SCSI interface controller on a Channel Adapter for SCSI (CHS). There are four (4) SCPs on each CHS.
SCSI Multi-port	The Hitachi RAID subsystem provides the capability to allow I/O through multiple SCSI port connections.
SVP laptop	A built-in laptop computer that enables configuration and monitoring of the subsystem.

Multiported Devices

Certain disk storage units provide multiple physical connections (ports) to multiple I/O buses, and then configure their logical volumes to allow access from those ports simultaneously. This feature is called *multiporting*.

Multiporting provides redundant access to individual logical volumes, increasing their availability. If a particular port or I/O bus becomes unavailable, an unaffected port and/or I/O bus can be used to access the disk. Locating multiple ports for a logical volumes on separate buses provides protection from both an I/O bus failure and the failure of a port within a storage unit. Locating multiple ports for a logical volume on the same I/O bus is legal, but this method does not provide protection from the failure of that I/O bus. With multiporting:

- A port can have more than one logical volume connected to it.
- A single logical volume can be connected to more than one port. E.g., in Figure 4-1, the disk resource is connected to Port 0 and Port 1, and has a path through both fibre channel bridges *Bridge 0* and *Bridge 1*.

The operating system also takes advantage of multiporting to better distribute its I/O load. Doing so *may* reduce the latency of individual I/O requests (although this is not guaranteed).

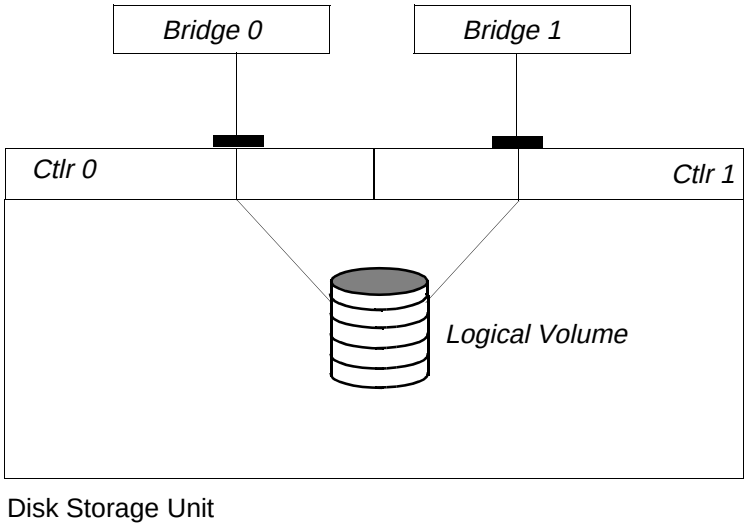

Figure 4-1.
Sample Multiported Disk Storage Unit.

RAID200 Integration Overview

To integrate the RAID200 subsystem with the DYNIX/ptx system, you must complete the following steps.

- Step A: Install the ptx/HITACHI-RAIDMP software.
- Step B: Prepare the RAID hardware for integration.
- Step C: Check or adjust the subsystem's configuration settings.
- Step D: Make cable connections between the host system and the subsystem.
- Step E: Verify the integration.

The following discussions cover these steps in detail or refer you to other documentation for details.

Step A: Install ptx/HITACHI-RAIDMP

To install the ptx/HITACHI-RAIDMP software, refer to Chapter 2 of this document. After the software has been installed, return here to continue the integration process at Step B.

Step B: Prepare for Integration

To prepare the subsystem for integration, do the following:

1. Unpack the unit.

2. Immediately after the unit is unpacked, inspect the unit for signs of shipping damage. If the unit appears damaged, contact the carrier in charge of shipping the unit and the vendor from which the unit was shipped.

3. Make sure that all parts and subassemblies were received.

4. Set aside and save the packing carton and all packing materials, in case the unit should ever need to be repackaged.

5. If the unit appears undamaged, locate the unit where it will be used.

6. Collect all cables and terminators, etc. that will be required.

7. Collect all required documentation.

Step C: Configure the RAID200 Subsystem

CAUTION

The following procedures assume that service personnel for Hitachi or a Hitachi distributor have first configured and formatted RAID groups on the subsystem. This section only defines how to configure paths.

CAUTION

Do not attempt to change the subsystem's configuration while user activity is occurring.

CAUTION

Incorrect configuration of the hardware could lead to data corruption or loss. ***Before any hardware configuration*** *is done, the customer and all personnel involved in the installation should carefully define what the RAID configuration for this site will be. The following procedure covers* ***only*** *the use of the Online Maintenance menu system to perform the actual configuration tasks, and assumes that configuration planning has already been completed.*

This section describes the steps necessary to configure a RAID200 subsystem for SCSI connection to a DYNIX/ptx system.

Getting to the Configuration Screens

To change the hardware configuration settings, use the SVP program on the built-in laptop computer attached to the subsystem.

1. Connect the subsystem to an appropriate AC power source, and bring the subsystem up.
2. On the laptop, bring up the SVP program.
3. From the SVP main menu bar, double-click the `ONLINE` icon.
4. From the resulting display, double-click the `MAINTENANCE` icon. When the subsystem has finished reading current configuration data (this will take several seconds), the `Maintenance Online` window and menu bar will be displayed.

5. Pull down the `Mode` menu and select `Modify`.
6. Pull down the `Maintenance` menu and select `Change Subsystem`. An additional menu will pop up; from that menu, select `Device Structures`.

   The SVP program will again read subsystem configuration data; this will take several seconds.

7. Now, several windows appear in succession.

a. At the `Device Structure Setup` window, select `OK`.

b. At the `Channel Configuration` window, select `OK`.

c. At the `SCSI PCB Detail` window, select `OK`.

d. At the `Channel Configuration` window, select `OK`.

e. On the `Drive Configuration` window, select `OK`.

f. On the `RAID Group` window, select `OK`.

g. At the `"Drive Configuration OK?"` message window, select `OK`.

The `SCSI Path - LDEV Configuration` window now appears:

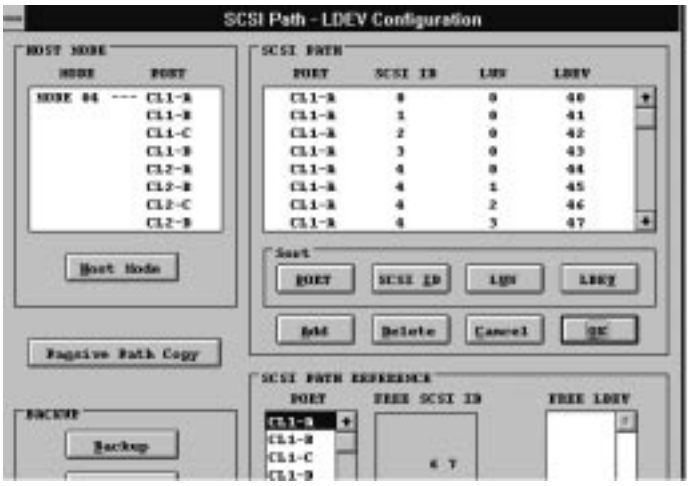

From this screen you can do the following:

- The `HOST MODE` area lets you set the Host Mode for one or more ports.
- The `SCIS PATH` area lets you add, change, or delete path definitions that associate ports on the subsystem, SCSI IDs and / or LUNs, and logical devices.

To select an item, simply click on the item with the mouse.

> **CAUTION**
> *SCSI IDs 6 and 7 are reserved for use by the DYNIX/ptx system. Do not use these SCSI IDs.*

Note, on this window, that the `SCSI PATH REFERENCE` area lists the available ports and shows the SCSI IDs and LUNs and the LDEVs available for assignment to each port. Also note that LDEVs that are already assigned a SCSI path have an asterisk (*) beside them; it they have been assigned more than one SCSI path, they will have two asterisks (**) beside them.

Setting the Host Mode

The Host Mode must be set to **04**. To set the host mode, on the `SCSI Path - LDEV Configuration` window, click `Host Mode`. The `Set Mode` window appears:

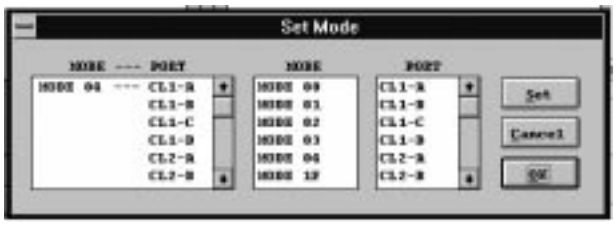

In the Mode and Port listings, highlight the ports for connection to the DYNIX/ptx system and highlight Mode 04, then click Set.

Adding, Changing, or Deleting a Path Definition

To **add or change** a path definition, select the `Add` button. The `Set SCSI Path Parameter` window is displayed:

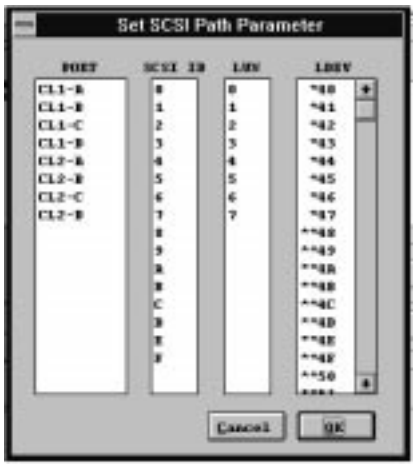

In this window, you associate paths with SCSI IDs/LUNs and logical devices (LDEVs) by clicking the appropriate elements (they are highlighted when clicked), then clicking OK.

To **delete** a path definition, select the `Delete` button. The `Delete SCSI Path` window appears:

Select the path(s) you want to delete, then click OK.

Renewing the Subsystem (Making New Settings Take Effect)

To make the new settings take effect:

1. When all changes in the SCSI Path - LDEV Configuration window have been made, click OK.
2. At the Drive Configuration window, click OK.
3. At the Device Configuration message window, click OK.
4. At the Change Subsystem message window, click OK. **Note** that this does not change the subsystem configuration, but merely approves the changes you made in the SCSI Path - LDEV Configuration window. You must *renew* the subsystem, as follows, for the changes to take effect.
5. At the Maintenance Online menu window, select Maintenance, then select Change Subsystem, then select Renew Subsystem.

The subsystem will now put the new configuration settings into effect.

Step D. Make Cable Connections

When the software installation has been completed and the RAID hardware has been configured, do the following:

1. If you are not already logged in as **root**, do so at this time.

2. Warn all users that the system will be coming down, and make sure that all user activity has stopped.

> **CAUTIONS**
> *ptx/HITACHI-RAIDMP supports **only Fibre Channel Bridge (FCB), differential SCSI connection** of the RAID200 subsystem.*
>
> *Any bridge to which a Hitachi RAID device will be connected must be **dedicated** to Hitachi RAID devices; no other types of devices can be connected to that bridge.*
>
> *Any bridge to which a Hitachi RAID device will be connected must be running the correct version of bridge firmware. Connecting to a bridge running non-supported versions of firmware may lead to unpredictable results.*
>
> *For SCSI connections, be sure to use **only** SCSI cables and terminators listed in the discussion titled **RAID200 Requirements for Use**, earlier in this chapter.*
>
> *SCSI IDs 6 and 7 are reserved for use by the DYNIX/ptx host system. Do not assign either of these SCSI IDs to the subsystem.*

3. Refer to the figure below and make SCSI cable connections between the host system's FCB and the RAID subsystem. This figure shows a (logical) cabling configuration for a multiport RAID subsystem with two (2) four-port CHSs (channel adapters for SCSI). You may need to make appropriate modifications in cabling for your particular installation.

4. If you are relocating any other SCSI devices that were disconnected from the FCBs dedicated to Hitachi RAID devices, connect those devices.

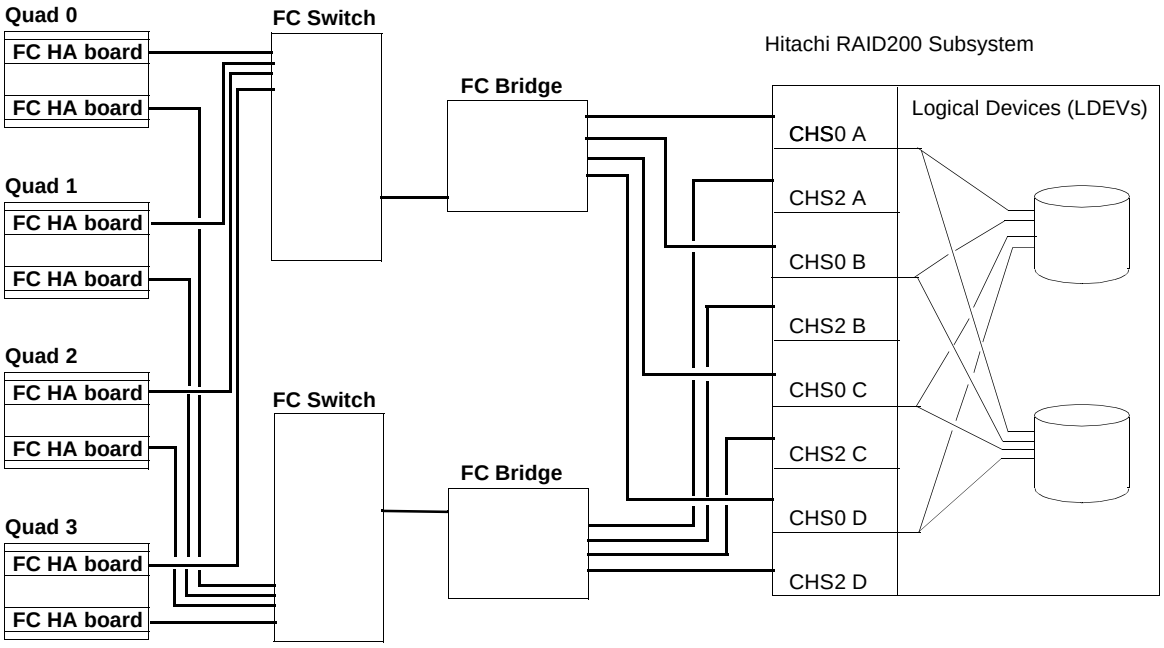

Figure 4-2.
Multiport RAID200 Configuration With Two 4-Port Channel Adapters for SCSI (CHSs).

Step E: Verify the Integration

To verify the integration:

1. Boot the host system.
2. Enter:

   /etc/dumpconf -d | grep sd

   and verify that the appropriate **sd** devices were found at system boot. You should see one entry for each data path to each Hitachi **sd** device. Each entry for the same Hitachi **sd** device will have different SCSI bus assignments.
3. For each **sd** device, use the **diskid** command to determine which **sd** devices are Hitachi devices:

   diskid sd<*n*>

   where <*n*> is the id number of the device. Hitachi RAID200 devices should appear as one of the following:

   OPEN-3 (2.4GB capacity devices)
   OPEN-9 (7.3GB capacity devices)

   Figure 4-3 shows sample **dumpconf** and **diskid** commands and output for a RAID200 subsystem. Figure 4-4 shows a typical connection and CHS and port mappings to cluster connectors.

   NOTE
   *The **dumpconf** and **infodev** commands report unit information in hex; the **diskid** command reports the **cport** information in decimal. You will need to translate between the output of these commands.*

   Port information for the RAID200 can be deciphered from the **diskid** command. A summary of CHS/Port mappings is shown in Table 4-2.
4. Once the devices have been verified, perform I/O to each device to verify its operation.
5. Complete the preparation of the RAID subsystems for use. This may include building VTOCs and filesystems or configuration using ptx/SVM. If you are using VTOCs, refer to the discussion titled *Notes on Usage and Administration*, at the end of this chapter, for details on the VTOCs supplied for RAID200 devices.
6. Bring the system back to multi-user mode and notify users that the system is now available for use.

Table 4-2
CHS/Port Mappings

	Cluster 1		Cluster 2	
CHS #	**0**	**1**	**2**	**3**
Slot Label	**E**	**F**	**Q**	**R**
Port 0				
Port Name	A	E	A	E
cport	O	16	32	48
Port 1				
Port Name	B	F	B	F
cport	1	17	33	49
Port2				
Port Name	C	G	C	G
cport	2	18	34	50
Port 3				
Port Name	D	H	D	H
cport	3	19	35	51

```
# diskid sd12
sd12: scsibus8 cport 0 target 9 lun 0, No Pbay on scsibus8
  Vendor HITACHI, product OPEN-3, revision 0316, serial # R200-8A2C-0049
  capacity 2347 Mbytes
```

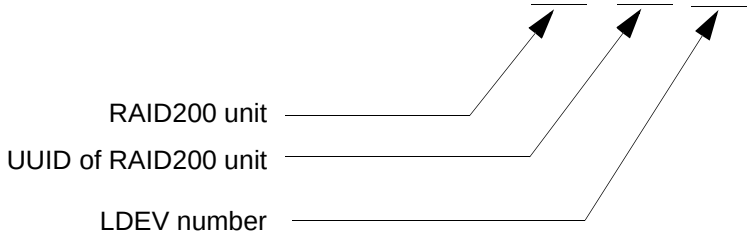

This **sd** device has 2 SCSI paths to the LDEV

```
# dumpconf -d|grep sd12
sd12          sd            12  0x00000090  SM     scsi     scsibus8
sd12          sd            12  0x20000090  SM     scsi     scsibus9
```

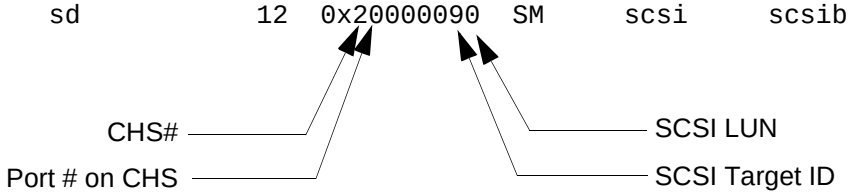

Figure 4-3.
Sample dumpconf and diskid output

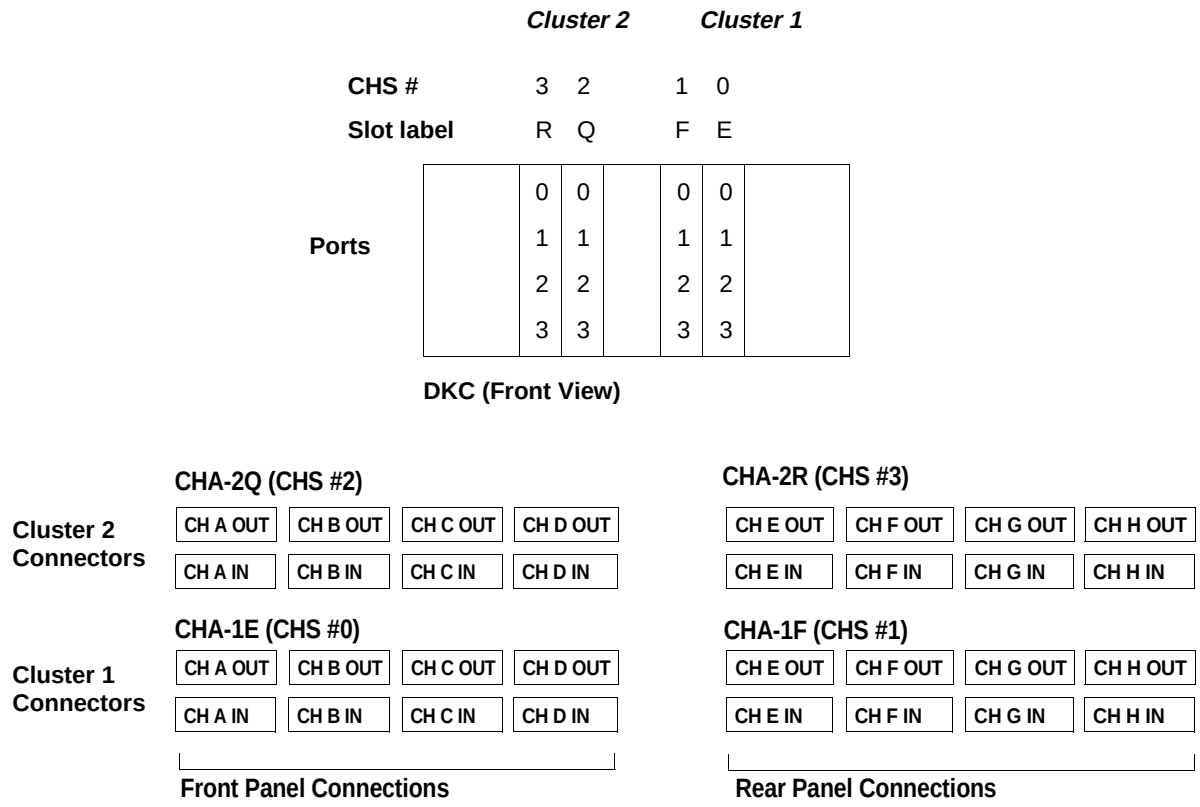

NOTE: The association of ports to connectors depends on the CHS, according to the following scheme:

Port	Connectors for CHA-1E, CHA-2Q	Connectors for CHA-1F, CHA-2R
0	A	E
1	B	F
2	C	G
3	D	H

Figure 4-4.
CHS, Port Mappings and Typical Connection.

Notes on Usage and Administration

Obtaining Subsystem Status

You can use the subsystem's SVP program, available on the built-in laptop, to obtain status on all major subsystem modules, including the location of malfunctioning components. You can also examine the configuration of LDEVs (logical devices) and assignments of these to ports and SCSI IDs/LUNs.

1. On the laptop, bring up the SVP program.
2. From the SVP main menu bar, select the STATUS icon.
3. On the General Status Display, click the icons corresponding to the subsystem modules or to the kind of status information you want to see. Some of the icons may be grayed out. However, the items listed in Table 4-3 should be available.

The icons labeled **DKC** and **DKU-*xx*** represent the subsystem cabinets:

- **DKC** represents the cabinet that contains the controllers, SCSI ports to the host, etc.
- **DKU-L1**, **DKU-L2** represent the left cabinets, containing individual physical disk modules.
- **DKU-R1**, **DKU-R2** represent the right cabinets, containing individual physical disk modules.

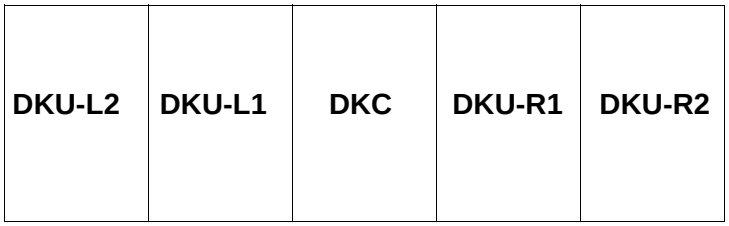

Clicking on either of these icons will display a screen showing the major functional submodules at that level. The layout of icons on the screen will *approximate* the physical layout of physical submodules in the current module. The lowest level screens will show the status of individual functional modules.

At each level, if a component or some sub-component is malfunctioning, the corresponding icon will blink. Continue to click on the blinking icons until you locate the specific component that is malfunctioning.

The other icons will display configuration status information, as indicated in Table 4-3.

Table 4-3
Status Icons and Contents

Icon	Information Displayed
DKC	Displays modules and status information for the left-hand subsystem cabinet. This cabinet contains the SCSI controllers, ports for connection to the host system, etc.
DKU-*xx*	Displays modules and status information for one of the subsystem cabinets that contain the DASs, individual disk modules, and redundancy components for the DASs (fans, power supplies, etc.). DKU-R1 represents the first cabinet to the right of the DKC. DKU-R2 represents the second cabinet to the right of the DKC. DKU-L1 represents the first cabinet to the left of the DKC. DKU-L2 represents the second cabinet to the left of the DKC.
Copy Status	Not available.
Logical Device	Displays the status of logical devices (LDEVs): device status, RAID group membership, physical devices associated with the LDEV, etc.
RAID Status	Displays the status of RAID groups and the individual disk modules within a RAID group: status, disk type, emulation mode.
Version	Displays firmware version information for each component in the subsystem that contains firmware.
LCP/MCP Path	Display path information for logical and physical paths.
SCSI Path	Displays SCSI path information: assignments of ports, SCSI IDs, and LUNs to LDEVs. Also shows available SCSI IDs and LDEVs for each port.

Decoding CHS/Port Errors Reported in ktlog

If there is an error on a CHS and / or port, a message similar to one of the following should appear in **ktlog**:

```
364cbd7d 15:15:09 tolog/warn q1/e6 sd4(port 17): I/O timeout, hung bus?
364cbd7d 15:15:09 tolog/warn q1/e6 sd4(port 17): disabling route
```

When the path is restored, a message similar to the following should appear in **ktlog**:

```
364cbe85 15:19:33 tolog/warn q0/e2 sd4(port 17): resuming route
```

Refer back to Table 4-2, under the discussion titled *Step E: Verify the Integration*, to identify the CHS port.

Sample VTOCs Supplied With ptx/HITACHI-RAIDMP

The ptx / HITACHI-RAIDMP distribution media includes the following VTOC files for use with RAID200 devices:

open-3	For OPEN-3 volumes
open-9	For OPEN-9 volumes

For use with ptx / SVM, the following RAID200 VTOC files are included:

open-3.vxsl	OPEN-3 volumes
open-9.vxsl	OPEN-9 volumes

Refer to the *ptx/SVM Administration Guide* for detailed information on ptx / SVM.

The make_vtoc(1M) Utility

The **make_vtoc**(1M) utility allows you to define custom VTOCs. Refer to the on-line man page or to the hard copy of the man page in Appendix A of this manual for details. Refer to the DYNIX / ptx **mkvtoc** man page for additional details on VTOC creation.

Chapter 5

RAID300 Subsystem Integration

This chapter describes the steps needed to integrate a Hitachi RAID300 RAID disk array subsystem with a DYNIX/ptx host system running ptx/HITACHI-RAIDMP. The following topics are included:

- **Requirements for use.** Gives hardware requirements for the use of this subsystem with ptx/HITACHI-RAIDMP.
- **Restrictions.** Includes restrictions specific to the functioning and use of this subsystem with ptx/HITACHI-RAIDMP.
- **RAID300 RAID concepts.** Gives a brief summary of the RAID terminology and concepts used in association with this subsystem.
- **RAID300 subsystem integration overview.** Gives a high-level overview of the steps required to integrate this subsystem with the host system.
- **Integration procedures.** Includes detailed procedures for ensuring that the subsystem has the required configuration settings and for installing and connecting the subsystem to the host system.
- **Notes on Use and Administration.** Includes information that may be useful in the ongoing administration and use of the subsystem.

RAID300 Requirements for Use

The following are required to integrate a RAID300 subsystem with a DYNIX/ptx host system and this release of ptx/HITACHI-RAIDMP.

NOTE

This release of ptx/HITACHI-RAIDMP supports the following RAID300 host connectivities:

For IBM NUMA-Q systems, switched fabric (FC-SW) and direct connect arbitrated loop (FC-AL) Fibre Channel connections, and FCB SCSI connections.

For IBM Symmetry S5000 systems, QCIC-E SCSI connections.

DYNIX/ptx System Requirements

The following are the DYNIX/ptx system requirements for integration of the RAID300 subsystem:

- The DYNIX/ptx host system must conform to the requirements listed in the software chapter of this manual.

IBM NUMA-Q Requirements

- If connection is via FC-SW Fibre Channel topology, the IBM NUMA-Q system must have enough ports available on the Fibre Channel switch for dedicated connection of the RAID300 devices.
- If connection is via direct connect FC-AL Fibre Channel topology, the IBM NUMA-Q system is restricted to a **maximum** of two (2) quads, with a **maximum** of two (2) paths to each subsystem LDEV (each path through a single FCHA in each quad).
- If connection is via FCB SCSI, the IBM NUMA-Q system must have enough FCB SCSI channels available for *dedicated* connection of the RAID300 devices.

Symmetry S5000 Requirements

- The S5000 system must have enough QCIC-E SCSI channels available for *dedicated* connection of the RAID300 devices.

RAID300 Subsystem Requirements

- Any SCSI connection to a DYNIX/ptx host system must be via *differential SCSI*.
- The subsystem must be running a minimum microcode revision of 52-47-71-00/00 or higher; contact the IBM NUMA-Q Customer Support Center for information on the latest supported microcode release(s).

- Multiport connectivity to the RAID300 subsystem requires that certain RAID300 modes be set; if these modes are not set properly, RAID300 subsystem performance in a multiport configuration will be severely degraded. *These settings can only be changed by Hitachi field service personnel.* If you are connecting the RAID300 subsystem in a multiport configuration, *contact your Hitachi distributor's service representative* to request that these modes be set properly.

RAID300 Host and System Modes

Host Modes. On RAID300 subsystems, each port is assigned a *host mode*, depending on the type of connection the host supports.

- Ports connected to DYNIX/ptx systems *only* require Host Mode 04, the Numaq_Unix (formerly Sequent) mode. This specifies fast/wide SCSI and multiport.
- Ports connected to both DYNIX/ptx and NT® systems (e.g., through a common switch) also require Host Mode 04.
- Ports connected to NT® systems *only* require Host Mode 00.

System Modes. System modes settings are specified for a given Host Mode, and are used for control throughout the entire subsystem.

> **NOTE**
> *The host mode determines which system modes are recognized and active at any given time. For example, if Host Mode 04 is not set for a particular port, System Mode 56 is ignored on that port. First make sure that the correct host mode is set* ***for each port****, then set the system modes.*

The following explains the main functionality of System Modes 56 and 57.

- System Mode 56 improves multiport sequential READ/WRITE performance on multiport configurations. It is only enabled for host connections with Host Mode 04.
- System Mode 57 enables RAID300 FRU error reporting through a SCSI UNIT ATTENTION. An I/O command will be aborted (sense key 0xB) with ASC/ASCQ 0xC300. The driver reports this as "internal FRU/power failure". A subsequent retry of the I/O will complete successfully. There should be one of these messages for each LDEV that is configured on the system.

Required Modes. The following table summarizes the required Host and System Modes.

Table 5-1
Host and System Modes

Operating System	Host Mode (by port)	System Modes
DYNIX/ptx	04	56 & 57
NT only	00	57
DYNIX/ptx & NT	04	56 & 57

Fibre Cabling Requirements

For Fibre Channel (FC-SW or FC-AL) connection to an IBM NUMA-Q host system, the standard fibre cabling available from IBM or Hitachi should be used.

SCSI Cabling and Termination Requirements

NOTE
The Hitachi disk array subsystem's SCSI connectors use a smaller screw thread size (2-56) than the standard SCSI connectors (4-40) on DYNIX/ptx host systems. The cables and terminator listed below use the smaller, 2-56 thread size.

IBM NUMA-Q Connections

SCSI cables to connect an IBM NUMA-Q host system to the subsystem must be chosen from the following:

09P7837	15.0m cable
09P7839	7.5m cable
09P7799	3.0m cable

Symmetry S5000 Connections

SCSI cables to connect a Symmetry S5000 host system to the subsystem must be chosen from the following:

09P7550	10.0m cable
09P7551	20.0m cable

SCSI Termination—FCB SCSI and QCIC-E SCSI

Each SCSI termination on the RAID300 subsystem requires the following SCSI terminator:

09P7552

This terminator is required for both IBM NUMA-Q FCB SCSI and IBM Symmetry S5000 QCIC-E SCSI connections.

RAID300 Restrictions

The following restrictions apply to the connection of RAID300 subsystems to a DYNIX/ptx host system. These are in addition to the general and subsystem-specific restrictions listed in Chapter 2 of this document.

General Restrictions

- Only the host systems and connectivities listed in this document earlier are supported.
- Single-port and multiport configurations are supported.
- RAID300 disks cannot be used as boot or dump devices.

Fibre Channel Connectivity Restrictions

- Fibre Channel switched fabric (FC-SW) and direct connect arbitrated loop (FC-AL) connections are supported.
- There is *no clusters support* for FC-AL connections.
- The only FC switch(es) supported are the ones listed in Chapter 1 of this manual. No other switches or firmware versions are supported, and no switches from other vendors are supported
- No connections through any other switches or through hubs are supported.
- Only FC multimode cables supplied by Hitachi or IBM are supported.
- For fibre connections, up to eight (8) paths are supported.

SCSI Connectivity Restrictions

The following restrictions apply to RAID300 SCSI connections to *all* DYNIX/ptx platforms:

- Only differential SCSI connections are supported.
- External SCSI cable length is limited to 20m on a single bus.
- SCSI IDs 6 and 7 are reserved for QCIC-E and FCB SCSI ports. *Do not* configure the RAID300 subsystem with SCSI ID 6 or 7.
- For SCSI connections, up to two (2) paths are supported.

NUMA-Q SCSI Restrictions

- Multi-initiator SCSI is not supported with the RAID300. For path redundancy in IBM NUMA-Q configurations, it is recommended that the multiport capabilities of the RAID300 be used, along with FC (fibre channel) multipathing with fibre channel switches or bridges.
- The IBM NUMA-Q Fibre Channel Bridge (FCB) does not support Ultra-Wide SCSI protocol. However, it negotiates SCSI-2 protocol.

Symmetry S5000 SCSI Restrictions

- SCSI connectivity on S5000 systems *must be* via QCIC-E SCSI.
- Multi-initiator SCSI is not supported with the RAID300. For path redundancy in S5000 configurations, it is recommended that the multiport capabilities of the RAID300 be used, along with different QCIC board connections on the DYNIX/ptx system.
- On Symmetry S5000 host systems, QCIC-E SCSI does not support Ultra-Wide SCSI protocol. However, it negotiates SCSI-2 protocol.
- There is no ptx/Clusters support for RAID300 devices on Symmetry systems.

RAID300 RAID Concepts

The following discussions give an overview of RAID300 terminology and concepts.

Terminology

The following terms are used in the remainder of this document to describe the Hitachi RAID300 subsystem and software.

Table 5-2
RAID300 Terminology

CHA	Channel Adapter. The MVS (ESCON) connection.
CHS	Channel Adapter for SCSI. Each CHS supports four (4) SCSI buses.
CHF	Channel Adapter for FC. Each CHF supports two (2) FC connections.
DKA	Disk Adapter. The interface between DKC cache and the internal SCSI disk drives.
DKC	Disk Controller. The main chassis that houses the CHA, CHS, cache, and SVP laptop.
DKU	Disk Unit. A frame that houses SCSI disk drives.
frame	Physical chassis that houses one or more hardware components.
LDEV	Logical Device. This is a 3390-3, 3390-8, or 3390-9 compatible volume, in terms of capacity. Each LDEV is assigned and addressable as a SCSI target ID/LUN.
LUSE	LUN Size Expansion. Allows concatenation of multiple volumes into one logical unit.
SCP	SCSI Control Processor. A SCSI interface controller on a Channel Adapter for SCSI (CHS). There are four (4) SCPs on each CHS.
SCSI Multi-port	The Hitachi RAID subsystem provides the capability to allow I/O through multiple SCSI port connections.
SVP laptop	A built-in laptop computer that enables the configuration and monitoring of the subsystem.

Multiported Devices

Certain disk storage units provide multiple physical connections (ports) to multiple I/O buses and then configure their logical volumes to allow access from those ports simultaneously. This feature is called *multiporting*.

Multiporting provides redundant access to individual logical volumes, increasing their availability. If a particular port or I/O bus becomes unavailable, an unaffected port and/or I/O bus can be used to access the logical volume. Locating multiple ports for a logical volume on separate buses provides protection from both an I/O bus failure and the failure of a port within a storage unit. Locating multiple ports for a logical volume on the same I/O bus is legal, but this method does not provide protection from the failure of that I/O bus. With multiporting:

- A port can have more than one logical volume connected to it.
- A single logical volume can be connected to more than one port. E.g., in Figure 5-1, the disk resource is connected to Port 0 and Port 1, and has a path through both fibre channel bridges *Bridge 0* and *Bridge 1*.

The operating system also takes advantage of multiporting to better distribute its I/O load. Doing so *may* reduce the latency of individual I/O requests (although this is not guaranteed).

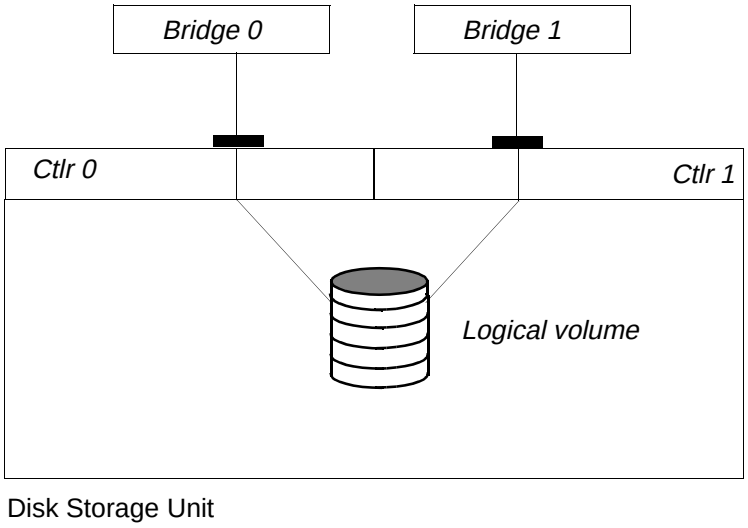

Figure 5-1.
Sample Multiported Disk Storage Unit.

RAID300 Integration Overview

To integrate the RAID300 subsystem with the DYNIX/ptx system, you must complete the following steps.

- Step A: Install the ptx/HITACHI-RAIDMP software.
- Step B: Prepare the RAID hardware for integration.
- Step C: Check or adjust the subsystem's configuration settings.
- Step D: Make cable connections between the host system and the subsystem.
- Step E: Verify the integration.

The following discussions cover these steps in detail or refer to other documentation for details.

Step A: Install ptx/HITACHI-RAIDMP

If you have not already installed the current release of ptx/HITACHI-RAIDMP, refer to Chapter 2 of this manual and install the software at this time.

Step B: Prepare for Integration

To prepare the subsystem for integration, do the following:

1. Unpack the unit.
2. Set aside and save all packing materials in case the unit ever needs to be repackaged.
3. Immediately after the unit is unpacked, inspect the unit for signs of shipping damage. If the unit appears damaged, contact the carrier in charge of shipping the unit and the vendor from which the unit was shipped.
4. Make sure that all parts and subassemblies were received.
5. If the unit appears undamaged, locate the unit where it will be used.
6. Collect all cables and terminators, etc. that will be required.
7. Collect all required documentation.

Step C: Configure the RAID300 Subsystem

CAUTION

The following procedures assume that service personnel for Hitachi or a Hitachi distributor have first configured and formatted RAID groups on the subsystem. This section only defines how to configure paths.

CAUTION

Do not attempt to change the subsystem's configuration while user activity is occurring.

CAUTION

Incorrect configuration of the hardware could lead to data corruption or loss. ***Before any hardware configuration*** *is done, all personnel involved in the installation should carefully define and review the RAID configuration for this site. The following procedure covers* ***only*** *the use of the Online Maintenance menu system to perform the actual configuration tasks, and assumes that configuration planning has already been completed.*

NOTE

This release of ptx/HITACHI-RAIDMP supports Fibre Channel FC-SW and FC-AL connections and FCB SCSI connection of RAID300 subsystems to IBM NUMA-Q host systems, as well as QCIC-E SCSI connection of RAID300 subsystems to Symmetry S5000 host systems. The RAID300 configuration procedure differs with the type of connection. Be sure to choose the correct configuration procedure.

Configuration for Switched Fabric Fibre Channel Connection

This section describes the steps necessary to configure a RAID300 subsystem for operation with a DYNIX/ptx computer system using a switched fabric Fibre Channel connection.

Changing Configuration Settings

To change the hardware configuration settings, use the SVP program on the built-in laptop computer attached to the subsystem.

1. Connect the subsystem to an appropriate AC power source, and bring the subsystem up.
2. Boot the laptop. The SVP program should come up automatically. If it does not, start the SVP program. The main SVP window appears:

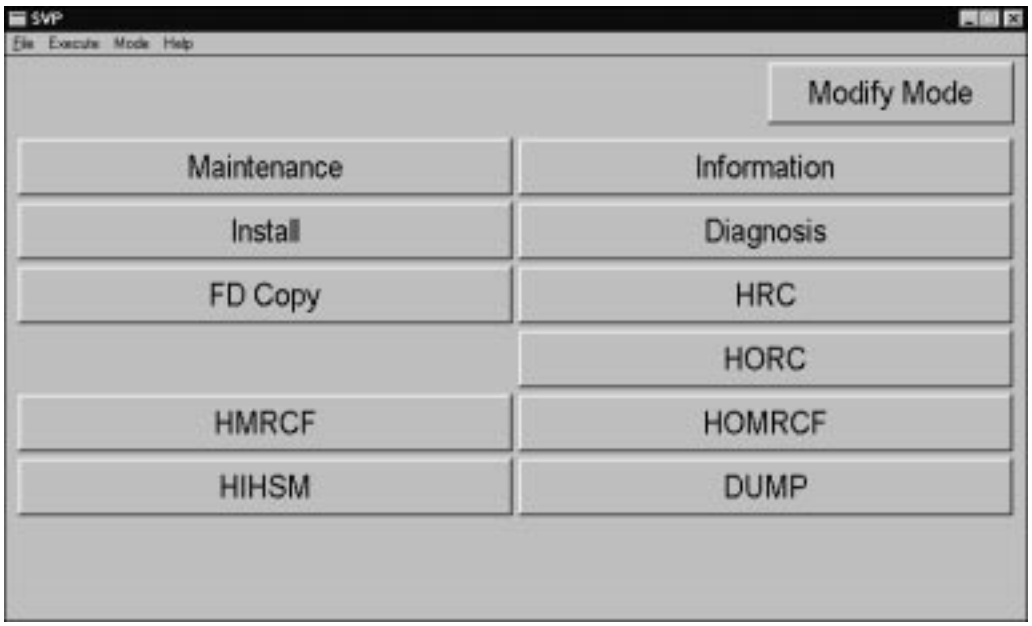

3. Select Modify Mode. You can do this by pulling down the Mode menu and selecting Modify Mode. Or, you can click the View Mode/Modify Mode icon in the upper right corner of the SVP window; this is a toggle switch between these two modes.
4. On the SVP window, click the Install icon. The Install window is displayed.

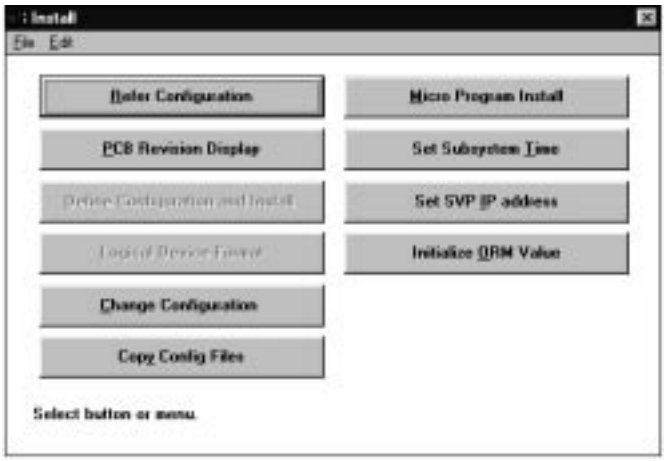

5. Select `Change Configuration`. The SVP program will read the current configuration (this will take a few seconds), and then a `Menu Dialog` box appears.

6. In the `Menu Dialog` box, select `SCSI I/F Configuration` and click `OK`. After another brief delay, the `SCSI Path-LDEV Configuration` window appears. In this window, click the Port button near the upper left corner of the window, then click the arrow to drop down a list of port choices. The screen should appear similar to the following:

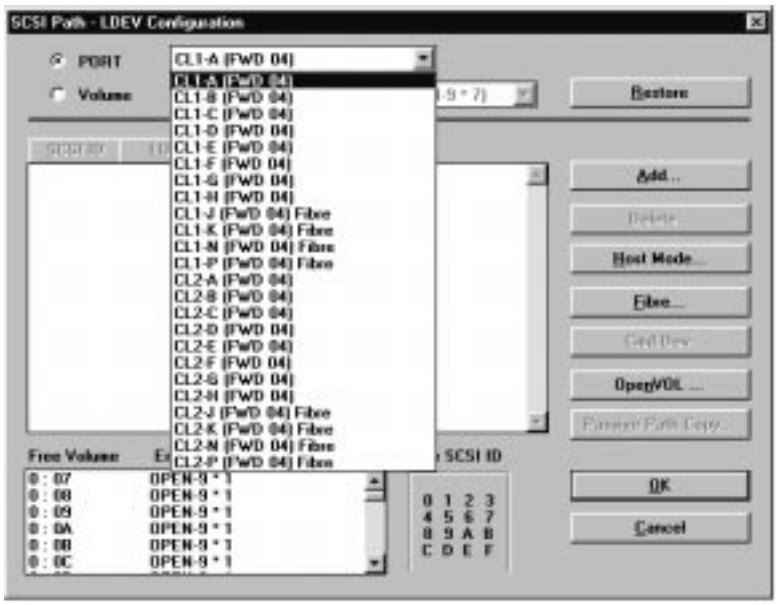

7. From the drop down list, select the desired port. The `SCSI Path-LDEV Configuration` window should now appear similar to the following:

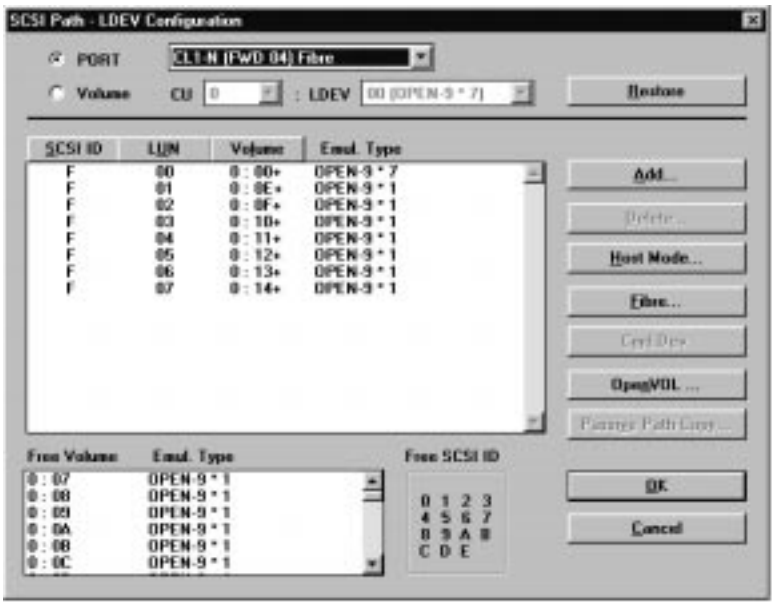

From this `SCSI Path-LDEV Configuration` window, you can add or delete path configurations (you cannot edit a path configuration, only delete it and add a new one with the desired configuration), or change the Host mode. The following discussions cover these tasks.

> **NOTE**
> *If the RAID300 LUN Security feature is enabled on this RAID300 subsystem, each time you click OK on the SCSI Path-LDEV Configuration window, the LUN Security window will appear. You **must** click OK on the LUN Security window for any changes to take effect, even if you are not changing LUN security information. See the discussion titled **Completing the Configuration Process for Fibre Channel Connections**, later in this chapter, for details. Also refer to the **LUN Security** discussion, later in this chapter, for background information on the RAID300's LUN Security feature.*

Changing the Host Mode

> **NOTE**
> *The Host Mode for each port connected to a DYNIX/ptx system must be set to **04**.*

To set the Host Mode for a port, do the following:

1. Get to the `SCSI Path - LDEV Configuration` window.
2. In that window, click `Host Mode`. The `Port Mode Information` window is displayed.

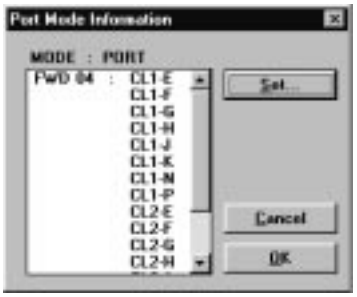

3. In this window, click `Set`. The `Set Port Mode` window is displayed.

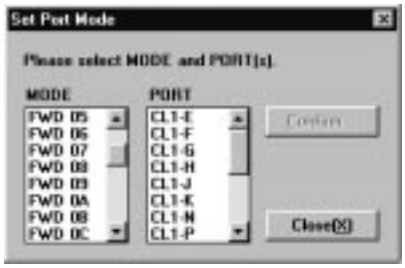

4. In the `Set Port Mode` window, select the desired Host Mode and Port, then click `Confirm`. A verification window appears; if the configuration is correct, click `OK`.

5. The SVP program returns to the `SCSI Path - LDEV Configuration` window. You can perform other tasks or click `OK` to complete the configuration process.

Setting the Fibre Topology

To set the fibre topology information, do the following

1. From the SCSI Path - LDEV Configuration screen, click the `Fibre` icon. A `Fibre Information` screen will be displayed:

2. On the `Fibre Information` screen, click `Topology`. The `Fibre - Topology Information` screen is then displayed:

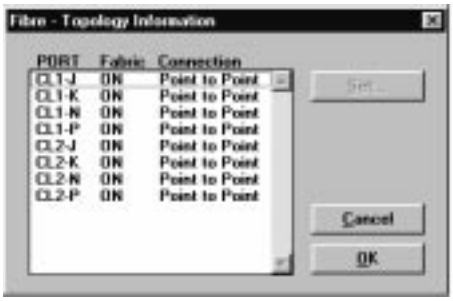

3. On this screen, select the connections whose topology you wish to set. The `Set...` icon will become active.

4. Click the `Set...` icon. The `Set Fibre Topology` screen is displayed:

5. On this screen:
For FC-SW connections, set Fabric to `ON`, Connection to `Point to Point`.
For FC-AL connections, set Fabric to `OFF`, Connection to `FC-AL`.

NOTE: The Host Port ID should always be set to EF. **Do not change this setting.**

Adding a Path Configuration

This discussion shows how to add a dual path configuration. The procedure for adding a single path configuration is similar.

> **NOTE**
> *If you are configuring a multipath, you must configure **all paths at the same time**. If you set only a single path configuration, you cannot edit this configuration later; you will have to delete the single path configuration, then define (add) the desired multi-path configuration.*

To add a path configuration:

1. Go to the `SCSI Path-LDEV Configuration` window and click the `Add` icon. The `Set SCSI Path Parameter - Fibre Channel` window comes up.

NOTE

*For Fibre Channel configurations (FC-SW and FC-AL), the SCSI ID will always be **F**.*

2. In the `Set SCSI Path - Fibre Channel` window, in the Channel section, click the `Fibre` radio button to display information for the fibre port.

3. ***Note:*** If you are defining a multi-path configuration, you must configure *all paths at the same time*.
   To do this, in the PORT section, highlight the ports you want to configure. *Note* that you can use the usual Ctl-Click and Shift-Click features of the mouse button to select multiple ports, etc. on the same screen. This allows you to define multiple paths from a single invocation of this window.

4. Use the mouse to highlight the LUNs to be defined.

5. In the LDEV ID section, use the mouse to highlight the LDEVs to be defined. At this point, the Set `SCSI Path Parameter - Fibre Channel` window should appear similar to the following:

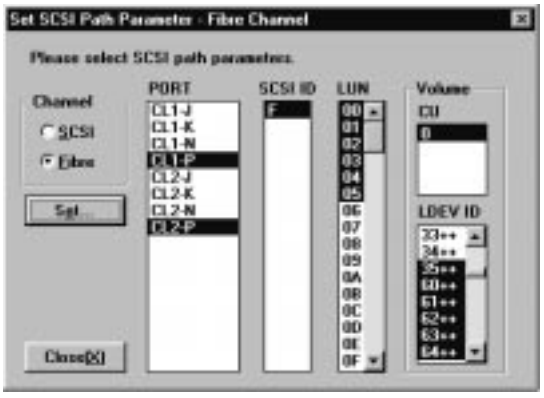

6. Click the `Set` icon. The program will display the `Add SCSI Path` window, listing the path configuration you have defined:

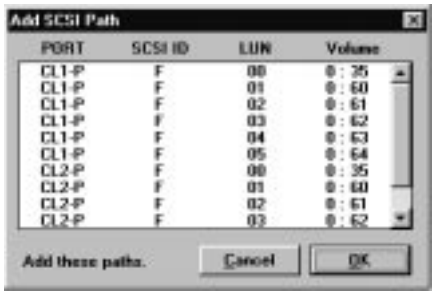

7. If this information is correct, click `OK`. The `Set SCSI Path - Fibre Channel` window reappears, showing the path configuration defined. If this configuration is okay, click `Close`. The SCSI Path - LDEV Configuration window now reappears.

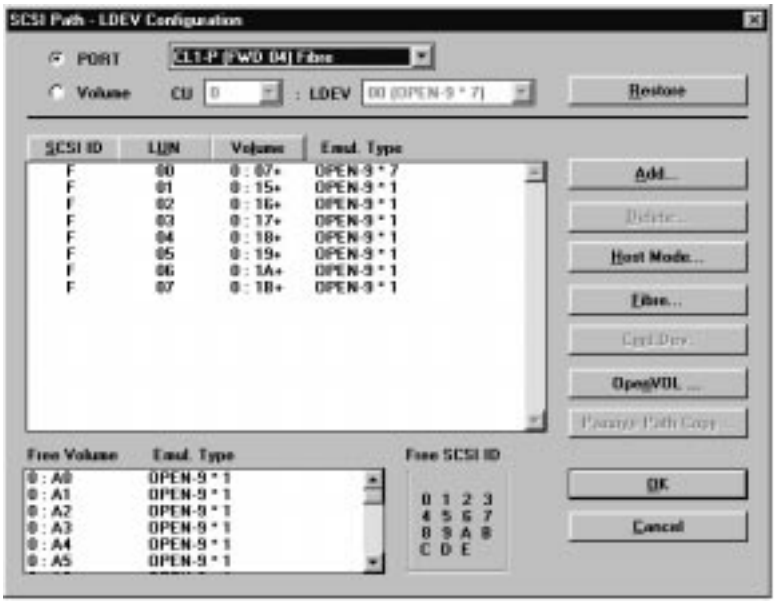

8. If you have no other tasks to perform here, click `OK` to complete the configuration process.

Configuring Volumes for LUSE

This discussion shows how to configure volumes for LUSE.

1. From the `SCSI Path - LDEV Configuration` screen, click `OpenVOL`. The `LU Expansion Define` screen appears:

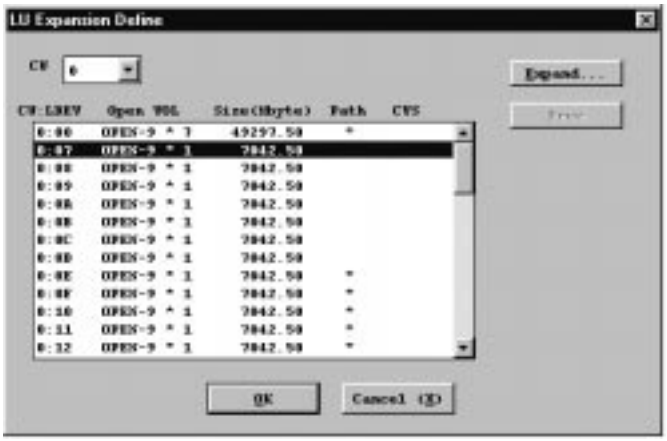

2. On this screen, highlight the LDEV you want to configure for LUSE, then click `Expand`. The `LU Expansion` screen now appears:

CAUTION

*Whatever number, **N**, of volumes you enter to configure for LUSE, the system will configure the **next N consecutive volumes** for LUSE. You cannot control this (i.e., you cannot specify non-consecutive volumes).*

3. On the `LU Expansion` screen, in the `Enter VOL count` field, enter the number of *sequential* volumes you want to configure for LUSE (the sample screen below shows seven (7) volumes entered), then click `OK`. The `LU Expansion Define` screen comes back, showing the number of volumes to be configured and the logical device size; e.g.:

4. If this configuration is acceptable, click `OK`. The `SCSI Path - LDEV Configuration` screen is re-displayed.

NOTE

*You **must** assign the LUSE LDEV to a port **before exiting this screen**.*

5. On the `SCSI Path - LDEV Configuration` screen, select the `Port` button, then select the desired port from the drop down list to the right of this button. Then click `Add` to assign the LUSE LDEV to the selected Port.

6. The `Set SCSI Path - Fibre Channel` screen appears. On this screen, you must do the following:

- In the Channel section, select the `Fibre` button.
- In the Port section, select the ports to be used (this example shows a *dual path* configuration).
- The SCSI ID is *always set to F*; it cannot be changed.
- Select the LUN (here, 00).
- In the Volume section, select the volume number (here, 0).
- In the LDEV ID section, select the LUSE volume ID (here, 07).

This screen should now appear similar to the following:

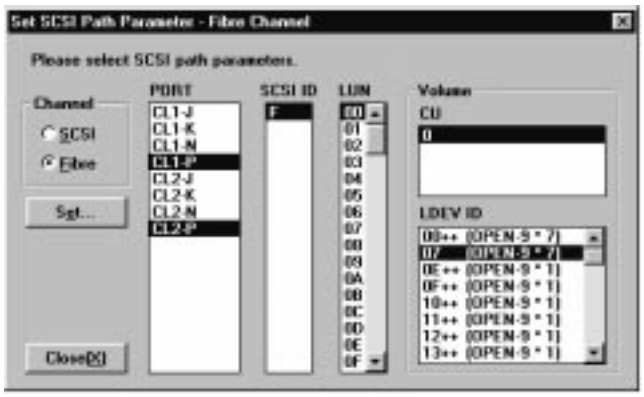

7. When the information on this screen is as you want it, click `Set`. The `Add SCSI Path` screen appears:

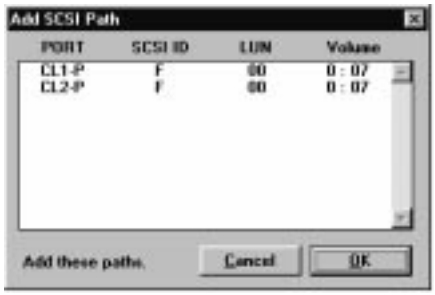

8. On this screen, click `OK`. The `Set SCSI Path Parameter - SCSI Channel` screen appears, showing the LDEV ID (here, 07) that you have configured for LUSE.

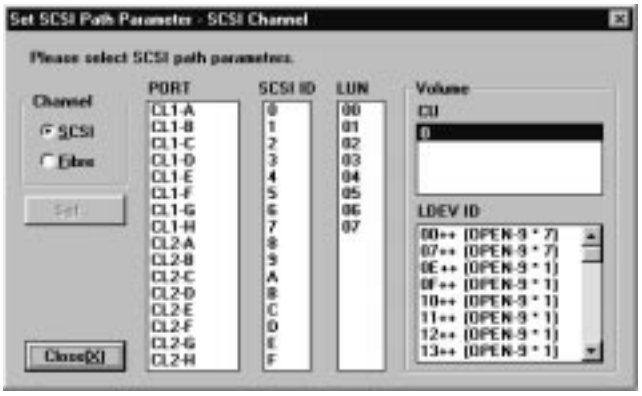

Deleting Path-LDEV Configurations

The following discussion shows how to delete path information associated with a particular LDEV.

1. In the upper left corner of the `SCSI Path-LDEV Configuration` window, there are two buttons: `PORT` and `Volume`. Select `PORT`.
2. From the PORT drop-down list, select the port for which you want to delete path configuration data.
3. Use the mouse to select the desired path configurations to remove. In the example shown here, all LUNs associated with the given path are highlighted for deletion. ***Note*** that you can use the usual Ctl-Click and Shift-Click mouse button features to highlight more than one element to delete in a given section of the screen.

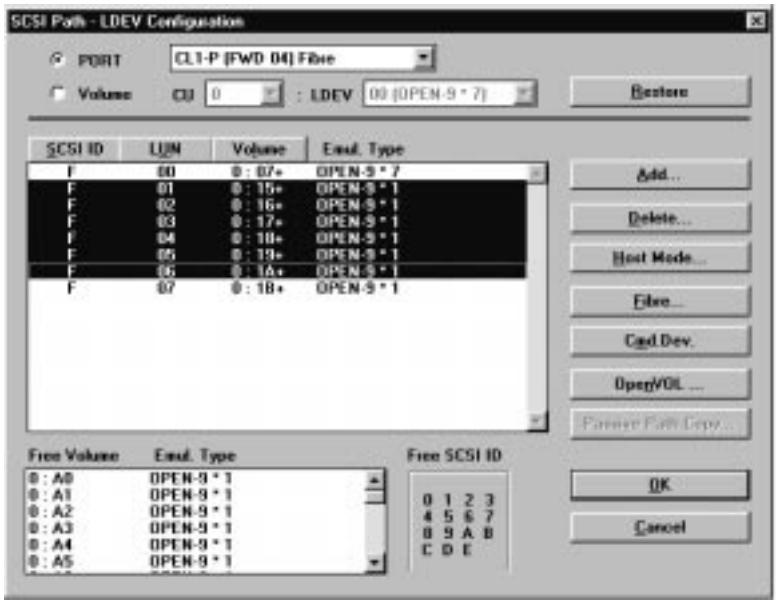

4. When the desired information has been highlighted, click `Delete`. The SVP program will check the paths and display a verification window.

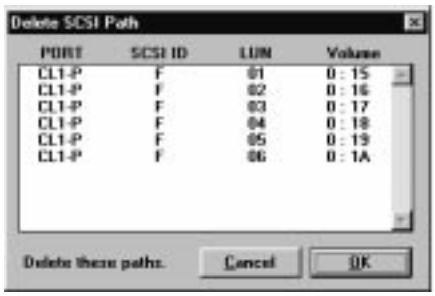

5. In the verification window, click `OK` to delete the path configuration(s). You can then perform other configuration tasks, or click `OK` in the `SCSI Path - LDEV Configuration` window to complete the configuration process.

Completing the Configuration Process for Fibre Channel Connections

When you have completed all configuration tasks, do the following:

1. At the `SCSI Path - LDEV Configuration` window, click `OK`.
2. If you are ***not*** using the LUN Security feature, ignore the remainder of this step and go directly to Step 3 to complete the configuration process.

   If you are using the LUN Security feature, complete the remainder of this step.

   LUN Security Operations

   If the LUN Security feature is enabled, each time you click `OK` on the `SCSI Path - LDEV Configuration` window, the `LUN Security` window will display:

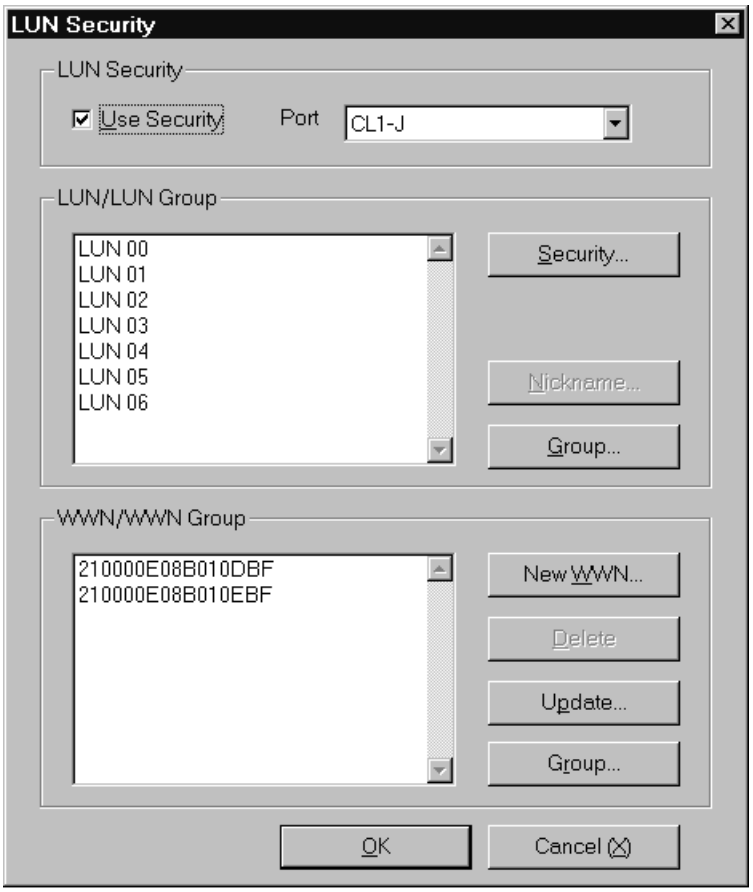

CAUTION

*If this screen appears, you **must click OK** on this screen to make any configuration changes you have initiated from the SCSI Path - LDEV Configuration window take effect, even if you are not changing LUN Security information. If you click Cancel on the LUN Security window, all changes will be lost, **including path and LDEV changes** you initiated from the SCSI Path - LDEV Configuration window.*

If you have modified port or path definition information for ports or paths that are part of one or more LUN Security configurations, you will need to update the LUN Security information for *each of these ports and LUN*s. To do this:

A. List or note *all* ports and LUNs for which the LUN Security information needs to be updated.

B. On the `LUN Security` window, use the drop-down `Port:` list to select the first port whose LUN Security information needs to be updated.

C. Make sure that the `Use Security` box is checked for the port displayed.

At this point, the `LUN Security` window will list all LUNs and all FCHA worldwide names (WWNs) associated with the selected port. You can now **delete** or **add** WWNs to the list, or you can change which WWNs are associated with one or more LUNs associated with this port.

Deleting a WWN. To **delete** a WWN, simply click the WWN to highlight it, then click the `Delete` button.

Adding a WWN. To **add** a WWN not listed for this port, do the following:

i. Click `New WWN`. The `WWN` window displays.

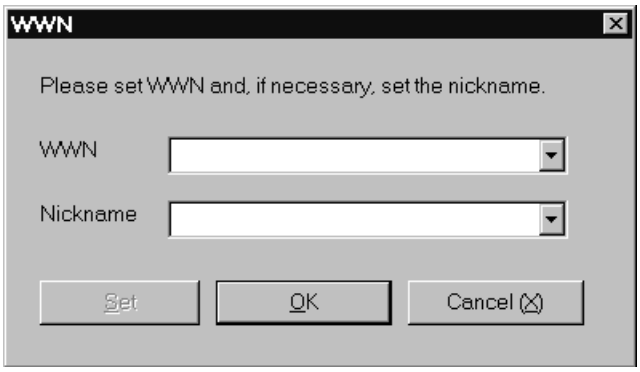

ii. In the WWN field, enter the 16-digit HEX WWN of the selected FCHA (*do not* enter spaces, colons, or other separators; enter the HEX digits of the WWN consecutively), then click `Set`. **Do not click `OK`!** When you click `Set`, the `LUN Security` window displays once again, and the new WWN will be listed in the lower area of the window.

D. Now, you must update the WWNs assigned to each LUN associated with this particular port. To do this:

i. On the `LUN Security` window, select the first LUN associated with the selected port and click `Security`. The `Security` window displays:

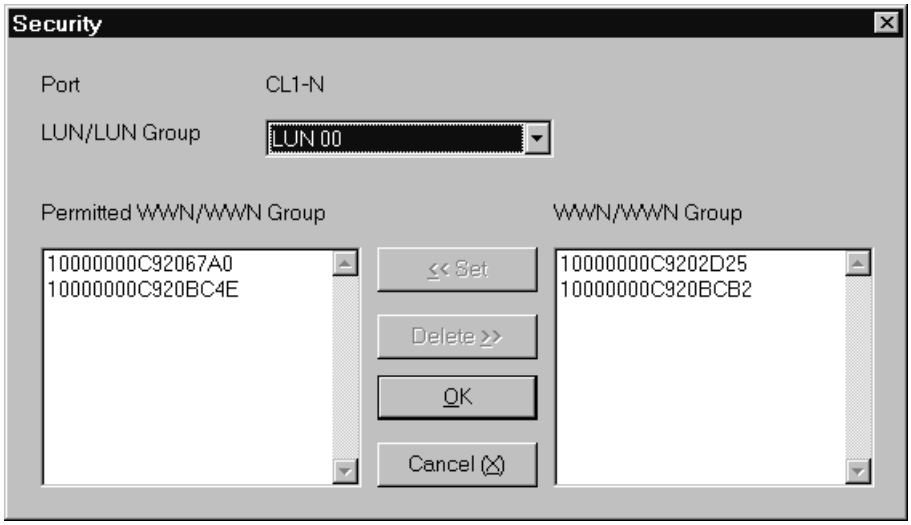

ii. The LUN listed near the top of the window will be visible to those and only those WWNs listed in the left "Permitted" area; the LUN will *not be visible* to WWNs listed on the right. To move a WWN from one area to the other, highlight the WWN, then click `Set` to move it to the "Permitted" area or click `Delete` to move it to the right-hand area.

iii. When the visibility of this LUN is the way you want it, click `OK`. The `LUN Security` window again displays.

iv. Repeat Steps i-iii above for each LUN associated with the selected port.

E. Repeat Steps B—D above for each port whose LUN Security information may have changed.

F. When done, at the `LUN Security` window, click `OK`. The `LUN Security` verification window now displays.

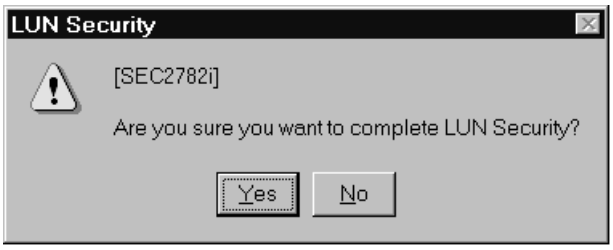

G. Click `OK` to complete the configuration process, or `No` to return to the `LUN Security` window and continue LUN Security configuration, or even go back to the `SCSI Path - LDEV Configuration` window and continue configuration tasks from there.

3. Once you have clicked OK at the `SCSI Path - LDEV Configuration` window (or `LUN Security` window, if LUN Security is enabled), a small number of additional verification windows may appear. At each one, you can cancel the operation or click `OK` to continue.

NOTE
On some subsystem firmware versions, you may see a screen asking for information about your network security setup. Network security is not supported by ptx/HITACHI-RAIDMP at this time. If this screen appears, just click OK.

4. When the `Changing SCSI Configuration` window appears, click `OK`. There will be a brief delay while the program reads configuration data, and then it will ask you to insert the Configuration FD (floppy disk). (The SVP program will make a backup copy of the subsystem configuration data to this floppy disk.)

5. Insert the Configuration floppy. An `FD Backup` window will appear, notifying you that the program is copying subsystem configuration data.

6. When you are prompted, remove the Configuration floppy disk.

This completes the RAID300 configuration process for Fibre Channel connection to the DYNIX/ptx host system.

Configuration for SCSI Connection

> **NOTE**
> *The RAID300 configuration procedure is the same for IBM NUMA-Q FCB SCSI connections and Symmetry S5000 QCIC-E SCSI connections.*

This section describes the steps necessary to configure a RAID300 subsystem for SCSI connection to a DYNIX/ptx computer system.

Changing Configuration Settings

To change the hardware configuration settings, use the SVP program on the built-in laptop computer attached to the subsystem.

1. Connect the subsystem to an appropriate AC power source, and bring the subsystem up.
2. Boot the laptop. The SVP program should come up automatically. If it does not, start the SVP program. The main SVP window appears:

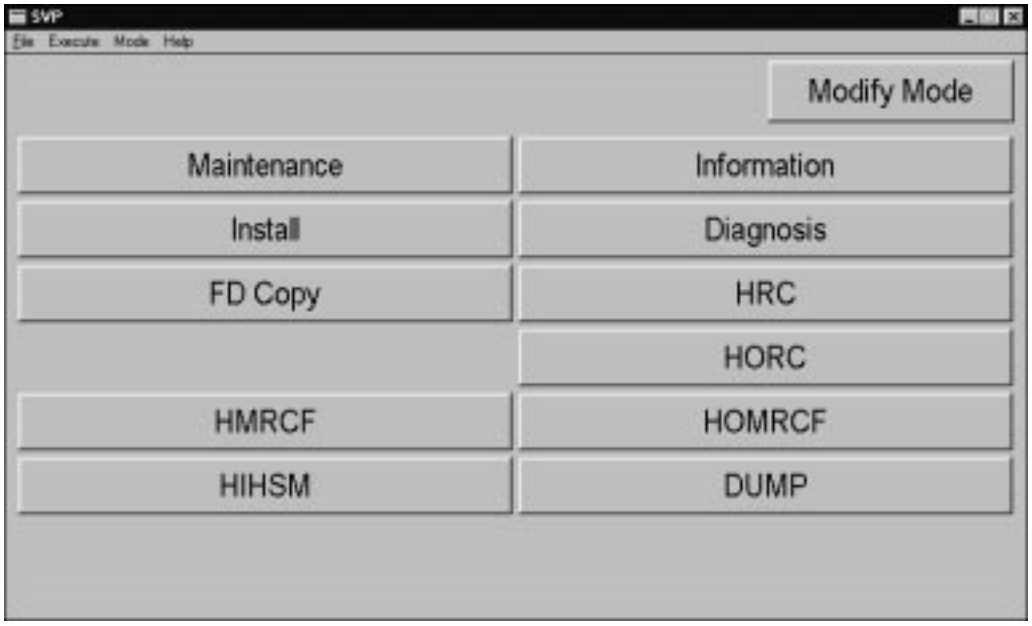

3. Pull down the Mode menu and select Modify Mode.
4. On the SVP window, click the Install icon. The Install window is displayed.

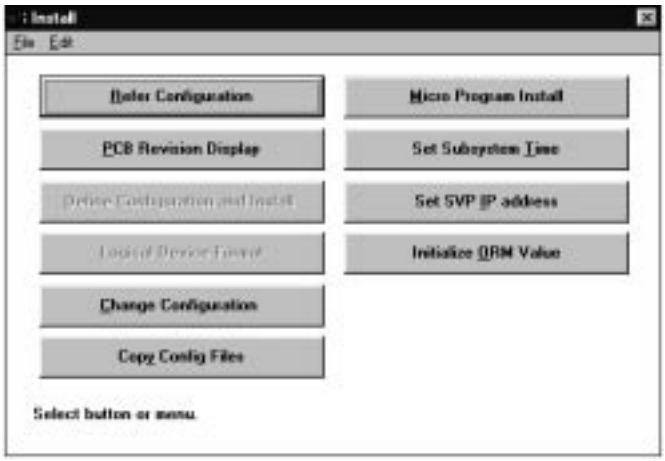

5. Select `Change Configuration`. The SVP program will read the current configuration (this will take a few seconds), then a `Menu Dialog` box appears.

6. In the `Menu Dialog` box, select `SCSI I/F Configuration` and click `OK`. After another brief delay, the `SCSI Path-LDEV Configuration` window appears:

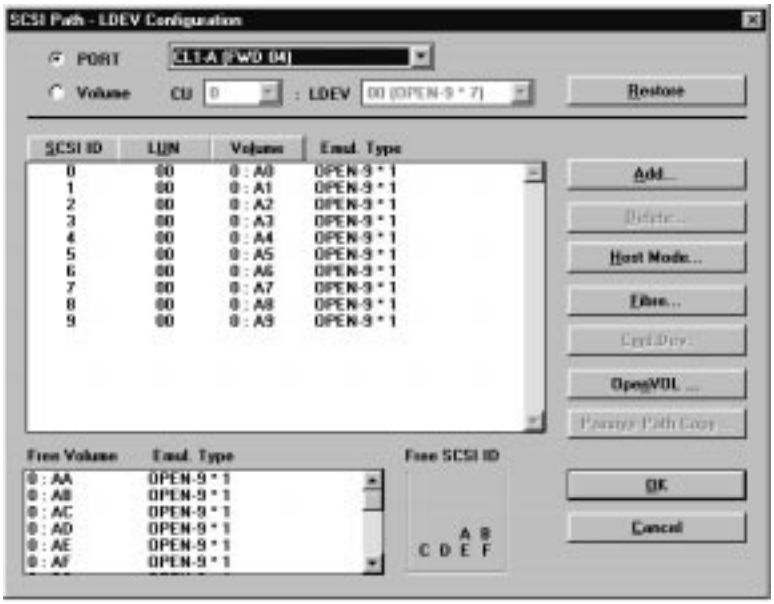

From this `SCSI Path-LDEV Configuration` window, you can add or delete path configurations (you cannot edit a path configuration, only delete it and add a new one with the desired configuration), or change the Host mode. The following discussions give guidelines for completing these tasks.

Changing the Host Mode

> **NOTE**
> *The Host Mode for each port connected to a DYNIX/ptx system must be set to* ***04****.*

To set the Host Mode for a port, do the following:

1. Get to the `SCSI Path - LDEV Configuration` window.
2. In that window, click `Host Mode`. The `Port Mode` Information window is displayed.

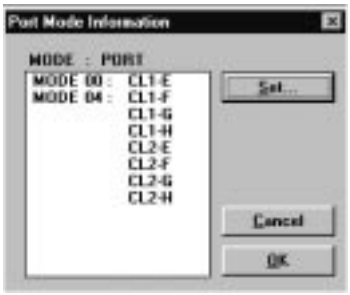

3. In this window, click `Set`. The `Set Port Mode` window is displayed.

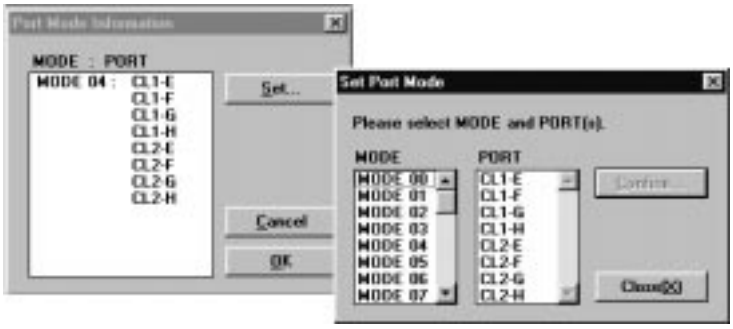

4. In the `Set Port Mode` window, click on the desired Host Mode and the Port, then click `Confirm`. A verification window appears; if the configuration is correct, click `OK`.
5. The SVP program returns to the `SCSI Path - LDEV Configuration` window. You can perform other tasks or click `OK` to complete the configuration process.

Deleting Path-LDEV Configurations

You can delete a path configuration (an association of an LDEV, SCSI port, SCSI ID, and LUN) in one of two ways:

- You can delete all path information associated with a particular *SCSI port*.
- You can delete all path information associated with a particular *volume*, i.e., with a particular *LDEV*.

The following discussion shows how to delete path information associated with a particular port.

1. In the upper left corner of the `SCSI Path-LDEV Configuration` window, there are two buttons: `PORT` and `Volume`. Select `Port`.

2. From the Port drop-down list, select the port for which you want to delete configuration data. The illustration below shows port CL1-A, with volumes A0—A9 selected.

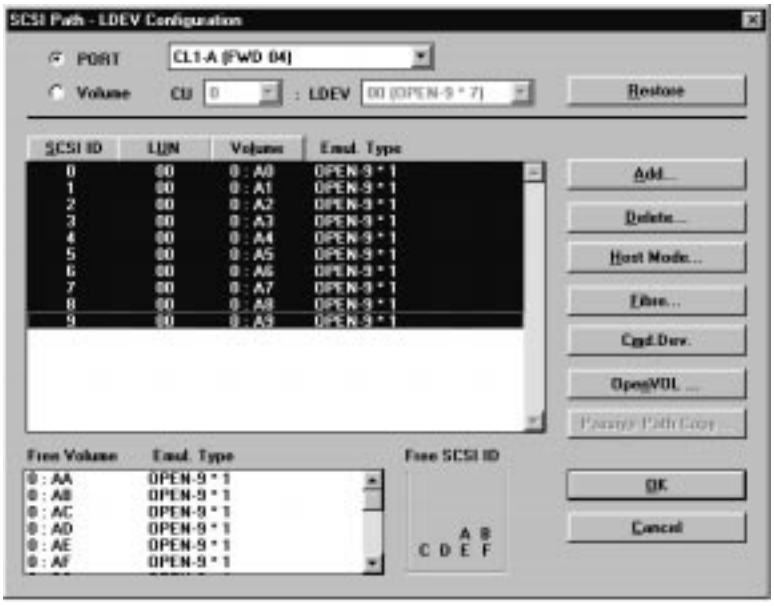

3. Use the mouse to select the desired volumes to remove. In this example, all 10 LUNs have been selected. *Note* that you can use the usual Ctl-Click and Shift-Click mouse button features to highlight more than one logical device (LDEV).

4. When all desired LDEVs have been highlighted, click `Delete`. The SVP program will check the paths and display a `Delete SCSI Path` verification window. Click `OK` to delete the path configuration(s). You can then perform other configuration tasks, or click `OK` in the `SCSI Path - LDEV Configuration` window to complete the configuration process.

Adding a Path Configuration

To add a path configuration:

1. Go to the `SCSI Path-LDEV Configuration` window and click the `Add` icon. The `Set SCSI Path Parameter - SCSI Channel` window comes up.

2. Use the mouse to select the Port, SCSI ID, LUN, and LDEV IDs that define the path configuration you want to add. *Note* that you can use the usual Ctl-Click and Shift-Click features of the mouse button to select multiple ports, etc. on the same screen. This allows you to define multiple paths from a single invocation of this window.

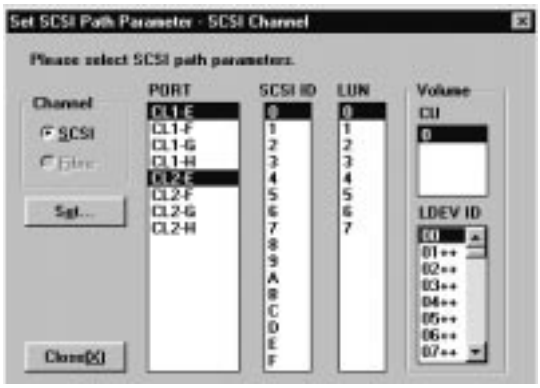

3. When the path configurations you want to add have been selected, click `Set`. The program will display a `Checking Volume Status` window while it verifies the information you have selected. If any path configuration you have attempted to define conflicts with existing paths (e.g., a SCSI ID / LUN combination is already in use), the SVP program will notify you.

   CAUTION
   SCSI IDs 6 and 7 are reserved for use by the host system's QCIC-E and/or FCB SCSI ports. Do not assign either of these IDS to the subsystem.

4. When the path configurations have been verified, in the `Checking Volume Status` window, click `OK`. The program returns to the `Set SCSI Path Parameter - SCSI Channel` window. In that window, if you have no other tasks to perform, click `Close`.

5. The SVP program now returns you to the `SCSI Path - LDEV Configuration` window. If you have no other tasks to perform here, click `OK` to complete the configuration process.

Completing the Configuration Process for SCSI Connection

When you have completed all configuration tasks, do the following to complete the configuration process:

1. At the `SCSI Path - LDEV Configuration` window, click `OK`. The program will display a few additional verification windows. At each one, you can cancel the operation or click `OK` to continue.
2. When the `Changing SCSI Configuration` window appears, click `OK`. There will be a brief delay while the program reads configuration data, and then it will ask you to insert the Configuration FD (floppy disk). (The SVP program will make a backup copy of the subsystem configuration data to this floppy disk.)
3. Insert the Configuration floppy. An `FD Backup` window will appear, notifying you that the program is copying subsystem configuration data.
4. When you are prompted, remove the Configuration floppy disk.

This completes the RAID300 configuration process for SCSI connection to the DYNIX/ptx host system.

Step D: Make Cable Connections

When the software installation has been completed and the RAID subsystem has been configured, do the following:

1. If you are not already logged in as **root**, do so at this time.

2. Warn all users that the system will be coming down, and make sure that all user activity has stopped.

> **NOTE**
> *ptx/HITACHI-RAIDMP supports switched fabric (FC-SW) and arbitrated loop (FC-AL) Fibre Channel connections and FCB SCSI connections of RAID300 subsystems to IBM NUMA-Q host systems, as well as QCIC-E SCSI connection of RAID300 subsystems to IBM Symmetry S5000 systems. Be sure to choose the correct cabling diagram, below.*
>
> **CAUTIONS for SCSI Connection**
> *Only **differential SCSI connections** are supported.*
>
> *Only **Fibre Channel Bridge (FCB) SCSI** connections to IBM NUMA-Q systems are supported.*
>
> *Only **QCIC-E SCSI** connections to IBM Symmetry S5000 systems are supported.*
>
> *For SCSI connection to any host, be sure to use **only** SCSI cables and terminators listed earlier in this document.*
>
> **NOTE**
> *The following figures are intended to provide cabling guidelines only. You may need to make appropriate modifications in cabling to match particular installation and subsystem configuration.*

3. Refer to the appropriate cabling diagram below and make cable connections between the subsystem and the DYNIX/ptx host system.

- Figure 5-2 shows cabling for a typical switched fabric Fibre Channel connection to a IBM NUMA-Q host system.
- Figure 5-3 shows cabling for a typical arbitrated loop (FC-AL) Fibre Channel connection to a IBM NUMA-Q host system.
- Figure 5-4 shows cabling for a typical FCB SCSI connection to an IBM NUMA-Q host system.
- Figure 5-5 shows cabling for a typical QCIC-E SCSI connection to an IBM Symmetry S5000 host system.

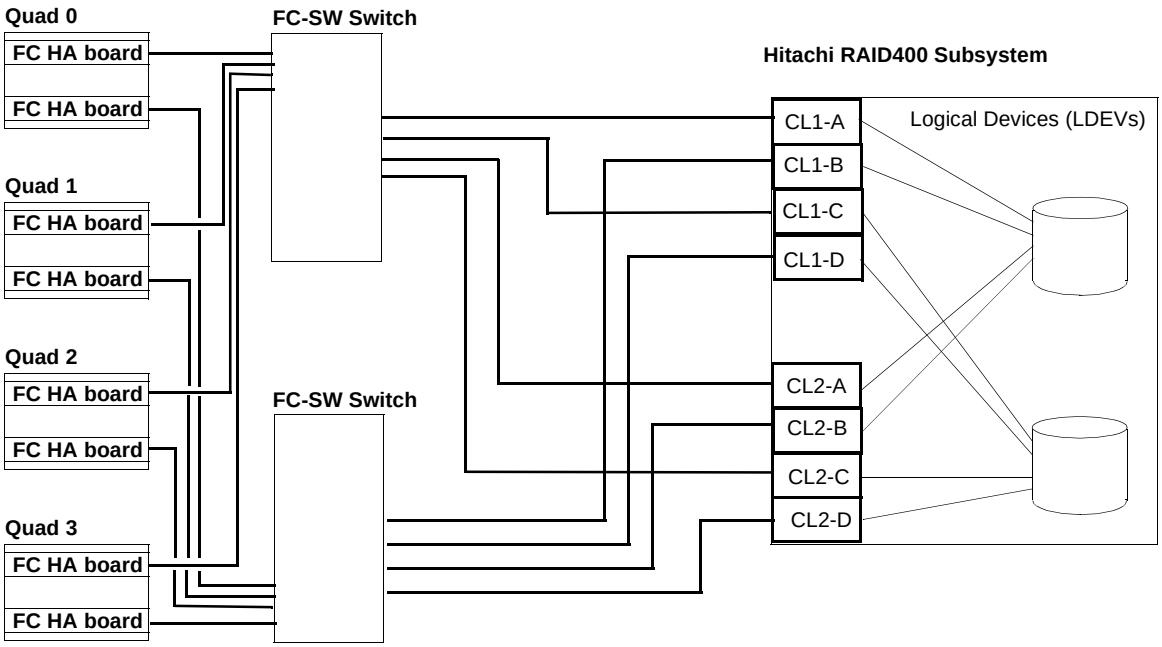

Figure 5-2.
Switched Fabric (FC-SW) Fibre Channel Connection to an IBM NUMA-Q System.

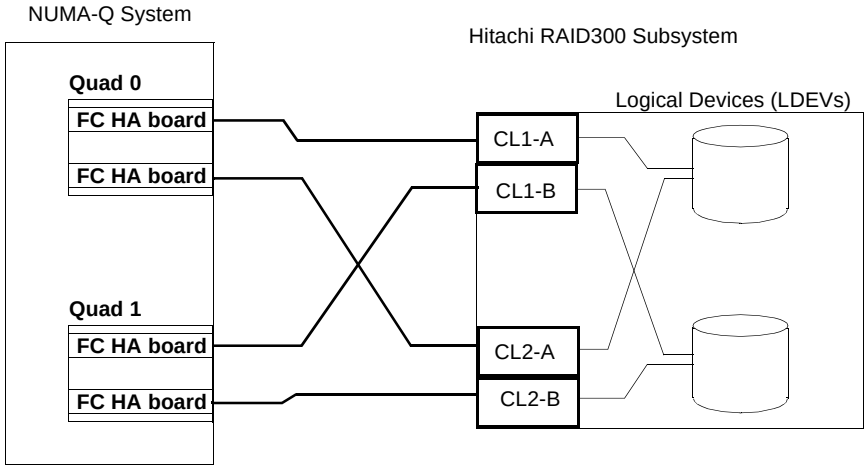

Figure 5-3.
Arbitrated Loop (FC-AL) Connection to an IBM NUMA-Q System.

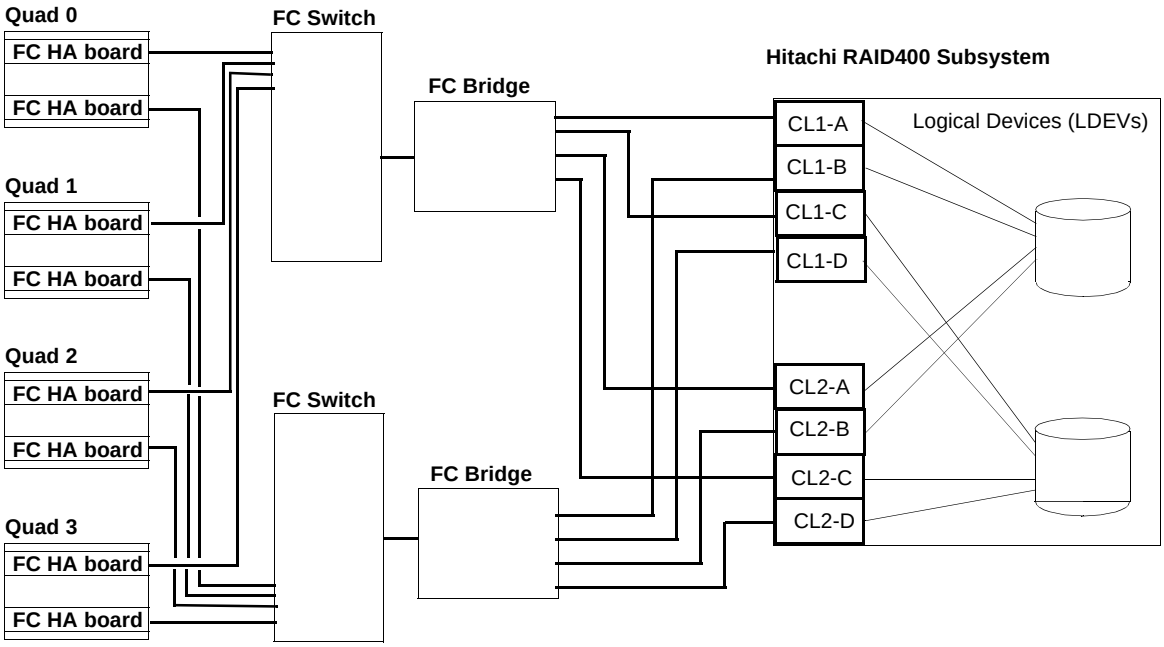

Figure 5-4.
FCB SCSI Connection to an IBM NUMA-Q System.

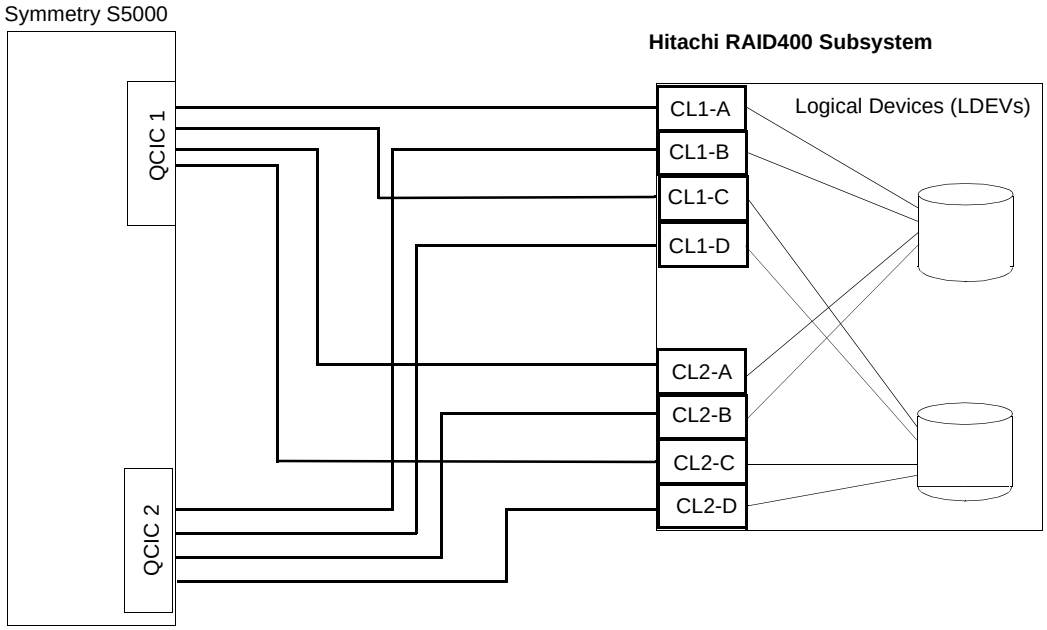

Figure 5-5.
QCIC-E SCSI Connection to an IBM Symmetry S5000 System.

Step E: Verify the Installation

To verify the installation:

1. If this is **not** an arbitrated loop (FC-AL) connection, skip this step and go directly to Step 2, below.

   Direct Connect FC-AL Connections *Only*
   If this is an arbitrated loop (FC-AL) connection, do the following at the DYNIX/ptx **root** prompt:

   a. Change (**cd**) to the directory */usr/conf/uts/io/ff*.

   b. Open the file *ff_space.c* for editing.

   c. In this file, find the line that reads:

      ```
      ff_topology=FC_TOPOLOGY_FABRIC
      ```

      and change this line to read:

      ```
      ff_topology=FC_TOPOLOGY_AL
      ```

   d. Save and exit the file.

   This sets up the DYNIX/ptx system for arbitrated loop communications.

2. Reboot the DYNIX/ptx system.

3. When the system has rebooted, enter:

   /etc/dumpconf -d | grep sd

   and verify that the appropriate **sd** devices were found at system boot. You should see one entry for each data path to each Hitachi **sd** device. Each entry for the same Hitachi **sd** device will have different SCSI bus assignments.

4. For each **sd** device, use the **diskid** command to determine which **sd** devices are Hitachi devices:

   diskid sd<*n*>

   where <*n*> is the id number of the device. Hitachi devices should appear as one of the following:

   OPEN-3 (2.4GB capacity devices)
   OPEN-8 (7.3GB capacity devices)
   OPEN-9 (7.3GB capacity devices)
   3390-3A (2.8GB capacity)
   3390-3B (2.8GB capacity)
   3390-3C (2.8GB capacity)
   3380-KA (1.8GB capacity)
   3380-KB (1.8GB capacity)
   3380-KC(1.8GB capacity)

OPEN-K (1.8GB capacity)
6586-KA
6586-KB
6586-KC
6588-3A
6588-3B
6588-3C

5. At this point you will want to determine path configuration information for the subsystem. You can do this using the **dumpconf**, **diskid**, and **infodev** utilities.

NOTE on Reporting of LUSE Expanded Devices

*Volumes configured with LUSE (LUN Size Expansion) will be reported as **device* number**, where device is the underlying (non-expanded) device type and number is the number of configured volumes. For expanded volumes, the volume capacity is reported as the capacity of the concatenated volume. For example, seven OPEN-9 devices would be reported with device id **OPEN-9* 7** and a capacity of 49297.*

NOTE

*The **dumpconf** and **infodev** commands report unit information in hex; the **diskid** command reports **cport** information in decimal. You will need to translate between the output of these commands. Table 5-2 provides a summary for decoding and translating between these various reporting outputs.*

The following figures and table summarize configuration information for a typical RAID300 subsystem:

Fibre Channel Connections
Figure 5-6 shows sample output from **dumpconf**, **infodev**, and **diskid** for a typical switched fabric (FC-SW) connection, including **diskid** output for LUSE-configured volumes.
Figure 5-7 shows sample output from **dumpconf**, **infodev**, and **diskid** for a typical arbitrated loop (FC-AL) connection.

SCSI Connections (FCB and QCIC-E)
Figure 5-8 shows sample output from **dumpconf** and **diskid** for a typical SCSI connection.
Figure 5-9 shows CHS/Port mappings.

CHS/Port Mappings
Table 5-2 shows CHS/CHF-to-Port mappings for *both SCSI and Fibre Channel connections.*

```
# /etc/dumpconf -d |grep sd12
sd12         sd           12  0x00210213  SMX    fcp      fabric0
sd12         sd           12  0x40210213  SMX    fcp      fabric1

* 5th and 6th digits show fabric port (both on port 2 of each of fabric here):
first two digits show 7700E CHS # and port#
```

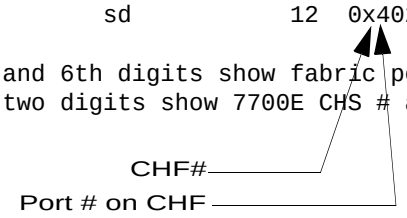


```
# infodev -a sd12
Vendor ID                       : HITACHI
Product ID                      : OPEN-9
Revision Level                  : 5247
Device Class                    : disk
Capacity (in blocks)            : 14423040
Block Size                      : 512
Fibre Channel Address           : 0xef
FCP LUN                         : 00:00:00:00:00:00:00:00
Located on                      : fabric5
```

diskid report for LUSE-configured volumes

```
# diskid sd12
sd12: fabric0 cport 0 fc id 0x210213, fcp lun 0x400
  No Pbay on fabric0
  Vendor HITACHI, product OPEN-3*5, revision 5246, serial # R300-79A4-0014
  capacity 11735 Mbytes
```

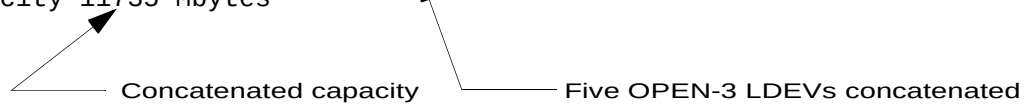

Figure 5-6.
Sample dumpconf, infodev, and diskid Output for Switched Fabric (FC-SW) Connections.

```
# diskid sd26
sd26: fabric5 cport 113 fc id 0xef, fcp lun 00:00:00:00:00:00:00:00
  No Pbay on fabric5
  Vendor HITACHI, product OPEN-9, revision 5247, serial # R300-76F0-0011
   capacity 7042 Mbytes


# dumpconf |grep sd26
sd26         sd           26  0x71fbf800  SM     fcp      fabric5
sd26          sd            26  0x31fbf800  SM      fcp       fabric6


# infodev -a sd26
Vendor ID                      : HITACHI
Product ID                     : OPEN-9
Revision Level                 : 5247
Device Class                   : disk
Capacity (in blocks)           : 14423040
Block Size                     : 512
Fibre Channel Address          : 0xef
FCP LUN                        : 00:00:00:00:00:00:00:00
Located on                     : fabric5
```

Figure 5-7.
Sample dumpconf, infodev, and diskid Output for Arbitrated Loop (FC-AL) Connections.

```
# diskid sd12
sd12: scsibus8 cport 0 target 9 lun 0, No Pbay on scsibus8
  Vendor HITACHI, product OPEN-3, revision 0316, serial # R300-8A2C-0049
  capacity 2347 Mbytes
```

RAID300 unit
UUID of RAID300 unit
LDEV number

This **sd** device has 2 SCSI paths to the LDEV

```
# dumpconf -d|grep sd12
sd12         sd           12  0x00000090  SM     scsi     scsibus8
sd12         sd           12  0x20000090  SM     scsi     scsibus9
```

CHS#
Port # on CHS
SCSI LUN
SCSI Target ID

diskid report for LUSE-configured volumes

```
# diskid sd6
sd6: scsibus9 cport 18 target 0 lun 0, No Pbay on scsibus9
  Vendor HITACHI, product OPEN-9*7, revision 5235, serial # R300-76F0-0028
  capacity 49297 Mbytes
```

Concatenated capacity
Seven OPEN-9 LDEVs concatenated

Figure 5-8.
Sample dumpconf, infodev, and diskid output for SCSI connection.

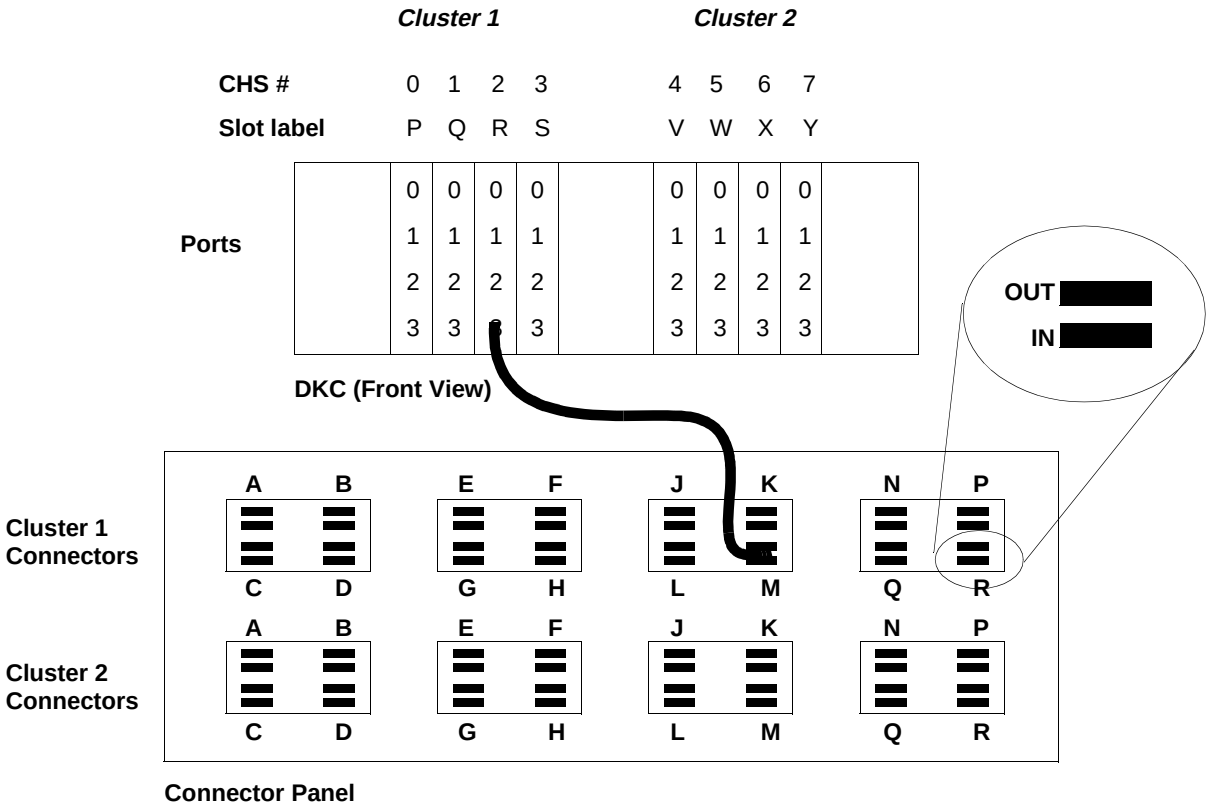

NOTE: The association of ports to connectors depends on the CHS, according to the following scheme:

Port	Connectors for CHSs 0, 4	Connectors for CHSs 1, 5	Connectors for CHSs 2, 6	Connectors for CHSs 3, 7
0	A	E	J	N
1	B	F	K	P
2	C	G	L	Q
3	D	H	M	R

Figure 5-9.
CHS, Port Mappings for SCSI Connection.

Table 5-3
CHS/CHF-to-Port Mappings

	Cluster 1				Cluster 2			
CHS/CHF #	**0**	**1**	**2**	**3**	**4**	**5**	**6**	**7**
Slot Label	**P**	**Q**	**R**	**S**	**V**	**W**	**X**	**Y**
Port 0								
Port Name	A	E	J	N	A	E	J	N
cport	O	16	32	48	64	80	96	112
Port 1								
Port Name	B	F	K	P	B	F	K	P
cport	1	17	33	49	65	81	97	113
Port2								
Port Name	C	G	L	Q	C	G	L	Q
cport	2	18	34	50	66	82	98	114
Port 3								
Port Name	D	H	M	R	D	H	M	R
cport	3	19	35	51	67	83	99	115

6. Once the devices have been verified, perform I/O to each device to verify its operation.

7. Complete the preparation of the RAID subsystem for use.

   - If you are using VTOCs, refer to the discussion titled *Notes on Usage and Administration*, at the end of this chapter, for details on the VTOCs supplied for RAID300 devices.

   - If you are using ptx/SVM to manage disk resources, refer to the ptx/SVM documentation supplied with the DYNIX/ptx host system.

8. Bring the system back to multi-user mode and notify users that it is now available for use.

Notes on Usage and Administration

Obtaining Subsystem Status

You can use the subsystem's SVP program, available on the subsystem's built-in laptop, to obtain status on all major subsystem modules, including the location of malfunctioning components. You can also examine the configuration of LDEVs (logical devices) and assignments of these to ports and SCSI IDs/LUNs.

1. On the laptop, bring up the SVP program.
2. From the SVP main menu bar, select the `Maintenance` icon.
3. On the Subsystem display, click the icons corresponding to the subsystem modules or to the kind of status information you want to see. Some of the icons may be grayed out. However, the items listed in Table 5-3 should be available.

The icons labeled **DKC** and **DKU-*xx*** represent the subsystem cabinets:

- **DKC** represents the cabinet that contains the controllers, SCSI ports to the host, etc.
- **DKU-L1**, **DKU-L2** represent the left cabinets, containing individual physical disk modules.
- **DKU-R1**, **DKU-R2** represent the right cabinets, containing individual physical disk modules.

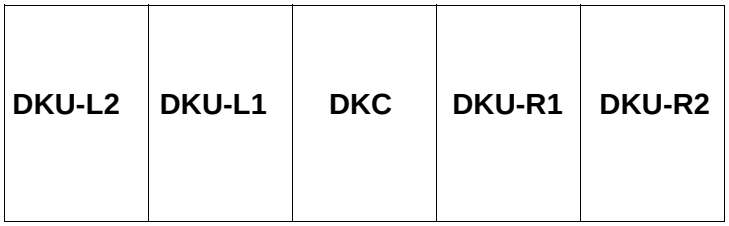

Clicking on either of these icons will display a screen showing the major functional submodules at that level. The layout of icons on the screen will *approximate* the physical layout of physical submodules in the current module. The lowest level screens will show the status of individual functional modules.

At each level, if a component or some sub-component is malfunctioning, the corresponding icon will blink. Continue clicking on blinking icons until you locate the specific component that is malfunctioning.

The other icons will display configuration status information, as indicated in Table 5-3.

Table 5-4
Status Icons and Contents

Icon	Information Displayed
DKC	Displays modules and status information for the left-hand subsystem cabinet. This cabinet contains the SCSI controllers, ports for connection to the host system, etc.
DKU-*xx*	Displays modules and status information for one of the subsystem cabinets that contain the DASs, individual disk modules, and redundancy components for the DASs (fans, power supplies, etc.). DKU-R1 represents the first cabinet to the right of the DKC. DKU-R2 represents the second cabinet to the right of the DKC. DKU-L1 represents the first cabinet to the left of the DKC. DKU-L2 represents the second cabinet to the left of the DKC.
Copy Status	Not available.
Logical Device	Displays the status of logical devices (LDEVs): device status, RAID group membership, physical devices associated with the LDEV, etc.
RAID Status	Displays the status of RAID groups and the individual disk modules within a RAID group: status, disk type, emulation mode.
Version	Displays firmware version information for each component in the subsystem that contains firmware.
LCP/MCP Path	Display path information for logical and physical paths.
SCSI Path	Displays SCSI path information: assignments of ports, SCSI IDs, and LUNs to LDEVs. Also shows available SCSI IDs and LDEVs for each port.

Typical ktlog Output Showing Route Failure and Recovery

Failure and recovery of a route will result in messages to **ktlog**. The following output is a typical sequence of messages show a route going down and coming back up.

```
3a9d7a12 14:22:10 note q1/e6/p2792 ff2: Node 0x210213 link down
3a9d7a12 14:22:10 tolog/warn q1/e7/p3156 sd289(port 1): disabling route
3a9d7a12 14:22:10 tolog/warn q1/e7/p3156 sd289(port 1): soft error on WRITE (unk
nown error code) buffer 0x2a632000 on route 0x328114a0
3a9d7a12 14:22:10 warn q1/e7/p3156 sd289(port 1): buffer 0x2a632000 being retrie
d on route 0x31b899c0 - sd289(port 65)

3a9d7aa5 14:24:37 note q1/e7/p3163 ff2: Node 0x210213 link up
3a9d7aa7 14:24:39 note q1/e7/p3780 ff2: Node 0x210313 link up
3a9d7aa9 14:24:41 note q1/e6/p3023 ff2: Node 0x210413 link up
3a9d7ab2 14:24:50 note q1/e5/p3174 ff2: Node 0x210513 link up
3a9d7ab3 14:24:51 note q1/e4 ff2: Node 0x210613 link up
3a9d7ab3 14:24:51 note q1/e4/p3774 ff2: Node 0x210713 link up

3a9d7ab4 14:24:52 tolog/warn q0/e3 sd289(port 1): resuming route
```

Sample VTOCs Supplied With ptx/HITACHI-RAIDMP

The following VTOCs are supplied for use with ptx/HITACHI-RAIDMP.

open-3	For OPEN-3 disk modules
open-8	For OPEN-8 disk modules
open-9	For OPEN-9 disk modules

RAID200/RAID300, one-slice disk

open-3.vxsl	OPEN-3 disk modules
open-8.vxsl	OPEN-8 disk modules
open-9.vxsl	OPEN-9 disk modules

Refer to the *ptx/SVM Administration Guide* for detailed information on ptx/SVM.

The make_vtoc(1M) Utility

The **make_vtoc**(1M) utility allows you to define custom VTOCs. Refer to the on-line man page or to the hard copy of the man page in Appendix A of this manual for details. Refer to the DYNIX/ptx **mkvtoc** man page for additional details on VTOC creation.

LUN Security

LUN Security Overview

The RAID300 LUN Security feature enables a fibre-connected RAID300 subsystem to be configured so that particular logical volumes (LUNs) can *only* be accessed by one or more specified Fibre Channel host adapters (FCHAs). An FCHA is identified to the RAID300 subsystem via its *worldwide name* or *WWN*, a 16-digit HEX number that uniquely identifies the FCHA. Figure 5-9 shows an example configuration; in this figure, the FCHAs in the IBM NUMA-Q system and those in the NT system are connected to the RAID300 subsystem through the same two switches. Although they share fiber links to the RAID300 subsystem, FCHAs in the IBM NUMA-Q system cannot see LDEVs assigned to the WWNs of the NT system, nor can FCHAs in the NT system see LDEVs assigned to the WWNs of the FCHAs in the IBM NUMA-Q system.

Details on setting up LUN security are covered in two Hitachi documents:

DKC3101 Disk Subsystem Remote Console Program Reference Manual
Hitachi Freedom Storage™ 7700E LUN Manager User's Guide

Basic procedures for setting up LUN Security on RAID300 subsystems with fibre connection to IBM NUMA-Q systems are covered earlier in this chapter, under *Step C: Configure the RAID300 Subsystem, Configuration for Fibre Channel Connection, Completing the Configuration Process for Fibre Channel Connection.*

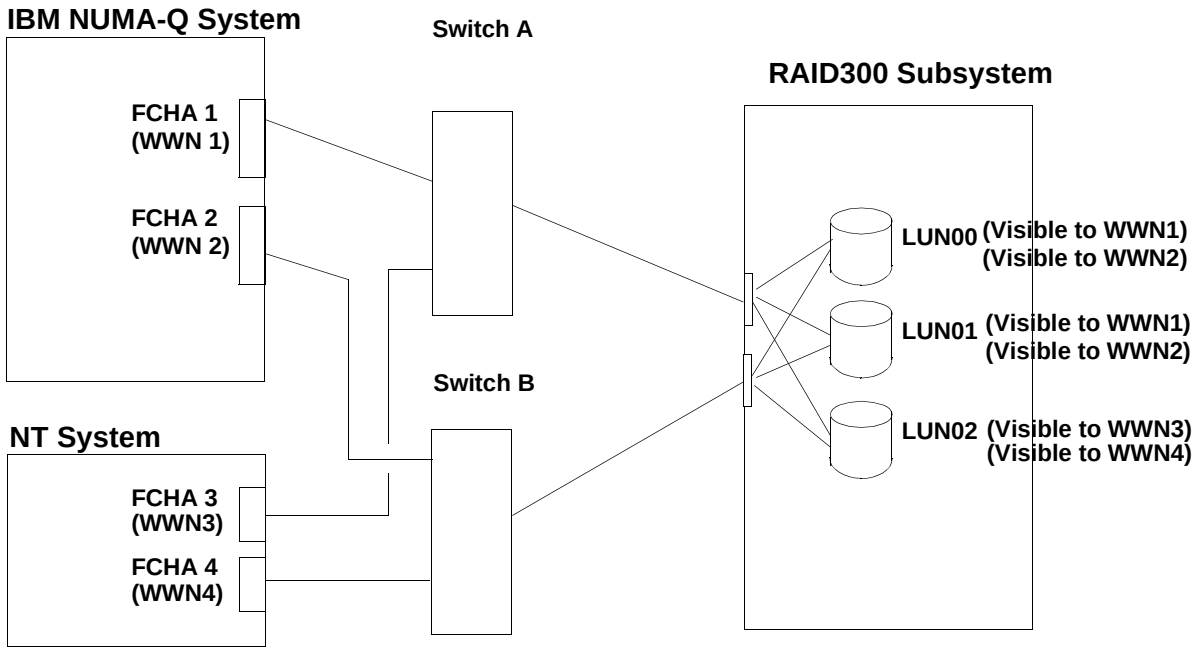

Figure 5-10.
Example Connection Scheme with LUN Security Feature.

Getting Ready to Use LUN Security

To use LUN security, you must do the following:

1. Install LUN Security using the instructions provided by Hitachi.

2. On the host system, identify the *worldwide names* of the Fibre Channel host adapters (FCHAs) that will be part of the LUN security scheme. To identify the FCHA worldwide names, you can do either of the following:

- Use the **dumpconf** or **showcfg** command to identify the FCHAs in this host system (FCHAs currently used in IBM NUMA-Q systems will be identified as "FF" devices). Then enter **infodev -a <*FCHA_name*>** for each FCHA to be included in the LUN security scheme; e.g., for FCHA ff1, you will see output similar to the following:

```
/etc/infodev -a ff1
Vendor ID                       : Emulex
Product ID                      : Firefly
Revision Level                  : FireFly-3
Device Class                    : ctlr
Node WWN                        : 10:00:00:00:c9:20:58:4f
Port WWN                        : 10:00:00:00:c9:20:58:4f
Located on                      : quad0
```

  Here, the Port WWN is 10:00:00:00:c9:20:58:4f. Note that WWNs for the Emulex adapters used in IBM NUMA-Q systems begin with the digits "10"; WWNs beginning with "50" indicate RAID300 subsystem adapters.

  ### CAUTIONS

  *With some host adapter firmware versions, the Node WWN may differ in appearance from the Port WWN; **be sure** to use the **Port WWN** to set up LUN security.*

  *In addition, in a LUN security configuration, the WWNs of the host system FCHAs that are registered in the array's controller ports must match **exactly** the WWNs of the probing FCHAs on the host system. If there is any mismatch, DYNIX/ptx will not recognize **any** logical units in that LUN security configuration.*

- Alternatively, you can **telnet** into the appropriate switch, and use the **switchShow** command, whose output includes port names *and* their worldwide names.

```
admin> switchShow
switchName:     switch
switchType:     3.2
switchState:    Online
switchRole:     Principal
switchDomain:   1
switchId:       fffc41
switchWwn:      10:00:00:60:69:20:15:a0
port  0: sw  No_Light
port  1: sw  Online     F-Port  10:00:00:00:c9:20:d2:26
port  2: sw  Online     F-Port  50:00:0e:10:00:00:54:52
port  3: sw  Online     F-Port  50:00:0e:10:00:00:47:67
port  4: sw  No_Light
port  5: sw  No_Light
port  6: sw  Online     F-Port  21:00:00:e0:8b:01:0d:bf
```

```
port  7: sw  No_Light
switch:admin>
```

3. Write down the WWN for each port that will be used for LUN security. You will need to enter these WWNs on the appropriate RAID300 configuration screen. For details on RAID300 configuration, refer back to *Step C: Configure the RAID300 Subsystem*, *Configuration for Fibre Channel Connection*, earlier in this chapter.

Chapter 6

RAID400 Subsystem Integration

This chapter describes the steps needed to integrate a Hitachi RAID400 RAID disk array subsystem with a DYNIX/ptx host system running ptx/HITACHI-RAIDMP. The following topics are included:

- **Requirements for use.** Gives hardware requirements for the use of this subsystem with ptx/HITACHI-RAIDMP.
- **Related Documentation.** Lists the documentation that you will find useful in completing the integration tasks described in this chapter.
- **Restrictions.** Includes restrictions specific to the functioning and use of this subsystem with ptx/HITACHI-RAIDMP.
- **RAID400 RAID concepts.** Gives a brief summary of the RAID terminology and concepts used in association with this subsystem.
- **RAID400 subsystem integration overview.** Gives a high-level overview of the steps required to integrate this subsystem with the host system.
- **Integration procedures.** Includes instructions for ensuring that the subsystem has the required configuration settings and for installing and connecting the subsystem to the host system.
- **Notes on Use and Administration.** Includes information that may be useful in the ongoing administration and use of the subsystem.
- **LUN Security.** Provides a high-level overview of the RAID400 LUN Security feature, from the viewpoint of the DYNIX/ptx host system.

RAID400 Requirements for Use

The following are required to integrate a RAID400 subsystem with a DYNIX/ptx host system and this release of ptx/HITACHI-RAIDMP.

DYNIX/ptx System Requirements

The following are the DYNIX/ptx system requirements for integration of the RAID400 subsystem:

- The DYNIX/ptx host system must meet the requirements listed in the software chapter of this manual.
- If connection is via FC-SW Fibre Channel topology, the IBM NUMA-Q system must have enough ports available on the Fibre Channel switch for dedicated connection of the RAID400 devices.
- If connection is via direct connect FC-AL Fibre Channel topology, the IBM NUMA-Q system is restricted to a **maximum** of two (2) quads, with a **maximum** of two (2) paths to each subsystem LDEV (each path through a single FCHA in each quad).
- Fastpatch FP# 254922 is required for connection and operation of RAID400 subsystems.

RAID400 Subsystem Requirements

The following are the integration requirements for the RAID400 subsystem itself:

- The subsystem must be running a minimum microcode revision of 01-11-23-00/04; contact the IBM NUMA-Q Customer Support Center for information on the latest supported microcode release(s).
- Multiport connectivity to the RAID400 subsystem requires that certain RAID400 modes be set; if these modes are not set properly, RAID400 subsystem performance in a multiport configuration will be severely degraded. *These settings can only be changed by Hitachi field service personnel.* If you are connecting the RAID400 subsystem in a multiport configuration, *contact your Hitachi distributor's service representative* to request that these modes be set properly.

RAID400 Host and System Modes

Host Modes. On RAID400 subsystems, each port is assigned a *host mode*, depending on the type of connection to the host system.

- Ports connected to DYNIX/ptx systems only require Host Mode 04, the *Sequent* or *Numaq_Unix* mode. This specifies fast/wide SCSI and multiport.
- Ports connected to NT® systems only require Host Mode 00.
- Ports connected to both DYNIX/ptx and NT systems (e.g., through a common switch) require Host Mode 04.

System Modes. System mode settings are specified for a given Host Mode, and are used for control throughout the entire subsystem.

> **NOTE**
> *System modes are dependent on the correct host mode being set. For example, if Host Mode 04 is not set for a particular port, System Mode 56 is ignored on that port. First make sure that the correct host mode is set* ***for each port****, then set the system modes.*

The following explains the main functionality of System Modes 56 and 57.

- System Mode 56 improves multiport sequential READ / WRITE performance on multiport configurations. It is only enabled for host connections with Host Mode 04.
- System Mode 57 enables RAID400 FRU error reporting through a SCSI UNIT ATTENTION. An I / O command will be aborted (sense key 0xB) with ASC / ASCQ 0xC300. The driver reports this as “internal FRU / power failure”. A subsequent retry of the I / O will complete successfully. There should be one of these messages for each LDEV that is configured on the system.

Required Modes. The following table summarizes the required Host and System Modes.

Table 6-1
Host and System Modes

Operating System	Host Mode (by port)	System Modes
DYNIX/ptx	04	56 & 57
NT only	00	57
DYNIX/ptx & NT	04	56 & 57

Fibre Cabling Requirements

Use only the standard fibre cabling available from IBM or Hitachi.

Related Documentation

The following Hitachi documentation will be especially useful in completing the integration tasks described in this chapter. You should have these manuals handy at all times.

- *Hitachi Freedom Storage™ Lightning 9900™ User and Reference Guide*
- *Hitachi Freedom Storage™ Lightning 9900™ LUN Manager User Guide*
- *Hitachi Freedom Storage™ Lightning 9900™ Remote Console User Guide*

RAID400 Restrictions

The following restrictions apply to the connection and operation of a RAID400 subsystem with a DYNIX/ptx host system running this release of ptx/HITACHI-RAIDMP. These are in addition to the general and subsystem-specific restrictions listed in Chapter 2 of this document.

- Only the host systems and connectivities listed earlier in this document are supported.
- Single-port and multiport configurations are supported.
- RAID400 disks cannot be used as boot or dump devices.
- Only switched fabric (FC-SW) and direct connect arbitrated loop (FC-AL) fibre channel connections are supported. SCSI connections and FL port arbitrated loop connections *are not supported*.
- There is no clusters support for direct connect FC-AL connections.
- The only FC switches supported are those listed in Chapter 1 of this manual, under the discussion titled *FC Switch Support and Operation*.
- No connections through any other switches or through hubs are supported.
- Only FC multimode cables supplied by Hitachi or IBM are supported.
- For fibre connections, up to eight (8) paths are supported.

RAID400 RAID Concepts

The following discussions give an overview of RAID400 terminology and concepts.

Terminology

The following terms are used in the remainder of this document to describe the Hitachi RAID400 subsystem and software.

Table 6-2
RAID400 Terminology

CHF	Channel Adapter for FC. Each CHF supports up to four (4) FC connections.
DKA	Disk Adapter. The interface between DKC cache and the internal SCSI disk drives.
DKC	Disk Controller. The main chassis that houses the CHA, CHF, cache, and SVP laptop.
DKU	Disk Unit. A frame that houses the SCSI disk drives.
frame	Physical chassis that houses one or more hardware components.
LDEV	Logical Device. This is a 3390-compatible or OPEN-compatible volume, in terms of capacity. Each LDEV is assigned and addressable as a SCSI target ID/LUN.
LUSE	LUN Size Expansion. Allows concatenation of multiple volumes into one logical unit.
SVP laptop	A built-in laptop computer that enables configuration and monitoring of the subsystem.
CHIP	Channel Host Interface Processor. Connects fibre ports to the system.
ACP	Array Control Processor. Interface from system to disk drives.

Multiported Devices

Certain disk storage units provide multiple physical connections (ports) to multiple I/O buses and then configure their logical volumes to allow access from those ports simultaneously. This feature is called *multiporting*.

Multiporting provides redundant access to individual logical volumes, increasing their availability. If a particular port or I/O bus becomes unavailable, an unaffected port and/or I/O bus can be used to access the logical volume. Locating multiple ports for a logical volume on separate buses provides protection from both an I/O bus failure and the failure of a port within a storage unit. Locating multiple ports for a logical volume on the same I/O bus is legal, but this method does not provide protection from the failure of that I/O bus. With multiporting:

- A port can have more than one logical volume connected to it.
- A single logical volume can be connected to more than one port. E.g., in Figure 6-1, the disk resource is connected to Port 1A and Port 2A, and has a path through both fabrics (*fabric0* or *fabric1*).

The operating system also takes advantage of multiporting to better distribute its I/O load. Doing so *may* reduce the latency of individual I/O requests (although this is not guaranteed).

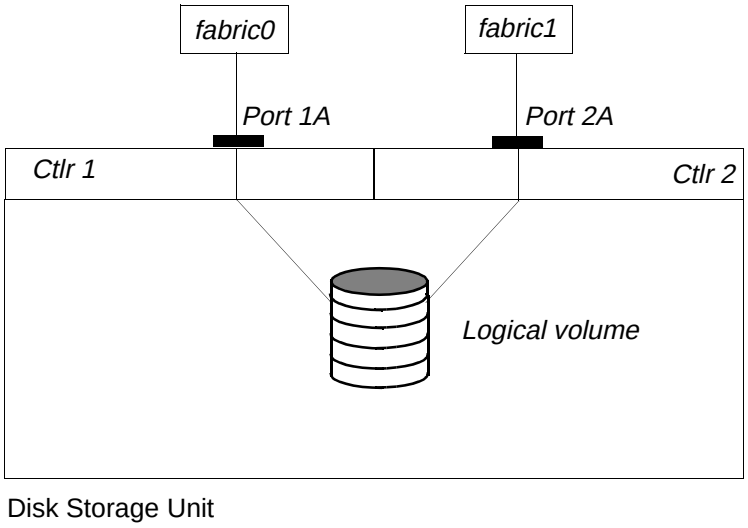

Figure 6-1.
Sample Multiported Disk Storage Unit.

RAID400 Integration Overview

To integrate the RAID400 subsystem with the DYNIX/ptx system, you must complete the following steps.

- Step A: Install the ptx/HITACHI-RAIDMP software.
- Step B: Prepare the RAID subsystem for integration.
- Step C: Check or adjust the subsystem's configuration settings.
- Step D: Make cable connections between the host system and the subsystem.
- Step E: Verify the integration.

The following discussions cover these steps in detail or refer to other documentation for details.

Step A: Install ptx/HITACHI-RAIDMP

If you have not already installed the current release of ptx/HITACHI-RAIDMP, refer to Chapter 2 of this manual and install the software at this time.

Step B: Prepare for Integration

Hitachi or Hitachi distributor personnel will unpack, install, and prepare the subsystem for integration. After the subsystem has been physically installed at the location where it will be used, do the following to prepare for the integration procedures:

1. Collect all cables and terminators, etc. that will be required.

2. Collect all required documentation.

Step C: Configure the RAID400 Subsystem

CAUTIONS

The procedures referenced in the following discussion assume that service personnel for Hitachi or a Hitachi distributor have first configured and formatted RAID groups on the subsystem. Only qualified Hitachi or Hitachi distributor personnel should attempt to configure or format RAID groups on the subsystem.

Incorrect configuration of the subsystem could lead to data corruption or loss. ***Before any subsystem configuration*** *is done, all personnel involved in the installation should carefully define and review the RAID configuration for this site.*

Do not attempt to change the subsystem's configuration while user activity is occurring.

Configuration Overview

The following general tasks must be completed to configure the RAID subsystem for use with ptx/HITACHI-RAIDMP:

1. Select LDEVs on each port (single or multipath).
2. Make the correct port settings for FC-SW or FC-AL operation.
3. Set the host mode.
4. If you are using LUN security, set up LUN security for all ports/LUNs involved. Refer to the discussion later in this chapter titled *LUN Security* for an overview of the LUN security feature from an IBM NUMA-Q host system point of view.

Procedures to accomplish these tasks are documented in detail in the subsystem's Remote Console and LUN Security user guides; see the discussion titled *Related Documentation*, earlier in this chapter.

Configuration Notes

The following specific notes or cautions should be kept in mind when configuring the RAID subsystem.

NOTE

For Fibre Channel configurations (FC-SW and FC-AL), the SCSI ID will always be ***F****.*

Step D: Make Cable Connections

When the software installation has been completed and the RAID subsystem has been configured, do the following:

1. If you are not already logged in as **root**, do so at this time.

2. Warn all users that the system will be coming down, and make sure that all user activity has stopped.

> **NOTE**
> *ptx/HITACHI-RAIDMP supports switched fabric (FC-SW) and direct connect arbitrated loop (FC-AL) Fibre Channel connections of RAID400 subsystems to IBM NUMA-Q host systems. Be sure to choose the correct cabling diagram, below.*

> **NOTE**
> *The following figures are intended to provide cabling guidelines only. You will need to make appropriate modifications in cabling to match this site's installation and subsystem configuration.*

3. Refer to the appropriate cabling diagram below and make cable connections between the subsystem and the DYNIX/ptx host system.

- Figure 6-2 shows cabling for a typical switched fabric (FC-SW) Fibre Channel connection to an IBM NUMA-Q host system.
- Figure 6-3 shows cabling for a typical direct connect arbitrated loop (FC-AL) Fibre Channel connection to an IBM NUMA-Q host system.

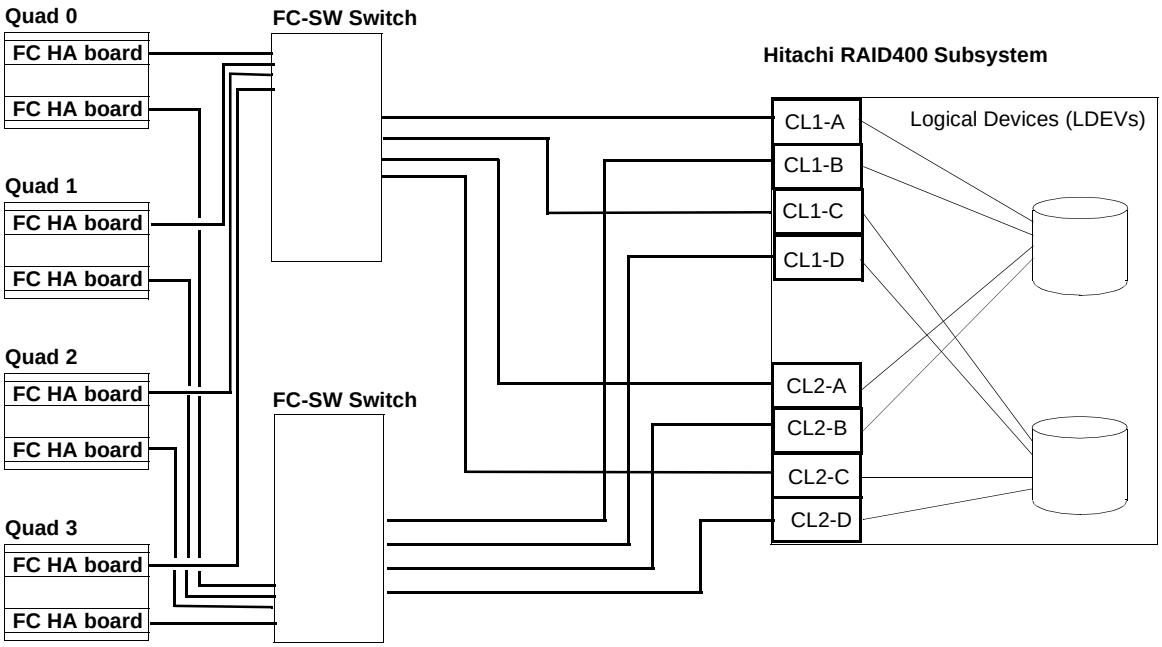

Figure 6-2.
Switched Fabric (FC-SW) Connection to an IBM NUMA-Q System.

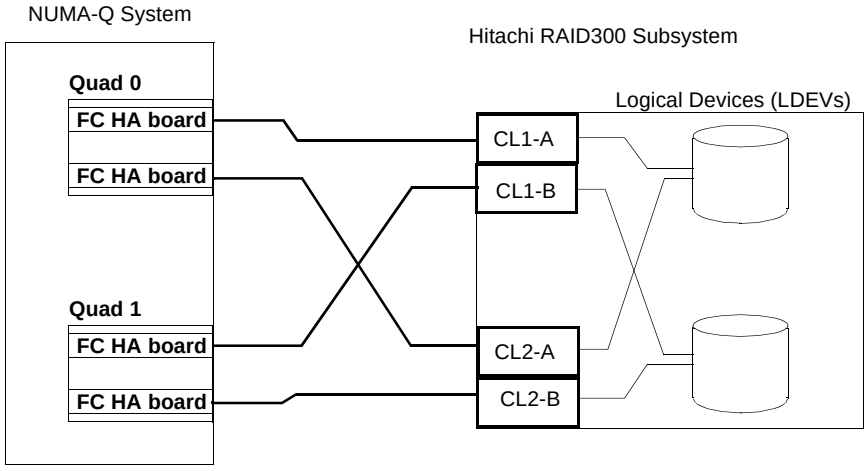

Figure 6-3.
Arbitrated Loop (FC-AL) Connection to an IBM NUMA-Q System.

Step E: Verify the Installation

To verify the installation:

1. If this is **not** an arbitrated loop (FC-AL) connection, skip this step and go directly to Step 2, below.

If this **is** an arbitrated loop (FC-AL) connection, you must configure the DYNIX/ptx host system for arbitrated loop communications. To do this, at the DYNIX/ptx **root** prompt:

a. If you are not already logged in as **root**, do so at this time.

b. Change (**cd**) to the directory */usr/conf/uts/io/ff*.

c. Open the file *ff_space.c* for editing.

d. In this file, find the line that reads:

```
ff_topology=FC_TOPOLOGY_FABRIC
```

and change this line to read:

```
ff_topology=FC_TOPOLOGY_AL
```

e. Save and exit the file.

f. Rebuild the kernel.

This sets up the DYNIX/ptx system for arbitrated loop communications.

2. Reboot the DYNIX/ptx system.

3. When the system has rebooted, enter:

/etc/dumpconf -d | grep sd

and verify that the appropriate **sd** devices were found at system boot. You should see one entry for each data path to each Hitachi **sd** device. Each entry for the same Hitachi **sd** device will have a unique SCSI bus assignment.

4. For each **sd** device, use the **diskid** command to determine which **sd** devices are Hitachi devices:

diskid sd<*n*>

where <*n*> is the id number of the device. Hitachi devices should appear as one of the following:

OPEN-3 (2.4GB capacity)
OPEN-8 (7.3GB capacity)
OPEN-9 (7.3GB capacity)
3390-3A (2.8GB capacity)
3390-3B (2.8GB capacity)
3390-3C (2.8GB capacity)

3380-KA (1.8GB capacity)
3380-KB (1.8GB capacity)
3380-KC(1.8GB capacity)
OPEN-K (1.8GB capacity)
OPEN-E (13.9GB capacity
6586-KA
6586-KB
6586-KC
6588-3A
6588-3B
6588-3C

5. At this point you will want to determine path configuration information for the subsystem. You can do this using the **dumpconf**, **diskid**, and **infodev** utilities.

NOTE on Reporting of LUSE Expanded Devices

*Volumes configured with LUSE (LUN Size Expansion) will be reported as **device*number**, where device is the underlying (non-expanded) device type and number is the number of configured volumes. For expanded volumes, the volume capacity is reported as the capacity of the concatenated volume. For example, seven OPEN-9 devices would be reported with device id **OPEN-9*7** and a capacity of 49297.*

NOTE

*The **dumpconf** and **infodev** commands report unit information in hex; the **diskid** command reports the **cport** information in decimal. You will need to translate between the output of these commands.*

The following figures and table summarize configuration information for a typical RAID400 subsystem:

Figure 6-4 shows sample output from **dumpconf**, **infodev**, and **diskid** for a typical Fibre Channel connection. Output for FC-SW and FC-AL connections will be similar.

Figure 6-5 shows sample output from the **hrutil** utility that reports device information for RAID400 subsystems.

Table 6-3 shows CHF-to-Port mappings for Fibre Channel connections.

```
# diskid sd101
sd101: fabric7 cport 18 fc id 0x200213, fcp lun 00:10:00:00:00:00:00:00
  No Pbay on fabric7
  Vendor HITACHI, product OPEN-E, revision 0111, serial # R400-787F-0096
  capacity 13893 Mbytes

# dumpconf |grep sd101
sd101        sd          101  0x12fe1010  SM     fcp      fabric7
sd101        sd          101  0x42feb810  SM     fcp      fabric7
sd101        sd          101  0x02fe7810  SM     fcp      fabric2
sd101        sd          101  0x52fec010  SM     fcp      fabric2

# infodev -a sd101
Vendor ID                      : HITACHI
Product ID                     : OPEN-E
Revision Level                 : 0111
Device Class                   : disk
Capacity (in blocks)           : 28452960
Block Size                     : 512
Fibre Channel Address          : 0x200213
FCP LUN                               : 00:10:00:00:00:00:00:00
Located on                     : fabric7
```

Figure 6-4.
Sample dumpconf, infodev, and diskid Output for Fibre Channel (FC-SW or FC-AL) Connections.

```
elm9a41>hrutil -ia sd101
hrutil: Inquiry page of sd101
--------------------------------
Logical Unit sd101 : Defined.
Media cannot be removed.
DKC does not comply with the ISO version.
DKC does not conform to the ECMA (Computer Manufacturer Association Version)
version.
ANSI Version: 0x02 (SCSI)
DKC Does not have AENC (Asynchronous Event Notification Capability.
DKC Does not support a TERMINATE I/O PROCESS message.
DKC complies with the NSI SCSI-3 specification.
Additional length requested: 63
DKC Does not support Relative Address.
DKC Does not support 32-bit data transfer.
DKC supports 16-bit data transfer.
DKC supports synchronous transfer.
DKC supports linked commands.
DKC supports tagged command queuing.
DKC uses only HARD RESET.
Vendor: HITACHI
Product: OPEN-E
Product Revision Level: 0111
Model: RAID400 HA
Product Serial Number: 30847 (0x787F)
LDEV Number: 0x0096 (150)
VolStat of CA: SMPL
VolStat of BC0(MU#0): SMPL
VolStat of BC1(MU#1): SMPL
VolStat of BC2(MU#2): SMPL
Position of LDV in RAID Group: 2-1
        RAID Group No. in DKC: 1
        Sub number of LDEV in RAID Group:2
RAID Level: RAID-5
Sub System ID of DKC: 4
CU number of DKC: 0
LDEV is not copy paired.
```

Figure 6-5.
Sample hrutil Output.

Table 6-3
CHF-to-Port Mappings

	Cluster 1				Cluster 2			
CHS/CHF #	**0**	**1**	**2**	**3**	**4**	**5**	**6**	**7**
Slot Label	**P**	**Q**	**R**	**S**	**V**	**W**	**X**	**Y**
Port 0								
Port Name	A	E	J	N	A	E	J	N
cport	O	16	32	48	64	80	96	112
Port 1								
Port Name	B	F	K	P	B	F	K	P
cport	1	17	33	49	65	81	97	113
Port2								
Port Name	C	G	L	Q	C	G	L	Q
cport	2	18	34	50	66	82	98	114
Port 3								
Port Name	D	H	M	R	D	H	M	R
cport	3	19	35	51	67	83	99	115

6. Once the devices have been verified, perform I/O to each device to verify its operation.

7. Complete the preparation of the RAID subsystem for use.

   - If you are using VTOCs, refer to the discussion titled *Notes on Use and Administration*, at the end of this chapter, for details on the VTOCs supplied for RAID400 devices.
   - If you are using ptx/SVM to manage disk resources, refer to the ptx/SVM documentation supplied with the DYNIX/ptx host system.
   - If you are using switched fabric (FC-SW) connection, check that all ports have logged in successfully. To do this, log into the switch and use the **switchShow** command to verify the connections.
   - If you are using the RAID400 subsystem's LUN Security feature, refer to the discussion titled *LUN Security*, later in this chapter, for an overview of LUN security from the viewpoint of the DYNIX/ptx host system. Also refer to Hitachi's *LUN Manager User's Guide* for details on setting up LUN security.

8. Bring the DYNIX/ptx system back to multi-user mode and notify users that it is now available for use.

Notes on Use and Administration

Obtaining Subsystem Status

You can use the subsystem's SVP program, available on the subsystem's built-in laptop, to obtain status on all major subsystem modules, including the location of malfunctioning components. You can also examine the configuration of LDEVs (logical devices) and assignments of these to ports and SCSI IDs/LUNs.

1. On the laptop, bring up the SVP program.
2. From the SVP main menu bar, select the `Maintenance` icon.
3. On the Subsystem display, click the icons corresponding to the subsystem modules or to the kind of status information you want to see. Some of the icons may be grayed out. However, the items listed in Table 6-4 should be available.

The icons labeled **DKC** and **DKU-*xx*** represent the subsystem cabinets:

- **DKC** represents the cabinet that contains the controllers, SCSI ports to the host, etc.
- **DKU-L1**, **DKU-L2** represent the left cabinets, containing individual physical disk modules.
- **DKU-R1**, **DKU-R2** represent the right cabinets, containing individual physical disk modules.

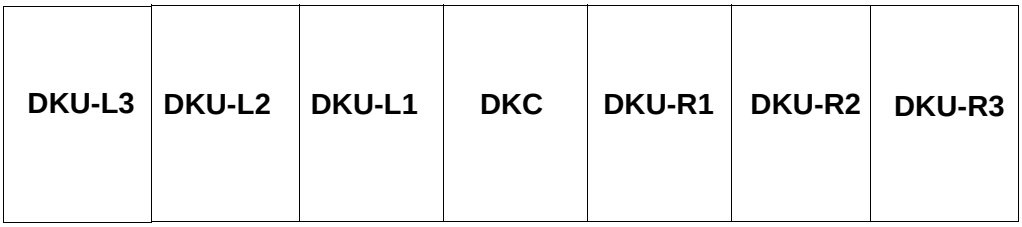

Clicking on either of these icons will display a screen showing the major functional submodules at that level. The layout of icons on the screen will *approximate* the physical layout of physical submodules in the current module. The lowest level screens will show the status of individual functional modules.

At each level, if a component or some sub-component is malfunctioning, the corresponding icon will blink. Continue clicking on blinking icons until you locate the specific component that is malfunctioning.

The other icons will display configuration status information, as indicated in Table 6-4.

Table 6-4
Status Icons and Contents

Icon	Information Displayed
DKC	Displays modules and status information for the left-hand subsystem cabinet. This cabinet contains the SCSI controllers, ports for connection to the host system, etc.
DKU-*xx*	Displays modules and status information for one of the subsystem cabinets that contain the DASs, individual disk modules, and redundancy components for the DASs (fans, power supplies, etc.). DKU-R1 represents the first cabinet to the right of the DKC. DKU-R2 represents the second cabinet to the right of the DKC. and so on.
Copy Status	Not available.
Logical Device	Displays the status of logical devices (LDEVs): device status, RAID group membership, physical devices associated with the LDEV, etc.
RAID Status	Displays the status of RAID groups and the individual disk modules within a RAID group: status, disk type, emulation mode.
Version	Displays firmware version information for each component in the subsystem that contains firmware.
LCP/MCP Path	Display path information for logical and physical paths.
SCSI Path	Displays SCSI path information: assignments of ports, SCSI IDs, and LUNs to LDEVs. Also shows available SCSI IDs and LDEVs for each port.

Typical ktlog Output Showing Route Failure and Recovery

Failure and recovery of a route will result in messages to **ktlog**. The following output is a typical sequence of messages show a route going down and coming back up.

```
3a9d7a12 14:22:10 note q1/e6/p2792 ff2: Node 0x210213 link down
3a9d7a12 14:22:10 tolog/warn q1/e7/p3156 sd289(port 1): disabling route
3a9d7a12 14:22:10 tolog/warn q1/e7/p3156 sd289(port 1): soft error on WRITE (unk
nown error code) buffer 0x2a632000 on route 0x328114a0
3a9d7a12 14:22:10 warn q1/e7/p3156 sd289(port 1): buffer 0x2a632000 being retrie
d on route 0x31b899c0 - sd289(port 65)

3a9d7aa5 14:24:37 note q1/e7/p3163 ff2: Node 0x210213 link up
3a9d7aa7 14:24:39 note q1/e7/p3780 ff2: Node 0x210313 link up
3a9d7aa9 14:24:41 note q1/e6/p3023 ff2: Node 0x210413 link up
3a9d7ab2 14:24:50 note q1/e5/p3174 ff2: Node 0x210513 link up
3a9d7ab3 14:24:51 note q1/e4 ff2: Node 0x210613 link up
3a9d7ab3 14:24:51 note q1/e4/p3774 ff2: Node 0x210713 link up

3a9d7ab4 14:24:52 tolog/warn q0/e3 sd289(port 1): resuming route
```

Sample VTOCs Supplied With ptx/HITACHI-RAIDMP

The following VTOCs are supplied for use with ptx/HITACHI-RAIDMP.

open-3	For OPEN-3 disk modules
open-8	For OPEN-8 disk modules
open-9	For OPEN-9 disk modules
open-e	For OPEN-E disk modules

RAID200/RAID400, one-slice disk

open-3.vxsl	OPEN-3 disk modules
open-8.vxsl	OPEN-8 disk modules
open-9.vxsl	OPEN-9 disk modules
open-e.vxsl	OPEN-E disk modules

Refer to the *ptx/SVM Administration Guide* for detailed information on ptx/SVM.

The make_vtoc(1M) Utility

The **make_vtoc**(1M) utility allows you to define custom VTOCs. Refer to the on-line man page for details. Refer to the DYNIX/ptx **mkvtoc** man page for additional details on VTOC creation.

LUN Security

LUN Security Overview

The RAID400 LUN Security feature enables a RAID400 subsystem to be configured so that particular logical volumes (LUNs) can *only* be accessed by one or more specified Fibre Channel host adapters (FCHAs). An FCHA is identified to the RAID400 subsystem via its *worldwide name* or *WWN*, a 16-digit HEX number that uniquely identifies the FCHA. Figure 6-6 shows an example configuration; in this figure, the FCHAs in the IBM NUMA-Q system and those in the NT system are connected to the RAID400 subsystem through the same two switches. Although they share fiber links to the RAID400 subsystem, FCHAs in the IBM NUMA-Q system cannot see LDEVs assigned to the WWNs of the NT system, nor can FCHAs in the NT system see LDEVs assigned to the WWNs of the FCHAs in the IBM NUMA-Q system.

Details on setting up LUN security are covered in two Hitachi documents:

Hitachi Freedom Storage™ Lightning 9900™ Remote Console User's Guide
Hitachi Freedom Storage™ Lightning 9900™ LUN Manager User's Guide

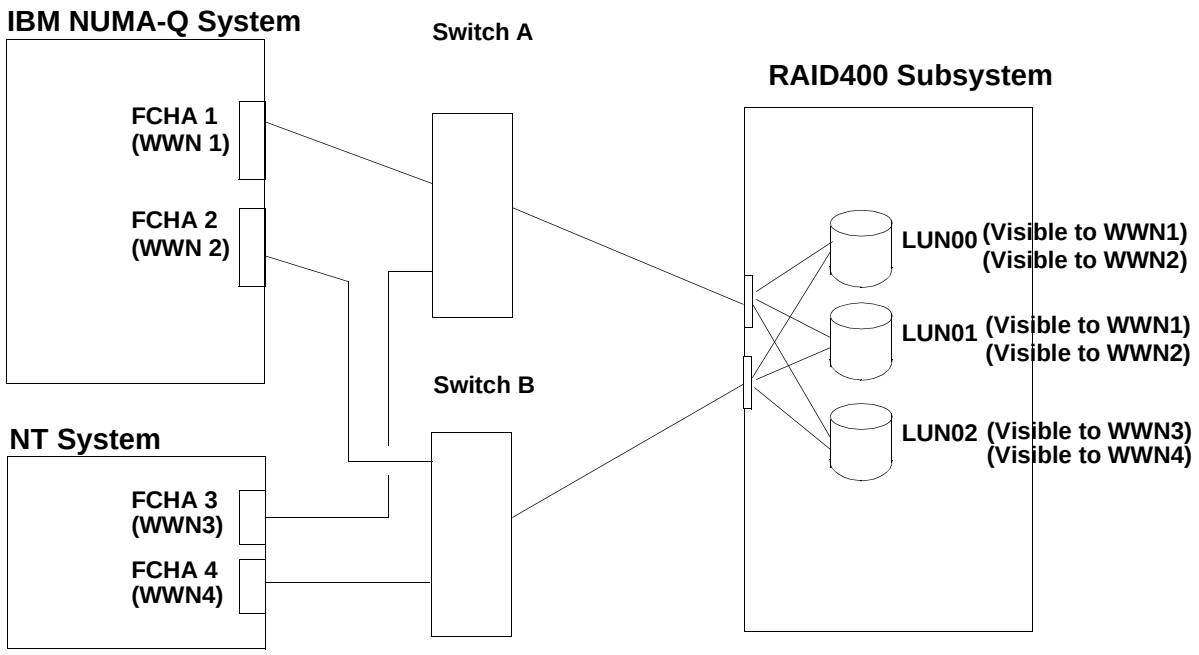

Figure 6-6.
Example Connection Scheme with LUN Security Feature.

Getting Ready to Use LUN Security

To use LUN security, you must do the following:

1. Install LUN Security per the instructions from Hitachi.

2. On the host system, identify the *worldwide names* (WWNs) of the Fibre Channel host adapters (FCHAs) that will be part of the LUN security scheme. To identify the FCHA worldwide names, you can do either of the following:

   - Use **dumpconf** or **showcfg** to identify the FCHAs in this host system (FCHAs currently used in IBM NUMA-Q systems will be identified as "FF" devices). Then enter **infodev -a <*FCHA_name*>** for each FCHA to be included in the LUN security scheme; e.g., for FCHA *ff1*, you will see output similar to the following:

     ```
     /etc/infodev -a ff1
     Vendor ID                      : Emulex
     Product ID                     : Firefly
     Revision Level                 : FireFly-3
     Device Class                   : ctlr
     Node WWN                       : 10:00:00:00:c9:20:58:4f
     Port WWN                       : 10:00:00:00:c9:20:58:4f
     Located on                     : quad0
     ```

     WWNs for the adapters used in IBM NUMA-Q systems begin with the digits "10"; WWNs beginning with "50" indicate RAID400 subsystem adapters.

     ### CAUTIONS

     *With some host adapter firmware versions, the Node WWN may differ in appearance from the Port WWN; **be sure** to use the **Port WWN** to set up LUN security.*

     *In addition, in a LUN security configuration, the WWNs of the host system FCHAs that are registered in the array's controller ports must match **exactly** the WWNs of the probing FCHAs on the host system. If there is any mismatch, DYNIX/ptx will fail to recognize **all** the logical units in that LUN security configuration.*

   - Alternatively, you can **telnet** into the appropriate switch, and use the **switchShow** command, whose output includes port names *and* their worldwide names.

     ```
     admin> switchShow
     switchName:     switch
     switchType:     3.2
     switchState:    Online
     switchRole:     Principal
     switchDomain:   1
     switchId:       fffc41
     switchWwn:      10:00:00:60:69:20:15:a0
     port  0: sw  No_Light
     port  1: sw  Online     F-Port  10:00:00:00:c9:20:d2:26
     port  2: sw  Online     F-Port  50:00:0e:10:00:00:54:52
     port  3: sw  Online     F-Port  50:00:0e:10:00:00:47:67
     port  4: sw  No_Light
     port  5: sw  No_Light
     port  6: sw  Online     F-Port  21:00:00:e0:8b:01:0d:bf
     port  7: sw  No_Light
     switch:admin>
     ```

3. Write down the WWN for each port that will be used for LUN security. You will need to enter these WWNs on the appropriate RAID400 configuration screen.

Chapter 7

DF500 Subsystem Integration

This chapter describes the steps needed to integrate a Hitachi DF500 RAID disk array subsystem with a DYNIX/ptx host system running ptx/HITACHI-RAIDMP. The following topics are included:

- **Requirements for use.** Gives hardware requirements for the use of this subsystem with ptx/HITACHI-RAIDMP.
- **Related Documentation.** Lists the Hitachi documentation that you will find most useful in completing the integration tasks described in this chapter.
- **Restrictions.** Includes restrictions specific to the functioning and use of this subsystem with ptx/HITACHI-RAIDMP.
- **DF500 RAID concepts and architecture.** Gives a brief summary of the DF500 RAID concepts and subsystem architecture.
- **DF500 subsystem integration overview.** Gives a high-level overview of the steps required to integrate this subsystem with the DYNIX/ptx host system.
- **Integration procedures.** Includes detailed procedures for ensuring that the subsystem has the required configuration settings and for installing and connecting the subsystem to the DYNIX/ptx host system.
- **Notes on Use and Administration.** Includes information that may be useful in the ongoing administration and use of the subsystem.

DF500 Requirements for Use

The following are required to integrate a DF500 subsystem with a DYNIX/ptx host system and this release of ptx/HITACHI-RAID:

- The DYNIX/ptx system must meet the requirements listed in the software chapter of this manual.
- Fibre Channel switched (FC-SW) connections must be via one of the switches supported by IBM, running a supported microcode version. Refer to the discussion in Chapter 1 titled *FC Switch Support and Operation*.
- The Hitachi DF500 RAID subsystem must meet the following requirements:
  - Configured for Fibre Channel switched fabric (FC-SW) or direct connect arbitrated loop (FC-AL) connectivity.
  - Fibre Channel interface cards installed.
  - Minimum microcode level 0555/C; contact the IBM NUMA-Q Customer Support Center for information on the latest supported microcode release(s).
- Fastpatch FP# 254922 is required for connection or operation of DF500 subsystems.

Related Documentation

The following Hitachi documentation will be especially useful in completing the integration tasks described in this chapter. You should have these manuals handy at all times.

- *Hitachi Freedom Storage™ Thunder 9200™ User and Reference Guide*
- *Hitachi DF-F500-WSEC LUN Security User's Guide*
- *Disk Array Management Program (for GUI) User's Guide*

DF500 Restrictions

The following restrictions apply to the integration and use of the DF500 subsystem with a DYNIX/ptx host system and this release of ptx/HITACHI-RAIDMP. These are in addition to the general and subsystem-specific restrictions listed in Chapter 2 of this document.

- Only switched fabric (FC-SW) and direct connect arbitrated loop (FC-AL) Fibre Channel connections are supported. SCSI connections and FL port arbitrated loop connections *are not supported*.
- There is no clusters support for direct connect FC-AL connections.
- DF500 disks cannot be used as boot or dump devices.
- The subsystem's RS-232 port is not an ASCII terminal port. It is for use only with the Hitachi Disk Array Management Program (DAMP), which runs on a PC with MicroSoft Windows™.
- Only the "single controller" or "dual controller/Dual-Active mode" configurations are supported.
- Hot standby mode is not supported.
- Only subsystems mounted in cabinets supplied by Hitachi or Hitachi distributors are supported in this release.
- The DF500 Turbo LU feature is supported.
- When using the subsystem with ptx/Clusters, "Data Share" mode *must be* ON.

DF500 RAID Concepts and Architecture

Basic DF500 Architecture

The following describe basic RAID concepts and basic system architecture of the DF500 subsystem. The terms and concepts discussed here should enable the reader familiar with RAID concepts and RAID architectures in general to complete the integration tasks described in the remainder of this section. For additional details on the DF500 subsystem architecture or RAID concepts, refer to the documentation shipped with the subsystem.

A single DF500 controller connects up to 100 hard disk drive units, in configurations similar to that shown in Figure 7-1. Each controller connects to four (4) Fibre Channel loops. Usually, the controller accesses a disk by two loops, the “normal” loops. The other two “alternate” loops can be used if a fault occurs in the normal loops. The normal loop for controller 0 is an alternate loop for controller 1, while the normal loop for controller 1 is an alternate loop for controller 0.

In a dual controller configuration, one controller accesses the disk connected by the alternate loop of the other controller. The connected disk unit is connected to two loops by a dual port route. One loop of the array unit connects up to 50 data or parity disks.

Cabinets are classified into *basic* cabinet and *expansion* cabinets. The basic cabinet contains two controllers and up to 10 disk units. Each expansion cabinet contains up to 10 disk units. The maximum configuration consists of a basic cabinet and nine expansion cabinets, for up to 100 disk modules in a single subsystem.

The spare disk setting position is variable. Except for the first two disk units of the basic cabinet, optional positions of the connected disks may be used as spares. You can set up to five (5) spare disks per 100-disk array.

A maximum of two (2) ports per controller can be connected to the host system.

A PC or workstation running the DAMP (Disk Array Management Program) can connect to the array unit via RS232C interface or LAN interface. This utility program allows you to define arrays and to check or change subsystem configuration. No other tools are provided to perform these tasks.

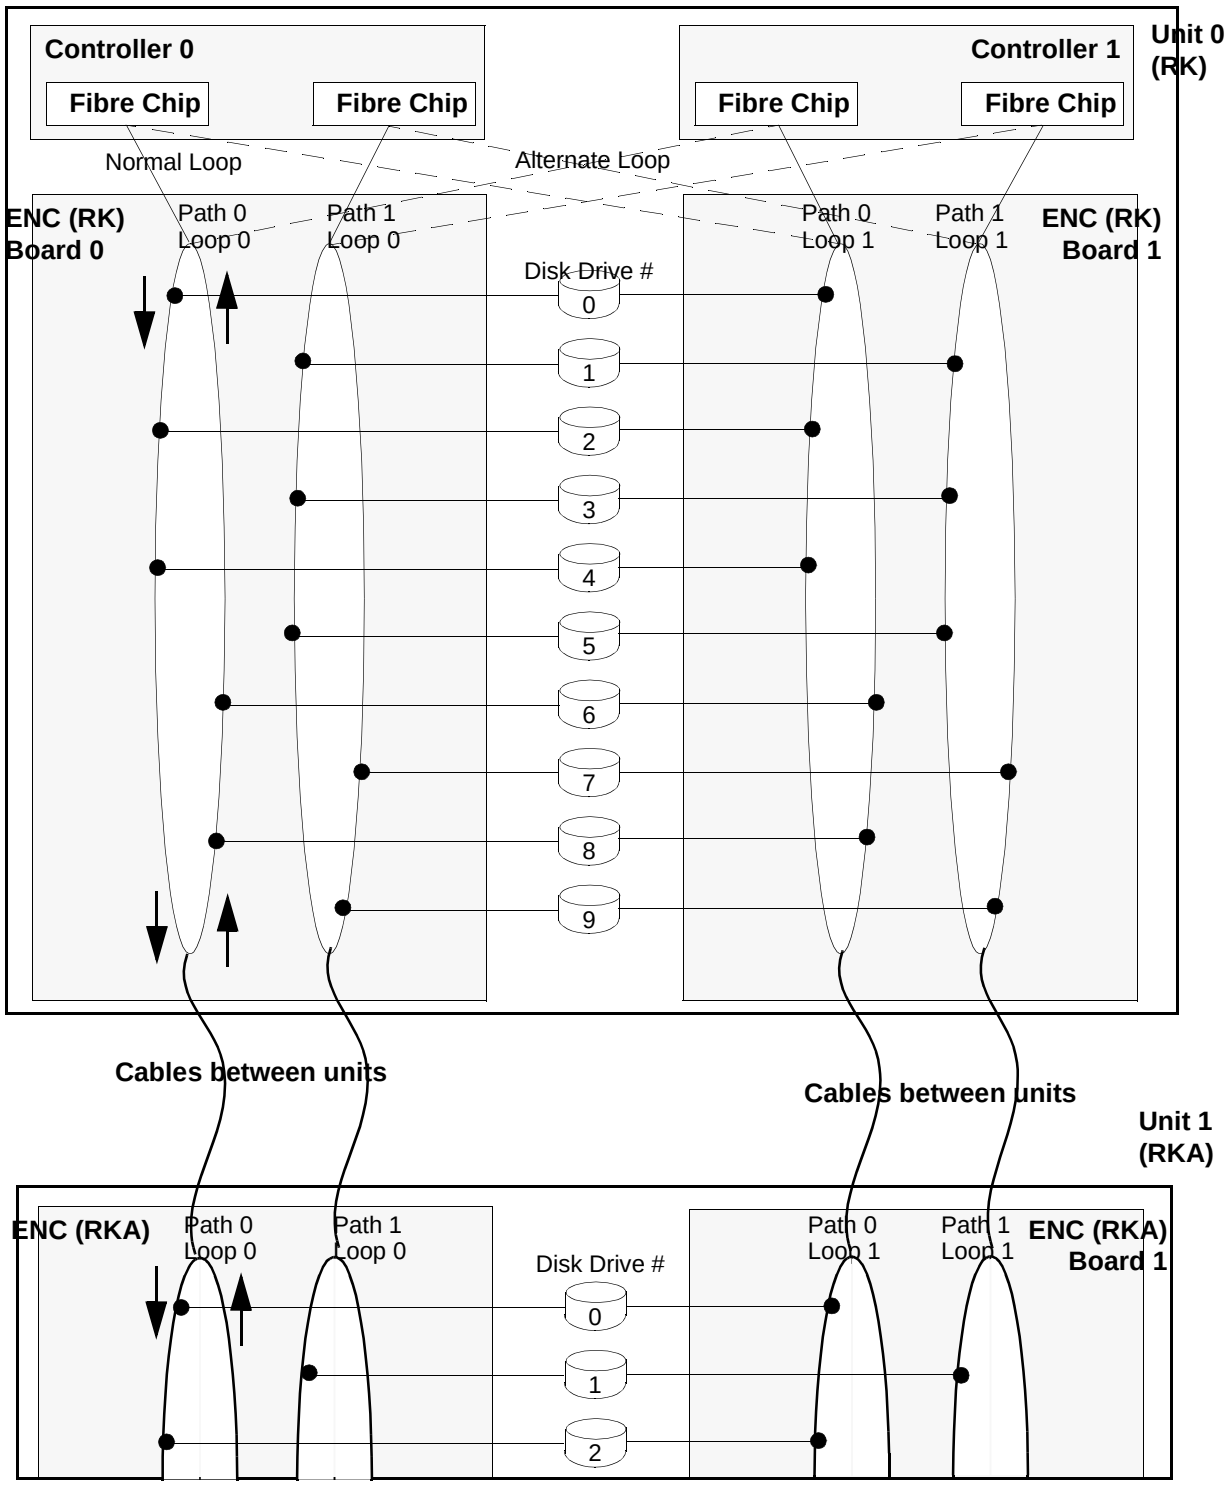

Figure 7-1.
DF500 System Architecture.

Multiporting on the DF500 Subsystem

Figure 7-2 shows a typical DF500 multiport configuration. In this configuration, each LU is "owned" by a single controller. A common configuration is for three LUs to be owned by each controller, with only one I/O route "active" for each LU at a given time. Ownership of each LU is defined at the time the LUs are configured, during hardware setup. After ownership of an LU has changed, the subsystem must be rebooted.

The DYNIX/ptx multiport (**mpt**) driver is able to identify "Primary" and "Secondary" routes from subsystem Inquiry data. It is critical that the subsystem have the correct configuration settings. If these settings are not correct, either the devices will not be identified and configured correctly by DYNIX/ptx, or performance may be degraded.

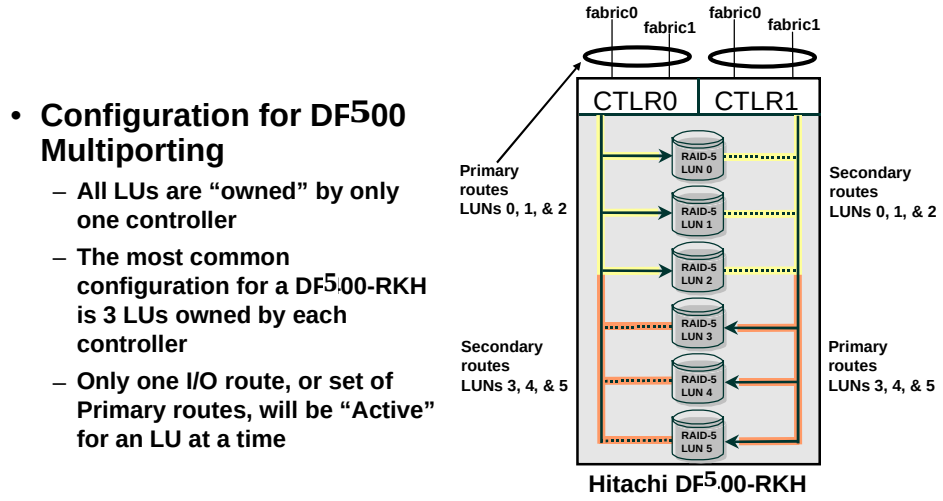

Figure 7-2.
Multiport Configuration.

In the event of a route failure (e.g., controller failure, bad cable, bad switch, etc.), the driver will automatically reroute I/O to the "Secondary" route. Because the Secondary route's controller is not the "owning" controller, performance will be degraded but will allow I/O to continue until the Primary route is restored and/or the cause of the failure is resolved. Once the failure is resolved, the driver will automatically switch back to the Primary route.

The following two figures show examples of multiport operation.

- In Figure 7-3, the Primary route on fabric 0 to Controller 0 has failed. Note that I/O continues to Controller 0 through the *Primary* route on fabric 1.
- In Figure 7-4, *both Primary routes* to Controller 0 have failed. Note that I/O *still continues to Controller 0*, but through the Secondary routes.

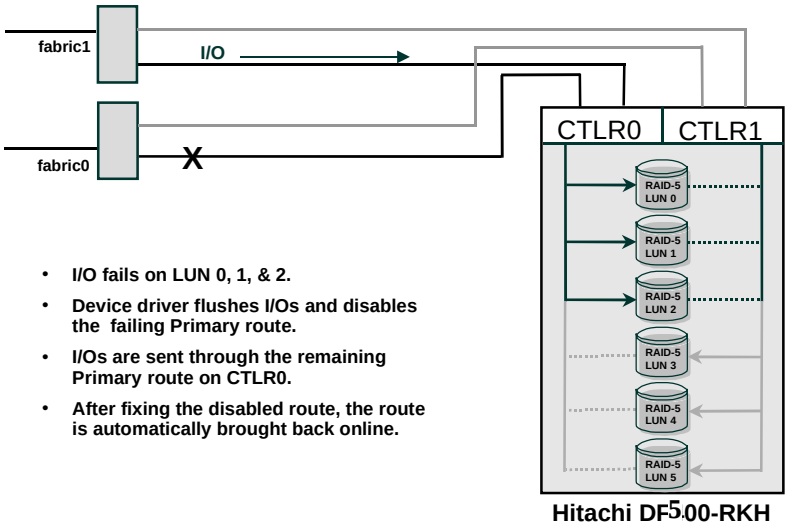

Figure 7-3.
Single Primary Path Failure.

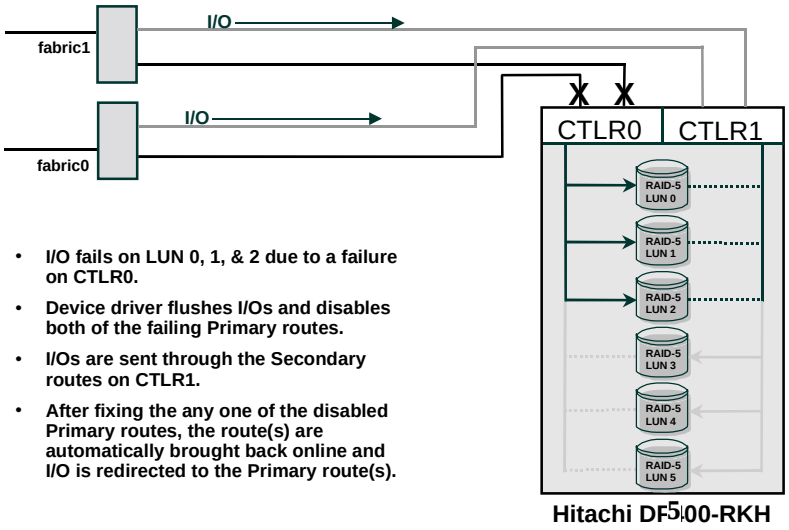

Figure 7-4.
Dual Primary Path Failure.

DF500 Integration Overview

To integrate the DF500 subsystem with a DYNIX/ptx host system, you must complete the following major steps:

- Step A: Install the ptx/HITACHI-RAIDMP software.
- Step B: Prepare the RAID subsystem for integration.
- Step C: Check or adjust the subsystem's configuration settings.
- Step D: Make cable connections between the host system and the subsystem.
- Step E: Verify the integration.

The following discussions cover these steps in detail or refer to other documentation for details.

Step A: Install ptx/HITACHI-RAIDMP

At this time, if you have not already done so, refer to Chapter 2 of this manual and install the ptx/HITACHI-RAIDMP software.

Step B: Prepare for Integration

Hitachi or Hitachi distributor personnel should unpack, install, and prepare the subsystem for integration

After the subsystem has been physically installed and located where it will be used, you should do the following:

1. Collect all cables and terminators, etc. that will be required.

2. Collect all required documentation.

Step C: Check and Adjust Configuration Settings

CAUTION

Use extreme caution when changing the subsystem's RAID configuration. Do not attempt to do this unless you are thoroughly familiar with RAID concepts as implemented in this subsystem. Errors in configuration can cause data corruption or loss. For example, removing a logical unit causes all data on that logical unit to be lost!

The subsystem's configuration settings are viewed and adjusted using the Disk Array Management Program (DAMP), which is provided with the subsystem. The following lists the critical configuration settings for the DF500 subsystem to operate with this release of ptx/HITACHI-RAIDMP.

Required Settings

The following settings are required for the DF500 subsystem to work properly with a DYNIX/ptx host system running this release of ptx/HITACHI-RAIDMP.

Controller 0 and 1, Common Parameters

- Dual active
- Data share ON (required for use with ptx/Clusters)
- Multipath controller
- Command queueing — ON
- Cache mode should be set to OFF
- Host connect mode (for each port)
- Trespass mode

Controller 0, Controller 1 Separately

- Controller Identification ON
- Connect LAN — YES
- DHCP — OFF
  (Set IP address, mask, and default gateway as required)

Fibre Channel

- Topology — set as required

NOTE

Cache mode should be set to all OFF; this should be the default setting. If you need to change this parameter, use the Disk Array Management Program (DAMP).

Step D: Make Cable Connections

Refer to the correct cabling diagram and connect the disk array subsystem(s) to the host system. The first figure shows a schematic diagram of the DF500 back panel connectors, including the Fibre connectors for both controllers. This diagram is provided only for reference locations of connectors; refer to the subsystem documentation for details on back panel connectors and switches.

The two following figures show typical switched fabric (FC-SW) and direct connect arbitrated loop (FC-AL) Fibre Channel cabling schemes. ***Note***, however, that these are example diagrams only; actual cabling configurations will depend on the particular site configuration and needs.

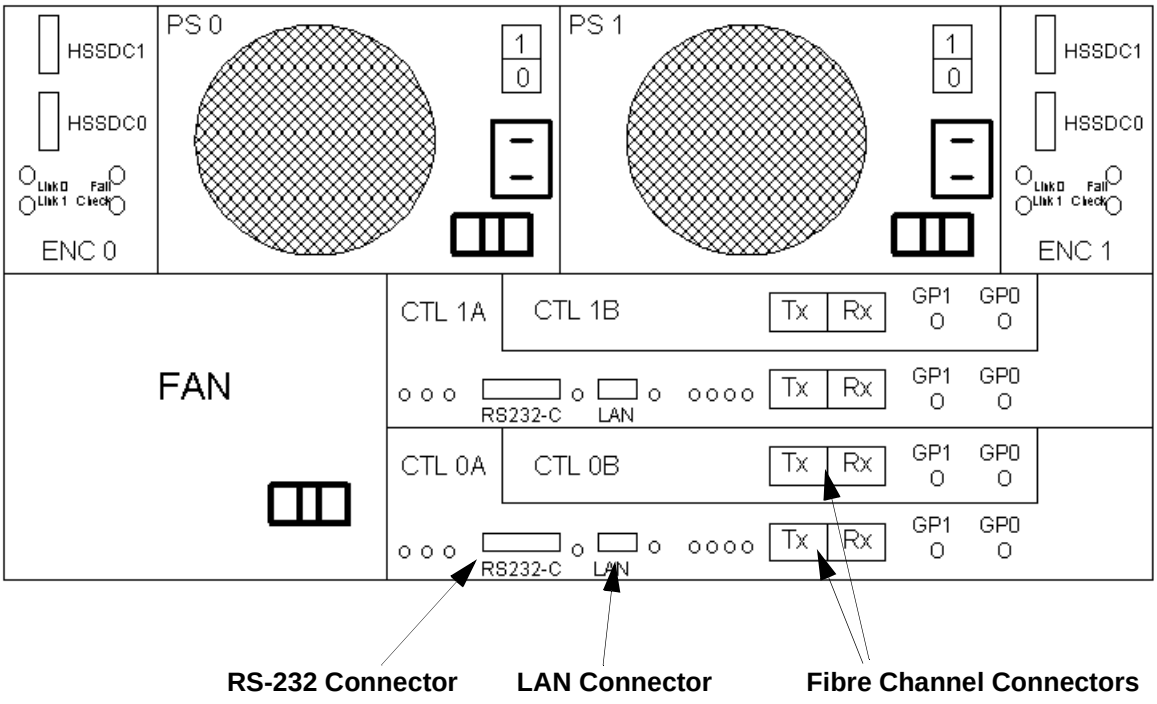

Figure 7-5.
DF500 Back Panel Connections.

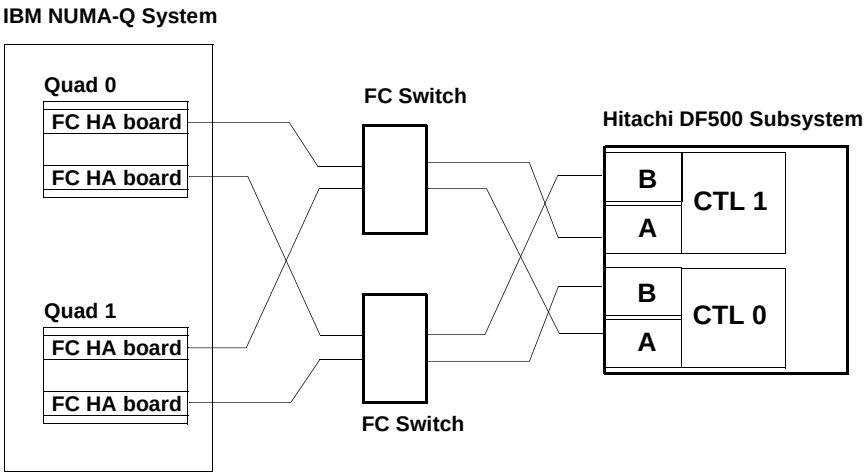

Figure 7-6.
Typical Switched Fabric (FC-SW) Fibre Channel Connection.

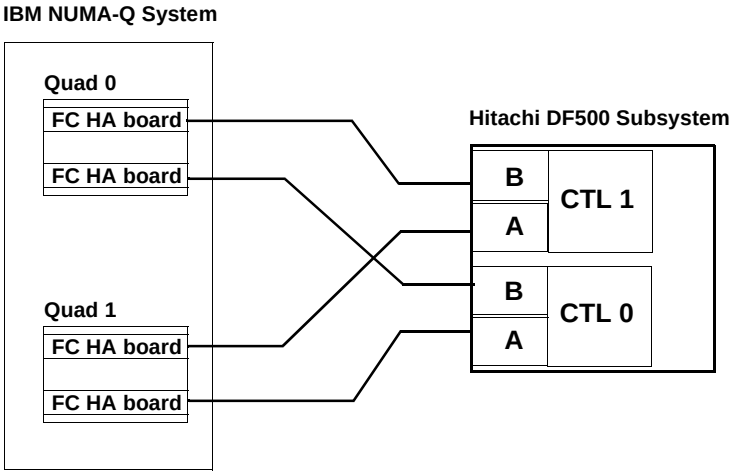

Figure 7-7.
Typical Direct Connect Arbitrated Loop (FC-AL) Fibre Channel Connection.

Step E: Verify the Integration

To verify the installation and integration of the subsystem:

1. Bring the subsystem up.
2. Warn users of the host system that the system will be coming down.
3. Reboot the host system with the new kernel that you compiled when you installed the ptx/HITACHI-RAIDMP software.
4. Use the **dumpconf**, **diskid** and **infodev** commands to verify that all **sd** devices were found at system boot, to determine which of these are DF500 devices, and to determine fabric and path information for the DF500 devices. Figure 7-5 shows sample output for these commands for a switched fabric (FC-SW) connection; Figure 7-6 shows sample output for these commands for a direct connect arbitrated loop (FC-AL) connection.
5. Once the devices have been verified, perform I/O to each device to verify its operation.
6. Complete the preparation of the RAID subsystem for use.
    - If you are using VTOCs, refer to the discussion of the **make_vtoc** utility, at the end of this chapter.
    - If you are using ptx/SVM to manage disk resources, refer to the ptx/SVM documentation supplied with the DYNIX/ptx host system.
    - If you are using switched fabric (FC-SW) connection, check that all ports have logged in successfully. To do this, log into the switch and use the **switchShow** command to verify the connections.
    - If you are using the LUN security feature of the DF500 subsystem, refer to the DF500 subsystem's LUN security user's guide for details.
7. Bring the DYNIX/ptx system back to multi-user mode and notify users that it is now available for use.

```
# dumpconf
NAME         CFGTYPE  DEVNUM  UNIT        FLAGS  OnBUS    OnDEVICE
quad0        quad          0  0x00000000  L      sci      sci0
asy0         asy           0  0x00000000  L      eisa     quad0
asy1         asy           1  0x00000001  L      eisa     quad0
mdc0         mdc           0  0x00000000  L      eisa     quad0
ff0          ff            0  0x00000002  SP     pci      quad0
fabric7      fabric        7  0x00fea800  SM     fc       ff0
sd10         sd           10  0x02fe8000  SM     fcp      fabric7
sd11         sd           11  0x02fe8001  SM     fcp      fabric7
sd12         sd           12  0x02fe8002  SM     fcp      fabric7
sd13         sd           13  0x02fe8003  SM     fcp      fabric7
sd14         sd           14  0x02fe8004  SM     fcp      fabric7
sd15         sd           15  0x02fe8005  SM     fcp      fabric7
sd10         sd           10  0x01fe9000  SM     fcp      fabric7
sd11         sd           11  0x01fe9001  SM     fcp      fabric7
sd12         sd           12  0x01fe9002  SM     fcp      fabric7
sd13         sd           13  0x01fe9003  SM     fcp      fabric7
sd14         sd           14  0x01fe9004  SM     fcp      fabric7
sd15         sd           15  0x01fe9005  SM     fcp      fabric7
sd10         sd           10  0x12fea000  SM     fcp      fabric7
sd11         sd           11  0x12fea001  SM     fcp      fabric7
sd12         sd           12  0x12fea002  SM     fcp      fabric7
sd13         sd           13  0x12fea003  SM     fcp      fabric7
sd14         sd           14  0x12fea004  SM     fcp      fabric7
sd15         sd           15  0x12fea005  SM     fcp      fabric7
sd10         sd           10  0x11feb000  SM     fcp      fabric7
sd11         sd           11  0x11feb001  SM     fcp      fabric7
sd12         sd           12  0x11feb002  SM     fcp      fabric7
sd13         sd           13  0x11feb003  SM     fcp      fabric7
sd14         sd           14  0x11feb004  SM     fcp      fabric7
sd15         sd           15  0x11feb005  SM     fcp      fabric7
ff1          ff            1  0x00000006  SP     pci      quad0
```

```
# diskid (partial output)
sd10: fabric7 cport 1 fc id 0x210b13, fcp lun 00:00:00:00:00:00:00:00
  No Pbay on fabric7
  Vendor HITACHI, product DF500F, revision 10400-00, serial # 10400-00
  capacity 136112 Mbytes
```

```
# infodev -n (partial output)
Fibre Channel Address           : 0x210613 on fabric7 (unit 0xfeb000)
Fibre Channel Address           : 0x210713 on fabric7 (unit 0xfeb800)
Fibre Channel Address           : 0x210913 on fabric7 (unit 0xfec800)
Fibre Channel Address           : 0x210b13 on fabric7 (unit 0xfed800)
```

Figure 7-8.
Sample dumpconf, diskid, and infodev Output for FC-SW Connection.

```
# dumpconf

NAME         CFGTYPE  DEVNUM  UNIT        FLAGS  OnBUS    OnDEVICE
quad0        quad          0  0x00000000  L      sci      sci0
asy0         asy           0  0x00000000  L      eisa     quad0
asy1         asy           1  0x00000001  L      eisa     quad0
mdc0         mdc           0  0x00000000  L      eisa     quad0
ff0          ff            0  0x00000002  SP     pci      quad0
fabric0      fabric        0  0x0001f800  SM     fc       ff0
sd10         sd           10  0x01fbf800  SM     fcp      fabric0
sd11         sd           11  0x01fbf801  SM     fcp      fabric0
sd12         sd           12  0x01fbf802  SM     fcp      fabric0
sd13         sd           13  0x01fbf803  SM     fcp      fabric0
sd14         sd           14  0x01fbf804  SM     fcp      fabric0
sd15         sd           15  0x01fbf805  SM     fcp      fabric0
ff1          ff            1  0x00000006  SP     pci      quad0
fabric1      fabric        1  0x0001f800  SM     fc       ff1
sd10         sd           10  0x02fbf800  SM     fcp      fabric1
sd11         sd           11  0x02fbf801  SM     fcp      fabric1
sd12         sd           12  0x02fbf802  SM     fcp      fabric1
sd13         sd           13  0x02fbf803  SM     fcp      fabric1
sd14         sd           14  0x02fbf804  SM     fcp      fabric1
sd15         sd           15  0x02fbf805  SM     fcp      fabric1
```

```
# diskid (partial output)

sd12: fabric0 cport 1 fc id 0xef, fcp lun 00:02:00:00:00:00:00:00
  No Pbay on fabric0
  Vendor HITACHI, product DF500F, revision --, serial # 10400-02
  capacity 136112 Mbytes
```

```
# infodev -n (partial output)

Fibre Channel Address           : 0xef on fabric0 (unit 0xfbf800)
Fibre Channel Address           : 0xef on fabric1 (unit 0xfbf800)
Fibre Channel Address           : 0xef on fabric2 (unit 0xfbf800)
Fibre Channel Address           : 0xef on fabric3 (unit 0xfbf800)
Fibre Channel Address           : 0xef on fabric0 (unit 0xfbf801)
Fibre Channel Address           : 0xef on fabric1 (unit 0xfbf801)
Fibre Channel Address           : 0xef on fabric2 (unit 0xfbf801)
Fibre Channel Address           : 0xef on fabric3 (unit 0xfbf801)
```

Figure 7-9.
Sample dumpconf, diskid, and infodev Output for FC-AL Connection.

The make_vtoc(1M) Utility

The **make_vtoc**(1M) utility allows you to define custom VTOCs. Refer to the on-line man page or to the hard copy of the man page in Appendix B of this manual for details. Refer to the DYNIX/ptx **mkvtoc** man page for additional details on VTOC creation.

LUN Security

LUN Security Overview

The DF500 LUN Security feature enables a DF500 subsystem to be configured so that particular logical volumes (LUNs) can *only* be accessed by one or more specified Fibre Channel host adapters (FCHAs). An FCHA is identified to the subsystem via its *worldwide name* or *WWN*, a 16-digit HEX number that uniquely identifies the FCHA. Figure 6-6 shows an example configuration; in this figure, the FCHAs in the IBM NUMA-Q system and those in the NT system are connected to the subsystem through the same two switches. Although they share fiber links to the subsystem, FCHAs in the IBM NUMA-Q system cannot see LDEVs assigned to the WWNs of the NT system, nor can FCHAs in the NT system see LDEVs assigned to the WWNs of the FCHAs in the IBM NUMA-Q system.

For details on setting up LUN security, see *Related Documentation*, earlier in this chapter.

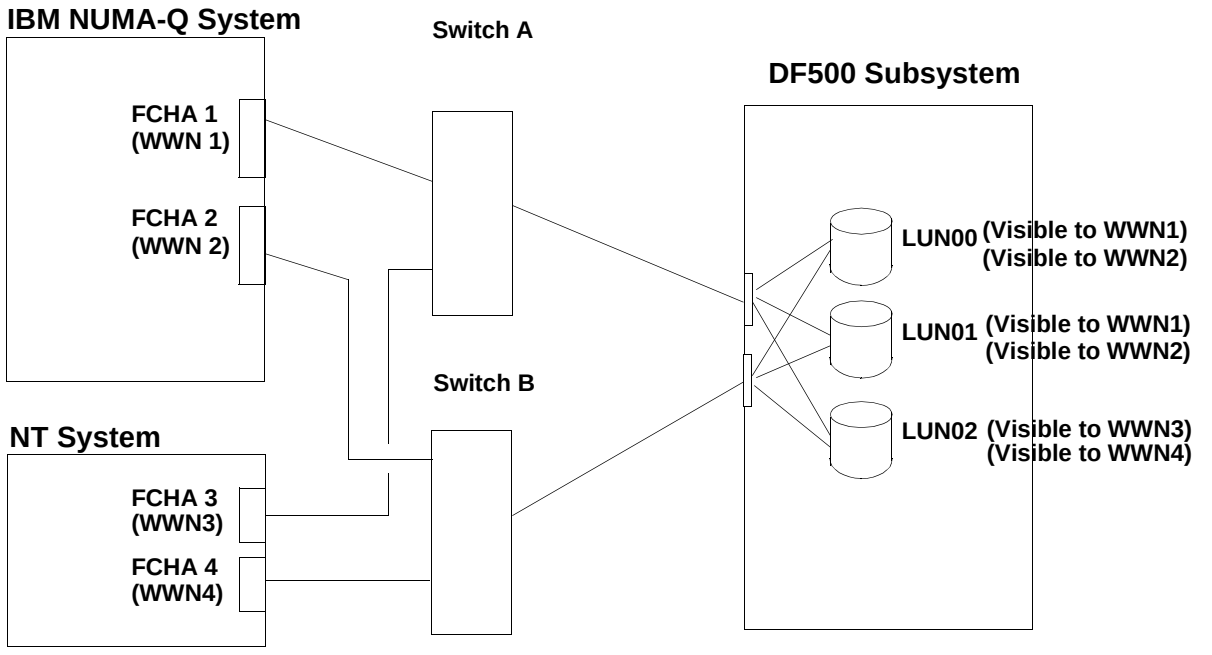

Figure 6-10.
Example Connection Scheme with LUN Security Feature.

Getting Ready to Use LUN Security

To use LUN security, you must do the following:

1. Install LUN Security per the instructions from Hitachi.

2. On the host system, identify the *worldwide names* (WWNs) of the Fibre Channel host adapters (FCHAs) that will be part of the LUN security scheme. To identify the FCHA worldwide names, you can do either of the following:

   - Use **dumpconf** or **showcfg** to identify the FCHAs in this host system (FCHAs currently used in IBM NUMA-Q systems will be identified as "FF" devices). Then enter **infodev -a <*FCHA_name*>** for each FCHA to be included in the LUN security scheme; e.g., for FCHA *ff1*, you will see output similar to the following:

     ```
     /etc/infodev -a ff1
     Vendor ID                      : Emulex
     Product ID                     : Firefly
     Revision Level                 : FireFly-3
     Device Class                   : ctlr
     Node WWN                       : 10:00:00:00:c9:20:58:4f
     Port WWN                       : 10:00:00:00:c9:20:58:4f
     Located on                     : quad0
     ```

     WWNs for the adapters used in IBM NUMA-Q systems begin with the digits "10"; WWNs beginning with "50" indicate DF500 subsystem adapters.

     ### CAUTIONS

     *With some host adapter firmware versions, the Node WWN may differ in appearance from the Port WWN; **be sure** to use the **Port WWN** to set up LUN security.*

     *In addition, in a LUN security configuration, the WWNs of the host system FCHAs that are registered in the array's controller ports must match **exactly** the WWNs of the probing FCHAs on the host system. If there is any mismatch, DYNIX/ptx will not recognize **any** logical units in that LUN security configuration.*

   - Alternatively, you can **telnet** into the appropriate switch, and use the **switchShow** command, whose output includes port names *and* their worldwide names.

     ```
     admin> switchShow
     switchName:     switch
     switchType:     3.2
     switchState:    Online
     switchRole:     Principal
     switchDomain:   1
     switchId:       fffc41
     switchWwn:      10:00:00:60:69:20:15:a0
     port  0: sw  No_Light
     port  1: sw  Online     F-Port  10:00:00:00:c9:20:d2:26
     port  2: sw  Online     F-Port  50:00:0e:10:00:00:54:52
     port  3: sw  Online     F-Port  50:00:0e:10:00:00:47:67
     port  4: sw  No_Light
     port  5: sw  No_Light
     port  6: sw  Online     F-Port  21:00:00:e0:8b:01:0d:bf
     port  7: sw  No_Light
     switch:admin>
     ```

3. Write down the WWN for each port that will be used for LUN security. You will need to enter these WWNs on the appropriate DF500 configuration screen.

Appendix A

Man Pages

Following are hard copies of the online man pages for this release of ptx/HITACHI-RAIDMP.

NAME

dfutil - Hitachi DF350 class Disk Array Subsystem configuration display and regressed condition logging routine

SYNOPSIS

/usr/lib/raidmp/dfutil -a sd# or
/usr/lib/raidmp/dfutil -d sd# or
/usr/lib/raidmp/dfutil -l[ghl] [sd#]

DESCRIPTION

dfutil displays configuration and log (FRU status) settings of the Hitachi DF350 disk subsystem.

The following options are recognized.

-d Displays configuration status, including the disk array geometry and information about current and default controllers associated with the device sd#.

-a Displays all the information given by **-d** option plus the log (FRU status) data.

-l Displays the status of the FRUs (field-replaceable units) to the log file /usr/adm/raidmp/dfXXXlog. If all units are found to be in normal condition, a GOOD status is reported. Those FRUs that require preventive maintenance are not reported, as the controller does not provide the host with this information.

The following options are recognized with the **-l** option

-l Causes status to be reported to the system console, as well as to the log file. Logging of events to the console is independent of logging to the log file /usr/adm/raidmp/dfXXXlog, which is done by default.
If the device (**sd**#) is not specified, **dfutil** attempts to open all the configured sd units.

-g **dfutil** displays GOOD status if all FRUs are found to be in normal condition. This flag disables the message that is output when all FRU’s are in normal operation.

-h Enables the log message display to stdout.

sd# The sd device from which to read messages. Default is all sd devices found by **dumpconf**.

NOTE: dfutil tries to open the sd# as /dev/rdsk/sd#; if this fails, it tries to open the diag device /dev/diag/rdsk/sd#.

The disk array has redundant disks, power supplies, and fans, allowing it to block out any section in which an error occurred. If an error occurs, the array enters the regressed condition, lights a warning lamp on the unit, and displays an error on the display panel. In the regressed condition, the array can continue processing host commands. To prevent a system halt or lost user data caused by overlapped failures, faulty FRUs should be replaced immediately.

Message Format

The format of the records logged by **dfutil** is:

> host month host day host time **dfutil: user <**login**>, host <**hostname**>**

followed by zero or more one-line messages, in the following format:

```
sd#: FRU UNIT FRU STATUS
```

where:

FRU UNIT is one of the following -
- Disk Drive
- LUN Data
- Power Supply
- Battery
- Cache
- Fan

and where:

FRU STATUS indicates one of the following -
- Not Available
- Uninstalled
- Blocked <regressed condition>
- Standby <if the FRU is a drive>

The output from **dfutil -d** reports the following information:

- Product Name
- Vendor Name
- Revision Number
- Array geometry Information -
  - Number of Rows
  - Number of Disks per Row
  - Number of Parity Disks per Row
  - Number of Spare Disk Drives
  - Spare Disk Mode
  - Drive configuration <Rows/Ports>

```
          Controller Configuration <Single/Dual>
          F/W Revision Number
          System Version Number
          FRU Condition Information -
               Battery Status
               Power Supply Condition
               Fan Condition
               Cache Module Status
               Drive Status
               Data Status
          RAID Groups Information
          Current Controller
          Default Controller
```

FILES

/usr/lib/raidmp/dfutil /usr/adm/raidmp/dfXXXlog

DIAGNOSTICS

dfutil produces no diagnostics; it simply reports the current configuration with the -d option. With the -l option, it attempts to lock the log file /usr/adm/raidmp/dfXXXlog before logging the data. This will cause **dfutil** to fail if another instance of **dfutil** is running.

NAME

hrutil - Hitachi RAID400 class Disk Array Subsystem information display routine.

SYNOPSIS

/usr/lib/raidmp/hrutil -i[a] sd#

DESCRIPTION

hrutil displays configuration and inquiry page information of the Hitachi RAID400 disk subsystem.

The following options are recognized.

-i Displays important fields of inquiry page information, including the Host Mode of sd#.

-ia Displays complete inquiry page information.

sd#
The sd device from which to read messages. Default is all sd devices found by **dumpconf**.

NOTE: hrutil tries to open the sd# as /dev/rdsk/sd#; if this fails, it tries to open the diag device /dev/diag/rdsk/sd#.

The inquiry page of the device has information like Vendor ID, Product ID, Product Revision level, Product serial number, location of the port, volume status and LDEV information and more. This utility tries to display all that and more in a readable format.

FILES

/usr/lib/raidmp/hrutil

DIAGNOSTICS

hrutil produces no diagnostics; it simply reports the inquiry data plus some more information.

NAME

make_vtoc - ptx/HITACHI-RAIDMP make a VTOC for a specified disk

SYNOPSIS

/usr/lib/raidmp/make_vtoc disk name altname description

DESCRIPTION

The **make_vtoc** will create the files necessary for a custom VTOC for the specified disk. **make_vtoc** will error out if /etc/diskinfo/name.geom or /etc/vtoc/name file exists. **make_vtoc** will ask the user to input the type and size of each slice.

Note: Some disk geometry parameters are fixed. For example, Slice 0 will always start at 256. Also, the last two slices are hardcoded: the next to last slice will be of type 2, start at 0, and of size 16; and the last slice will be of type 3, start at 16, and of size 16.

The **make_vtoc** utility allows the user to configure all other slices. These slices would typically be of types 1 or 8.

Typically, you should set the altname parameter to the same string as name, but in all uppercase. The following is an example invocation of **make_vtoc**:

```
/usr/lib/raidmp/make_vtoc sd15 raidmp_5_16GB
RAIDMP_5_16GB "16GB HITACHI RAID MP 5 array"
```

The command will then prompt the user for the required partitioning information.

The disk capacity used in the geom and the vtoc file will be retrieved from infodev.

Note: The **make_vtoc** utility creates a VTOC file. The user must run the **mkvtoc** command to populate the VTOC on the desired disks. See the **vtoc** and **mkvtoc** man pages for details.

FILES

/usr/lib/raidmp/make_vtoc

SEE ALSO

vtoc(1G), **mkvtoc**(1M)

DIAGNOSTICS

make_vtoc will exit if improperly invoked or if any of the files exist.