HITACHI
DATA SYSTEMS

HiCommand® Dynamic Link Manager (HDLM) for Solaris™ Systems

User's Guide

MK-92DLM114-13

© 2006, 2007 Hitachi, Ltd., Hitachi Data Systems Corporation, ALL RIGHTS RESERVED

Notice: No part of this publication may be reproduced or transmitted in any form or by any electronic or mechanical means, including photocopying and recording, or stored in a database or retrieval system for any purpose, without the express written permission of Hitachi Data Systems Corporation (hereinafter referred to as "Hitachi Data Systems").

Hitachi Data Systems reserves the right to make changes to this document at any time without notice and assumes no responsibility for its use. Hitachi Data Systems products and services can only be ordered under the terms and conditions of Hitachi Data Systems' applicable agreements, including license agreements. All of the features described in this document may not be currently available. Refer to the most recent product announcement or contact your local Hitachi Data Systems sales office for information on feature and product availability.

This document contains the most current information available at the time of publication. When new and/or revised information becomes available, this entire document will be updated and distributed to all registered users.

Trademarks

Hitachi Data Systems is a registered trademark and service mark of Hitachi, Ltd., and the Hitachi Data Systems design mark is a trademark and service mark of Hitachi, Ltd.

HiCommand, Hitachi TagmaStore, Lightning 9900, Thunder 9500, and Thunder 9200 are trademarks of Hitachi Data Systems Corporation in the United States and other countries.

AIX is a registered trademark of the International Business Machines Corp. in the U.S.

HP-UX is a product name of Hewlett-Packard Company.

HP StorageWorks is a trademark of Hewlett-Packard Company.

Linux is a registered trademark of Linus Torvalds.

Microsoft is a registered trademark of Microsoft Corp. in the U.S. and other countries.

Mozilla is a trademark of the Mozilla Foundation in the U.S and other countries.

OpenBoot PROM is a trademark or registered trademark of Sun Microsystems, Inc. in the United States and other countries.

ORACLE is a registered trademark of Oracle Corporation.

Oracle9i is a trademark of ORACLE Corporation.

Oracle 10g is a trademark of ORACLE Corporation.

Solaris is a trademark or registered trademark of Sun Microsystems, Inc. in the United States and other countries.

Solstice is a trademark or registered trademark of Sun Microsystems, Inc. in the United States and other countries.

All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc., in the United States and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

Sun is a trademark or registered trademark of Sun Microsystems, Inc. in the United States and other countries.

Sun StorageTek is a trademark or registered trademark of Sun Microsystems, Inc. in the United States and other countries.

SureStore is a registered trademark of Hewlett-Packard Company.

UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company Limited.

VERITAS Cluster Server is a trademark of Symantec Corporation.

VERITAS Volume Manager is a trademark of Symantec Corporation.

VERITAS, VERITAS logo are trademarks or registered trademarks of Symantec Software Corporation in the U.S. and other countries.

Windows is a registered trademark of Microsoft Corp. in the U.S. and other countries.

HiCommand Dynamic Link Manager includes some parts whose copyrights are reserved by Sun Microsystems, Inc.

HiCommand Dynamic Link Manager includes some parts whose copyrights are reserved by UNIX System Laboratories, Inc.

This product includes software developed by the JDOM Project (http://www.jdom.org/).

This product includes software developed by the Apache Software Foundation (http://www.apache.org/).

Notice of Export Controls

Export of technical data contained in this document may require an export license from the United States government and/or the government of Japan. Please contact the Hitachi Data Systems Legal Department for any export compliance questions.

Document Revision Level

Revision	Date	Description
MK-92DLM114-P	February 2002	Preliminary Release
MK-92DLM114-P1	March 2002	Revision 1 of Preliminary Release, supersedes and replaces MK-92DLM114-P
MK-92DLM114-00	April 2002	Initial release, supersedes and replaces MK-92DLM114-P1
MK-92DLM114-1P	October 2002	Revision 1P, supersedes and replaces MK-92DLM114-0
MK-92DLM114-02	November 2002	Revision 2, supersedes and replaces MK-92DLM114-1P
MK-92DLM114-03	January 2003	Revision 3, supersedes and replaces MK-92DLM114-2
MK-92DLM114-04	April 2003	Revision 4, supersedes and replaces MK-92DLM114-3
MK-92DLM114-05	August 2003	Revision 5, supersedes and replaces MK-92DLM114-4
MK-92DLM114-06	November 2003	Revision 6, supersedes and replaces MK-92DLM114-5
MK-92DLM114-07	March 2004	Revision 7, supersedes and replaces MK-92DLM114-6
MK-92DLM114-08	August 2004	Revision 8, supersedes and replaces MK-92DLM114-7
MK-92DLM114-09	October 2004	Revision 9, supersedes and replaces MK-92DLM114-8
MK-92DLM114-10	June 2005	Revision 10, supersedes and replaces MK-92DLM114-09
MK-92DLM114-11	April 2006	Revision 11, supersedes and replaces MK-92DLM114-10
MK-92DLM114-12	June 2006	Revision 12, supersedes and replaces MK-92DLM114-11
MK-92DLM114-13	June 2007	Revision 13, supersedes and replaces MK-92DLM114-12

Preface

This document describes and provides instructions for installing and using HiCommand® Dynamic Link Manager (HDLM) for Solaris software for Hitachi disk array subsystems. This document is intended for system administrators who use HDLM to operate and manage storage.

This document assumes that the user:

- Solaris and its management functionality
- Storage subsystem management functionality
- Cluster software functionality
- Volume management software functionality

For further information on the Hitachi storage subsystems, please refer to the user's guide for the subsystem (e.g., *Hitachi TagmaStore Universal Storage Platform User and Reference Guide*, MK-94RD231). For further information on Hitachi Data Systems products and services, please contact your Hitachi Data Systems account team, or visit Hitachi Data Systems online at http://www.hds.com.

Note: The use of HiCommand® Dynamic Link Manager (HDLM) and all other Hitachi Data Systems products is governed by the terms of your agreement(s) with Hitachi Data Systems.

Software Version

This document revision applies to HDLM for Solaris version 5.9

Convention for Storage Capacity Values

Storage capacity values displayed by HDLM are calculated based on the following values:

1 KB (kilobyte) = 1,024 bytes
1 MB (megabyte) = $1,024^2$ bytes
1 GB (gigabyte) = $1,024^3$ bytes
1 TB (terabyte) = $1,024^4$ bytes

Referenced Documents

- *Hitachi TagmaStore Universal Storage Platform User and Reference Guide*, MK-94RD231
- *Hitachi Lightning 9900 V Series User and Reference Guide*, MK-92RD100
- *Hitachi Lightning 9900 User and Reference Guide*, MK-90RD008
- *Hitachi TagmaStore Adaptable Modular Storage 1000 User and Reference Guide,* MP-95DF780
- *Hitachi TagmaStore Adaptable Modular Storage 500 User and Reference Guide,* MP-95DF714

- *Hitachi Thunder 9500 V Series User and Reference Guide*, MK-92DF601
- *Hitachi Thunder 9200 User and Reference Guide*, MK-90DF504
- *Hitachi Freedom Storage 5000 Series Software Configuration User's Guide*, BO-98DF376

Comments

Please send us your comments on this document. Make sure to include the document title, number, and revision. Please refer to specific section(s) and paragraph(s) whenever possible.

a) **E-mail:** doc.comments@hds.com

b) **Fax:** 858-695-1186

c) **Mail:**
Technical Writing, M/S 35-10
Hitachi Data Systems
10277 Scripps Ranch Blvd.
San Diego, CA 92131

Thank you! (All comments become the property of Hitachi Data Systems Corporation.

Contents

Chapter 1 Overview of HDLM 1

1.1 What is HDLM? 2
1.2 HDLM Features 3

Chapter 2 HDLM Functions 5

2.1 Devices Managed by HDLM 5
2.2 System Configuration 6
2.3 LU Configuration 8
2.4 Program Configuration 9
2.5 Position of the HDLM Driver and HDLM Device 11
2.6 Logical Device Files for HDLM Devices 12
2.7 Load Distribution Using Load Balancing 13
2.7.1 appliedappliedPaths to Which Load Balancing is Applied 15
2.7.2 Algorithms for Load Balancing 17
2.8 Failover and Failback Using Path Switching 18
2.8.1 Automatic Path Switching 19
2.8.2 Manual Path Switching 21
2.8.3 Path Status Transition 22
2.9 Monitoring Intermittent Errors (Functionality When Automatic Failback Is Used) ... 24
2.9.1 Checking Intermittent Errors 24
2.9.2 Setting Up Intermittent Error Monitoring 25
2.9.3 Actions for Intermittent Error Monitoring 25
2.9.4 When User Operations Change the Intermittent Error Information 28
2.10 Detecting Errors by Using Path Health Checking 29
2.11 Error Management 29
2.11.1 Types of Logs Collected 30
2.11.2 Filtering of Error Information 31
2.11.3 Collecting Error Information Using the Utility for Collecting HDLM Error Information (DLMgetras) 33
2.12 Cluster Support 33

Chapter 3 Creating an HDLM Environment 35

3.1 System Requirements 35
3.1.1 Applicable Hardware Models 35
3.1.2 Related programs 36
3.1.3 Memory and Disk Space Requirements 47
3.1.4 The number of paths guaranteed by HDLM 48
3.2 Flow for Creating an HDLM Environment 48
3.3 HDLM installation types 49
3.4 Notes on Creating an HDLM Environment 50
3.4.1 Notes on Hardware Settings 50
3.4.2 Notes on Installation 50
3.4.3 Notes on related software 54
3.4.4 Notes on command execution 55
3.4.5 Notes on the disk label 55
3.5 Installing HDLM 55

3.5.1 Preparations for a New Installation of HDLM 56
3.5.2 Performing a New Installation of HDLM (when Sun Cluster is not being used)63
3.5.3 Performing a New Installation of HDLM (when Sun Cluster is being used).... 68
3.5.4 Using the HDLM device unconfiguration function when performing a new installation of HDLM 75
3.5.5 Preparations for an Upgrade Installation or Re-installation of HDLM..... 77
3.5.6 Performing an update Performing an Upgrade Installation or Re-installation of HDLM..... 77
3.6 Configuring a Boot Disk Environment 80
3.6.1 Overview of Configuring a Boot Disk Environment..... 80
3.6.2 Procedure for Configuring a Boot Disk Environment 80
3.7 Migrating from a Boot Disk Environment to the Local Boot Disk Environment 92
3.8 Checking the Path Configuration 98
3.9 Setting Up HDLM Functions 99
3.9.1 Checking the current settings..... 99
3.9.2 Setting up HDLM functions 99
3.9.3 Checking the new settings..... 105
3.10 Setting up Integrated traces 106
3.10.1 Notes on Using the Hitachi Network Objectplaza Trace Library 107
3.10.2 Displaying the Hitachi Network Objectplaza Trace Library setup menu 107
3.10.3 Changing the size of integrated trace files 108
3.10.4 Changing the number of integrated trace files 109
3.10.5 Changing the buffer size per monitoring interval duration 109
3.10.6 Adjusting the number of messages to be output per monitoring interval .. 110
3.10.7 Finishing the Hitachi Network Objectplaza Trace Library settings 112
3.10.8 Applying the Hitachi Network Objectplaza Trace Library settings..... 112
3.11 Creating File Systems for HDLM (when volume management software is not used) 113
3.12 Setting Up VxVM..... 114
3.12.1 Creating a disk group..... 114
3.12.2 Creating VxVM volumes 117
3.12.3 Removing devices from VxVM 117
3.12.4 Introducing VxVM while using HDLM 131
3.12.5 Linking VxVM and Sun Cluster..... 131
3.13 Setting Up SDS 133
3.13.1 Notes..... 133
3.13.2 Registering HDLM devices 134
3.14 Setting Up SVM 136
3.14.1 Notes..... 136
3.14.2 Registering HDLM devices 137
3.15 Setting Up VCS..... 139
3.16 Uninstalling HDLM..... 140
3.16.1 Overview of HDLM uninstallation..... 140
3.16.2 Preparations for HDLM uninstallation..... 140
3.16.3 Uninstalling HDLM..... 147
3.16.4 Settings needed after HDLM uninstallation 149
3.16.5 Uninstalling Hitachi Network Objectplaza Trace Library (HNTRLib2) 151
3.16.6 Uninstalling Hitachi Network Objectplaza Trace Library (HNTRLib)..... 151

Chapter 4 HDLM Operation.....153

4.1 Notes on using HDLM 153

4.1.1 Displaying path information 153
4.1.2 When a path error is detected 153
4.1.3 iostat command..... 153
4.1.4 Storage subsystem..... 154
4.1.5 Command execution 154
4.1.6 Using Sun HBA..... 154
4.1.7 Starting the host, system, and HDLM in a Solaris 10 environment 155
4.1.8 Upgrading Solaris 155
4.1.9 Operation in single-user mode 155
4.1.10 Maintenance tasks on devices connected by paths in the boot disk environment..... 157
4.2 HDLM operations using commands 157
4.2.1 Notes on using commands..... 157
4.2.2 Viewing path information 157
4.2.3 Changing the status of paths 158
4.2.4 Viewing LU information 159
4.2.5 Displaying corresponding information about an HDLM device, sd or ssd device, and LDEV..... 160
4.2.6 Initializing statistical information for paths 161
4.2.7 Viewing and setting up the operating environment..... 161
4.2.8 Viewing license information 163
4.2.9 Updating the license 163
4.2.10 Viewing HDLM version information 163
4.2.11 Viewing HDLM component information 164
4.3 Starting and stopping the HDLM manager 165
4.3.1 Starting the HDLM manager..... 165
4.3.2 Stopping the HDLM manager..... 165
4.4 Changing the configuration of the HDLM operating environment..... 166
4.4.1 Precautions regarding changes to the configuration of an HDLM operating environment..... 166
4.4.2 Overview of reconfiguring the HDLM device 169
4.4.3 Adding a new logical unit 170
4.4.4 Configuration changes such as deleting a logical unit..... 184
4.4.5 Specifying whether a logical unit is to be managed by HDLM (when not using Sun Cluster) 189
4.4.6 Specifying whether a logical unit is to be managed by HDLM (when using Sun Cluster) 193

Chapter 5 Troubleshooting 209

5.1 Information collected by using the DLMgetras utility for collecting HDLM error information 209
5.2 CheckingError Information in Messages 209
5.3 Actions Taken for a Path Error 210
5.3.1 Examine the Messages 212
5.3.2 Obtain Path Information 212
5.3.3 Identify the Error Path 212
5.3.4 Narrow Down the Hardware That May Have Caused the Error 212
5.3.5 Identify the Error Location and Correct any Hardware Error..... 212
5.3.6 Place the Path Online..... 212
5.4 Actions to Take for a Path Error in a Boot Disk Environment..... 213

5.4.1 Path Errors During Boot Processing 213
5.4.2 Path Errors After Boot Processing Completes 214
5.5 Actions Taken for a Program Error 214
5.5.1 Examine the Messages..... 215
5.5.2 Obtain Program Information 215
5.5.3 Actions to Take for Program Errors 215
5.5.4 Contact Your HDLM Vendor or Maintenance Company 216
5.6 Actions Taken for Other Errors..... 216

Chapter 6 Command Reference 217

6.1 Overview of the HDLM Command (dlnkmgr) 217
6.2 clear (Clears the Path Statistics to the Initial Value) 218
6.2.1 Format..... 218
6.2.2 Parameters..... 218
6.3 help (Displays the Operation Format) 219
6.3.1 Format..... 219
6.3.2 Parameter..... 219
6.4 offline (Places a Path or Paths Offline)..... 221
6.4.1 Format..... 221
6.4.2 Parameters..... 222
6.5 online (Places a Path or Paths Online)..... 225
6.5.1 Format..... 225
6.5.2 Parameters..... 226
6.6 set (Sets the Operating Environment) 229
6.6.1 Format..... 229
6.6.2 Parameters..... 230
6.7 view (Displays Information) 239
6.7.1 Format..... 240
6.7.2 Parameters..... 241
6.8 monitor (Displays I/O information at a specified interval)..... 265
6.8.1 Format..... 265
6.8.2 Parameters..... 266

Chapter 7 Utility Reference..... 271

7.1 Overview of the utilities 271
7.2 DLMgetras (utility for collecting HDLM error information) 272
7.2.1 Format..... 272
7.2.2 Parameters..... 272
7.2.3 List of collected error information..... 274
7.3 dlmcfgmgr (utility for managing the HDLM configuration) 281
7.3.1 Format..... 281
7.4 dlmlisthdev (utility for assisting HDLM transitions)..... 283
7.4.1 Format..... 283
7.4.2 Parameters..... 283
7.5 dlmsetconf (utility for creating the HDLM driver configuration definition file) 285
7.5.1 Format..... 286
7.5.2 Parameters..... 286
7.6 dlmvxexclude (utility for assisting creation of the VxVM configuration file)..... 290
7.6.1 Format..... 290

7.6.2 Parameters 290

Chapter 8 Messages 293

8.1 Before Viewing the List of Messages 293
8.1.1 Format and Meaning of Message IDs 293
8.1.2 Terms Used in Messages and Message Explanations 294
8.1.3 Components That Output Messages to Syslog 294
8.2 HDLM Command (dlnkmgr and Operations) Messages 295
8.3 HDLM API Messages 312
8.4 HDLM Manager Messages 314
8.5 HDLM Driver (Filter Component) Messages or HDLM Nexus Driver Messages 320
8.6 HDLM Alert Driver Messages 326
8.7 HDLM Driver (Core Logic Component)Messages 328
8.8 HDLM Management Target Messages 329
8.9 HDLM Installation Program Messages 331
8.10 Messages from the DLMgetras Utility for Collecting HDLM Error Information 342
8.11 Messages from the dlmsetconf Utility for Creating the HDLM Driver Configuration Definition File or the dlmlisthdev Utility for assisting HDLM Transitions 349
8.12 Messages from the dlmcfgmgr Utility for Managing the HDLM Configuration 356
8.13 Messages from the dlmvxexclude Utility for Assisting the Creation of the VxVM Configuration File 363
8.14 Messages from HDLM remote access interface 368
8.15 Return Codes for the HDLM Remote Access Interface 371

Appendix A Differences in Functionality Between HDLM Versions 377

A.1 Differences in Functionality Between HDLM 5.9 and Earlier Versions 377
A.2 Precautions on Differences in Functionality Between HDLM 5.6.1 or earlier and HDLM 5.6.2 or later 377

Glossary 379

List of Figures

Figure 1.1 Connections between hosts and storage subsystems 2

Figure 2.1 HDLM System Configuration 7
Figure 2.2 LU configuration recognized by the host after HDLM installation 9
Figure 2.3 HDLM Program Configuration 10
Figure 2.4 Position of the HDLM Driver and HDLM Devices 12
Figure 2.5 I/O Flow When the Load Balancing Function is Not Being Used 14
Figure 2.6 I/O Flow When the Load Balancing Function Is Being Used 15
Figure 2.7 Overview of Load Balancing 16
Figure 2.8 Path Switching 20
Figure 2.9 Path Status Transitions 23
Figure 2.10 Action When an Intermittent Error Occurs in the Path 26
Figure 2.11 Action When an Intermittent Error Does Not Occur in the Path 26
Figure 2.12 Action When the Conditions for the Intermittent Error Are Changed During Error Monitoring 27
Figure 2.13 Data Flow When Collecting Error Information 30

Figure 3.1 Flow of HDLM Environment Setup 49
Figure 3.2 Execution Results of the mount -p Command 57
Figure 3.3 An example of how to edit the /etc/vfstab file 57
Figure 3.4 Example of an /opt/DynamicLinkManager/config/dlmfdrv.unconf file (before editing) 76
Figure 3.5 Example of an /opt/DynamicLinkManager/config/dlmfdrv.unconf file (after editing) 76
Figure 3.6 Flow of setting up the boot disk environment when using an existing HDLM environment as the pre-migration environment 81
Figure 3.7 Flow of setting up the boot disk environment when installing HDLM in the existing local boot disk environment 83
Figure 3.8 Flow of setting up the boot disk environment when using an existing boot disk environment without HDLM installed as the pre-migration environment 85
Figure 3.9 Flow of setting up the boot disk environment when building a new pre-migration environment 87
Figure 3.10 Migration flow when coping the file system from a boot disk environment to the local boot disk environment 94
Figure 3.11 An output example for the mount -p command 114
Figure 3.12 An example of how to edit the /etc/vfstab file 114
Figure 3.13 Devices That Are Recognized by VxVM before They Are Removed from VxVM 118
Figure 3.14 Devices Recognized by VxVM after Removal of Devices from VxVM 119
Figure 3.15 When HBA ports on the alternate path differ from those on the primary path connected to the Hitachi storage subsystem 120
Figure 3.16 Example of a description in the /etc/vx/vxvm.exclude (controllers section) 121
Figure 3.17 When the HBA Port of the Physical Path Connected to a Hitachi Storage Subsystem Is the Same As the One Connected to a Third-party Storage Subsystem 124
Figure 3.18 Example of a description in the /etc/vx/vxvm.exclude (paths section) 125
Figure 3.19 The /etc/vx/vxvm.exclude file (before editing) 130
Figure 3.20 The /etc/vx/vxvm.exclude file (after editing) 130

Figure 3.21 Overview of HDLM uninstallation 140

Figure 4.1 Example of executing the format command 172
Figure 4.2 Example of executing the format command 174
Figure 4.3 Example of executing the format command 180
Figure 4.4 Example of executing the format command 182
Figure 4.5 Example of executing the format command 191
Figure 4.6 The /opt/DynamicLinkManager/config/dlmfdrv.unconf file (before editing) . 192
Figure 4.7 Example of editing the /opt/DynamicLinkManager/config/dlmfdrv.unconf file192
Figure 4.8 The /opt/DynamicLinkManager/config/dlmfdrv.unconf file (before editing) . 194
Figure 4.9 The /opt/DynamicLinkManager/config/dlmfdrv.unconf file (after editing)... 194
Figure 4.10 Result output by executing the dlmvxexclude utility 195
Figure 4.11 Coding in the /etc/vx/vxvm.exclude file 195
Figure 4.12 Example of executing the format command 197
Figure 4.13 Example of executing the format command 199
Figure 4.14 The /opt/DynamicLinkManager/config/dlmfdrv.unconf file (before editing) . 200
Figure 4.15 The /opt/DynamicLinkManager/config/dlmfdrv.unconf file (after editing)... 200
Figure 4.16 Example of executing the format command 204
Figure 4.17 The/opt/DynamicLinkManager/config/dlmfdrv.unconf file (before editing).. 205
Figure 4.18 The/opt/DynamicLinkManager/config/dlmfdrv.unconf file (after editing).... 205

Figure 5.1 Error Location When the KAPL08022-E Message is Output 211
Figure 5.2 Troubleshooting Procedure When a Path Error Occurs 211
Figure 5.3 Troubleshooting Procedure When a Program Error Occurs 215

Figure 7.1 Coding example of a file for defining the information to be collected..... 273

List of Tables

Table 2.1 HDLM System Components 7
Table 2.2 Components Recognized by the Host 9
Table 2.3 Function of HDLM Programs 10
Table 2.4 Algorithm for Load Balancing 17
Table 2.5 When User Operations Change the Intermittent Error Information 28
Table 2.6 Types of Error Information 31
Table 2.7 Error Levels 32

Table3.1 Applicable OSs for the host 35
Table 3.2 Supported cluster software versions 36
Table 3.3 Combinations of related programs supported by HDLM 38
Table 3.4 Programs required to create an Oracle9i RAC environment (for Solaris 8) 41
Table 3.5 Programs required to create an Oracle9i RAC environment (for Solaris 9) 42
Table 3.6 Patches required to create an Oracle9i RAC environment 43
Table 3.7 Programs required to create an Oracle RAC 10g environment (For Solaris 8 or 9) 43
Table 3.8 Programs required to create an Oracle RAC 10g environment (For Solaris 10) 44
Table 3.9 Patches required to create an Oracle RAC 10g environment (For Solaris 8) 45
Table 3.10 Patches required to create an Oracle RAC 10g environment (For Solaris 9) 46
Table 3.11 Patches required to create an Oracle RAC 10g environment (For Solaris 10) 46
Table 3.12 Disk space requirements for a host 47
Table 3.13 Maximum number of LUs and paths 48
Table 3.14 Support for the Solaris zones 53
Table 3.15 Information displayed by the dlmlisthdev utility for assisting HDLM transitions 67
Table 3.16 Information displayed by the dlmlisthdev utility for assisting HDLM transitions 73
Table 3.17 The default and recommended values for each function 100
Table 3.18 Values for the error log collection level setting 103
Table 3.19 Values for the trace level setting 103
Table 3.20 Default and recommended values for the integrated trace file settings 107
Table 3.21 When setup of vxvm.exclude is unnecessary 123
Table 3.22 When setup of vxvm.exclude is unnecessary 127

Table 6.1 Operations of the dlnkmgr Command 217
Table 6.2 Default values and recommended values 230
Table 6.3 Values of the Error Log Collection Level 231
Table 6.4 Trace Level Values 232
Table 6.5 Relationship Between the Setting Status for Automatic Failback and Intermittent Error Monitoring and the Executable Operations 237
Table 6.6 License Key Types 238
Table 6.7 Displayed program information 242
Table 6.8 Items that can be selected by the `-path -item` parameter and parameter values that can be specified after the `-path -item` parameter 248
Table 6.9 Displayed path information 250

Table 6.10 Elements of a path name 253
Table 6.11 Product ID displayed by the view -path operation 254
Table 6.12 Items that can be added by using the -lu -item parameter and parameters that can be specified after the -lu -item parameter 256
Table 6.13 Items that can be added by using the -lu -c -item parameter and parameters that can be specified after the -lu -c -item parameter 258
Table 6.14 Displayed LU information 259
Table 6.15 Items displayed as HBA port information 262
Table 6.16 Items displayed as CHA port information 263
Table 6.17 To display corresponding information about an HDLM device, sd or ssd device, and LDEV 264
Table 6.18 Recommended values for the -intvl parameter when there are 500 or more paths 266
Table 6.19 Items displayed as I/O information for each HBA port 267
Table 6.20 Items displayed as I/O information for each CHA port 268

Table 7.1 Information stored in the getras.tar.Z file 274
Table 7.2 Configuration changes applicable and not applicable to the dlmcfgmgr utility for managing the HDLM configuration 281
Table 7.3 Information displayed by the dlmlisthdev utility for assisting HDLM transitions284

Table 8.1 Format and Meaning of the Message ID KAPL*mmnnn-l* 294
Table 8.2 Terms Used in Messages and Message Explanations 294

Table A.1 Differences in LDEV Display Names Between Versions 377

Chapter 1 Overview of HDLM

HDLM manages paths between a host and storage subsystem. HDLM distributes the load across paths and switches to another path if there is a failure in a path being used, thus improving system reliability.

This chapter gives an overview of HDLM and describes its features.

- 1.1 What is HDLM?
- 1.2 HDLM Features

1.1 What is HDLM?

The widespread use of data warehousing and increasing use of multimedia data have increased the need for high-speed processing of large volumes of data on networks. To satisfy this need, networks dedicated to data transfer, such as a SAN, are now being used to provide access to storage subsystems.

The HDLM software manages access paths to storage subsystems. HDLM provides functionality for distributing the load across paths and switching to another path if there is a failure in a path being used, thus improving system availability and reliability.

Figure 1.1 illustrates the connections between hosts and storage subsystems. A server on which HDLM is installed is called a *host* or *management-target host*.

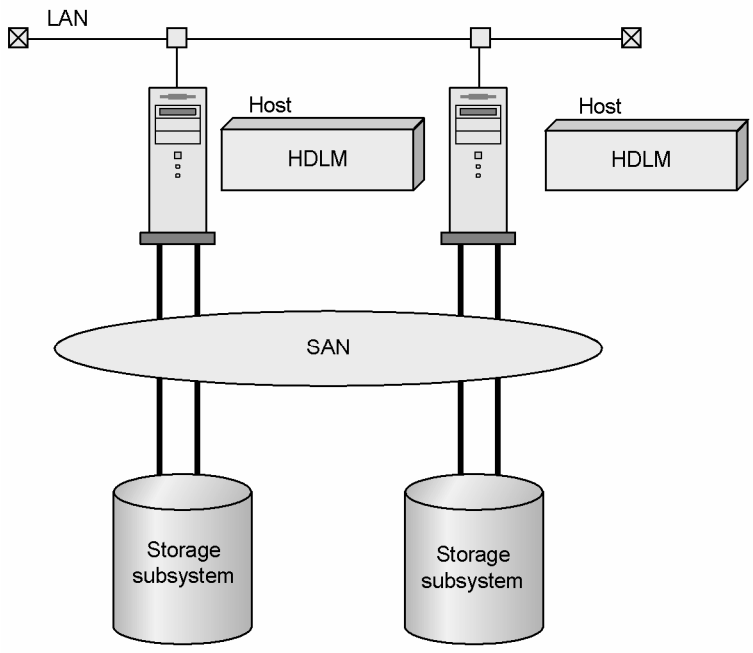

Figure 1.1 Connections between hosts and storage subsystems

HDLM supports the following storage subsystems:

- TagmaStore USP
- Universal Storage Platform V
- Lightning 9900V Series
- Lightning 9900 Series
- Freedom Storage 7700E

- TagmaStore AMS/WMS series
- Thunder 9500V Series
- Thunder 9200
- Freedom Storage 5800
- Freedom Storage 5700E

1.2 HDLM Features

HDLM has the following features.

Distributes the load across paths (*load balancing*)

When multiple paths connect a host and storage subsystem, HDLM distributes the load across the paths. This prevents a heavily loaded path from affecting processing speed.

For details on load balancing, see 2.7.

Allows processing to continue if there is a failure (*failover*)

When multiple paths connect a host and storage subsystem, HDLM automatically switches to another path if there is a failure in a path being used. This allows processing to continue without being affected by the failure.

For details on failover, see 2.8.

Allows you to place online a path recovered from an error (*failback*)

When a path recovers from an error, HDLM can place the path online. This enables the maximum number of paths to be online, which in turn enables HDLM to distribute the load for each path.

Route failback can be performed manually or automatically. In automatic failback, HDLM automatically restores the route to the active state after the user has corrected hardware problems in the route.

For details on failback, see 2.8.

Automatically checks the path status at regular intervals (*path health checking*)

HDLM can detect errors by checking the status of the paths at user-defined time intervals. This allows you to check for any existing path errors and to resolve them accordingly.

For details on path health checking, see 2.10.

Chapter 2 HDLM Functions

This chapter describes the HDLM functionality. Initially, the HDLM management targets, system configuration, and basic terms are described. The rest of this chapter describes the functionality including load distribution across paths and path switching.

- 2.1 Devices Managed by HDLM
- 2.2 System Configuration
- 2.3 LU Configuration
- 2.4 Program Configuration
- 2.5 Position of the HDLM Driver and HDLM Device
- 2.6 Logical Device Files for HDLM Devices
- 2.7 Load Distribution Using Load Balancing
- 2.8 Failover and Failback Using Path Switching
- 2.9 Monitoring Intermittent Errors (Functionality When Automatic Failback Is Used)
- 2.10 Detecting Errors by Using Path Health Checking
- 2.11 Error Management
- 2.12 Cluster Support

2.1 Devices Managed by HDLM

Devices that can and cannot be managed by HDLM are shown below. The devices that can be managed by HDLM are called *HDLM management-target devices*.

Devices that HDLM can manage:

HDLM can manage SCSI devices (sd or ssd devices) in Hitachi storage subsystems.

- Hitachi storage subsystem boot disks
- Hitachi storage subsystem dump devices
- Hitachi storage subsystem swap devices

Devices that HDLM cannot manage:

- SCSI devices (sd or ssd devices) other than those for Hitachi storage subsystems
- Built-in disks in a host
- Devices other than disks (tape devices, etc.)
- Hitachi storage subsystem command devices, such as Hitachi RAID Manager command devices

2.2 System Configuration

HDLM manages routes between a host and a storage subsystem by using the SCSI driver (sd or ssd driver). The host and storage subsystems are connected using SAN with fiber cables or SCSI cables. The cable port on the host is a *host bus adapter* (HBA). The cable port on the storage subsystem is a *port* (P) on a *channel adapter* (CHA).

A *logical unit* (LU) contained in a storage subsystem is the target of input to, or output from, the host. You can divide an LU into multiple areas. Each area after the division is called a *Dev*. The Dev is equivalent to a slice or partition. A route that connects a host and an LU is called a *physical path*, and a route that connects a host and a Dev is called a *path*. When an LU has been divided into multiple Devs, the number of paths set to the LU is equal to the number that is found by multiplying the number of physical paths by the number of Devs in the LU.

HDLM assigns an ID to each physical path and manages paths on a physical-path basis. Since you do not need to be aware of the difference between physical paths and paths to operate HDLM, the following descriptions may simply refer to paths, without distinguishing between physical paths and paths. The ID that HDLM assigns to each path (physical path) is called a *AutoPATH_ID*. A path is also sometimes called a *managed object*.

Figure 2.1 shows the HDLM system configuration.

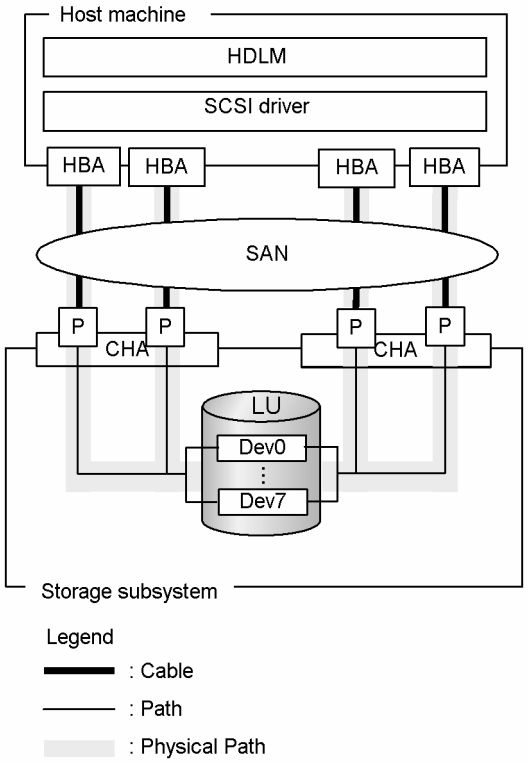

Figure 2.1 HDLM System Configuration

Table 2.1 lists the HDLM system components.

Table 2.1 HDLM System Components

Components	Description
HBA	A host bus adapter. This serves as a cable port on the host.
SAN	A dedicated network that is used for data transfer between the host machine and storage subsystems.
CHA	A channel adapter.
P	A port on a channel adapter. This serves as a cable port on a storage subsystem.
LU	A logical unit (a logical volume defined on the storage subsystem). This serves as the target of input or output operations from the host.
Dev	An area (slice or partition) that is created when an LU is divided.
Physical path	A route that connects a host and an LU.
Path	A route that connects a host and a Dev.

2.3 LU Configuration

When you install HDLM, the LU configuration recognized by the host changes as follows:

Before installing HDLM:

The host recognizes that an sd or ssd device is connected to each physical path.

Thus, a single LU in the storage subsystem is recognized as the same number of LUs as that of physical paths.

After installing HDLM:

An HDLM device that corresponds one-to-one with the Dev in an LU in the storage subsystem is created above an sd or ssd device.

Thus, from the host, LUs in the storage subsystem are also recognized as one LU regardless the number of physical paths.

An LU recognized by a host after HDLM installation, is called a *host LU* (HLU). The areas in a host LU that correspond to the Dev (slice or partition) in a storage subsystem LU are called *host devices* (HDev).

On a system using HDLM, the logical device file for the HDLM device is used to access the target LU instead of the logical device file for the sd or ssd device.

The logical device files for sd or ssd are deleted by HDLM.

Figure 2.2 show the LU configuration recognized by the host after HDLM installation.

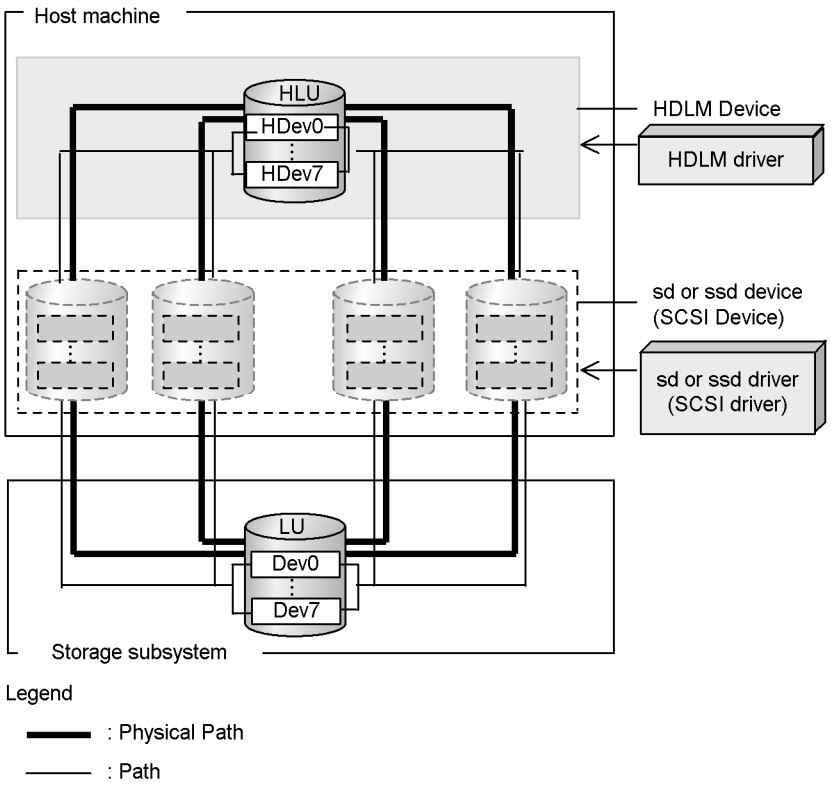

Figure 2.2 LU configuration recognized by the host after HDLM installation

Table 2.2 lists the components recognized by the host.

Table 2.2 Components Recognized by the Host

Components	Description
HLU	An LU that the host recognizes via the HDLM driver. It is called a *host LU*. No matter how many physical paths exist, one host LU is recognized for one LU in the storage subsystem.
HDev	A Dev (a slice or partition) in an LU that the host recognizes via the HDLM driver. It is called a *host device*. No matter how many physical paths exist, one host device is recognized for one Dev in the storage subsystem.

2.4 Program Configuration

HDLM runs as a combination of programs. Because these programs correspond to the HDLM operations, you need to understand the name, position, and role of each program.

Figure 2.3 shows the HDLM program configuration.

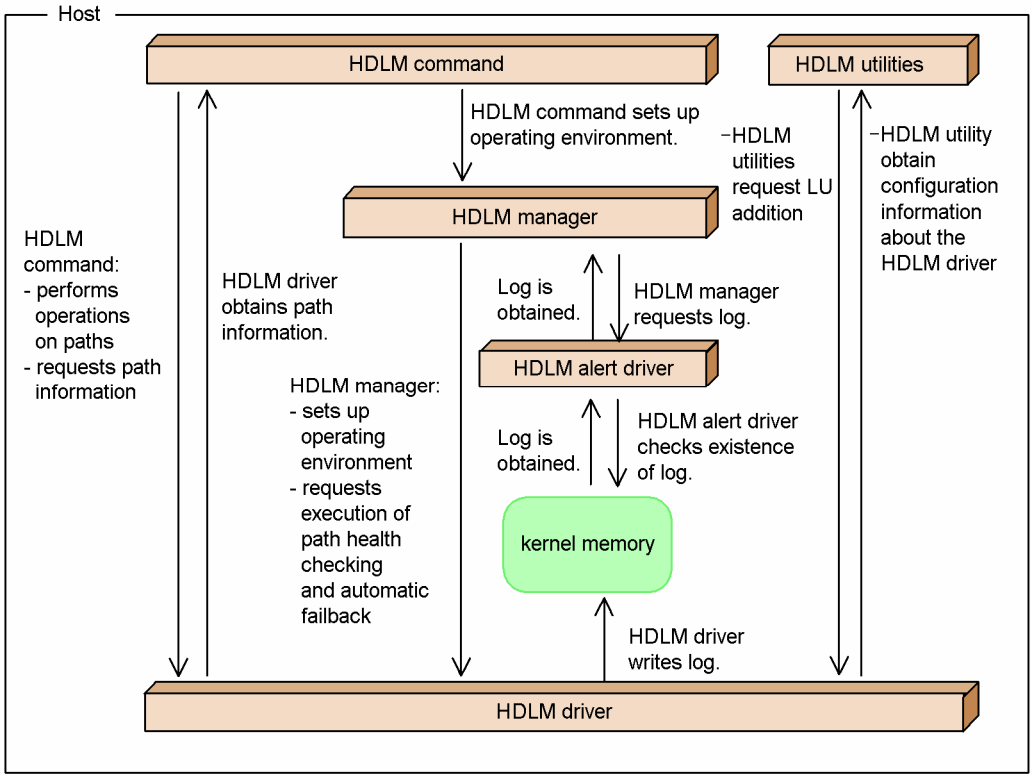

Figure 2.3 HDLM Program Configuration

Table 2.3 lists and describes the functionality of these programs.

Table 2.3 Function of HDLM Programs

Program name	Function
HDLM command	Provides the `dlnkmgr` command to enable you to: ▪ Manage paths ▪ Display error information ▪ Set up the HDLM operating environment
HDLM utility	Provides the HDLM utility to enable you to: ▪ Collect error information ▪ Add a new LU (reconfiguring an HDLM device dynamically) ▪ Create an HDLM driver configuration definition file (/kernel/drv/dlmfdrv.conf) ▪ Create a correspondence table of logical device files when migrating to HDLM 5.9 ▪ Support the creation of a VxVM configuration file
HDLM manager	Provides the HDLM manager to enable you to: ▪ Configure the HDLM operating environment ▪ Request the execution of path health checks and automatic failback ▪ Collect error log data
HDLM alert driver	Reports the log information collected by the HDLM driver to the HDLM manager. The driver name is `dlmadrv`.

Program name	Function
HDLM driver	Controls HDLM functionality, manages paths, and detects errors. The HDLM driver consists of the following: ▪ Core logic component Controls the basic functionality of HDLM. ▪ Filter component Sends and receives I/O. The driver name is `dlmfdrv`. ▪ HDLM nexus driver Performs operations such as reserving controller numbers for logical device files of the HDLM device, and managing HDLM driver instances for each HBA port. The driver name is `dlmndrv`.

2.5 Position of the HDLM Driver and HDLM Device

The HDLM driver is positioned above the SCSI driver. Each application on the host uses the HDLM device (logical device file) created by HDLM, to access LUs in the storage subsystem.

Figure 2.4 shows the position of the HDLM driver and HDLM devices.

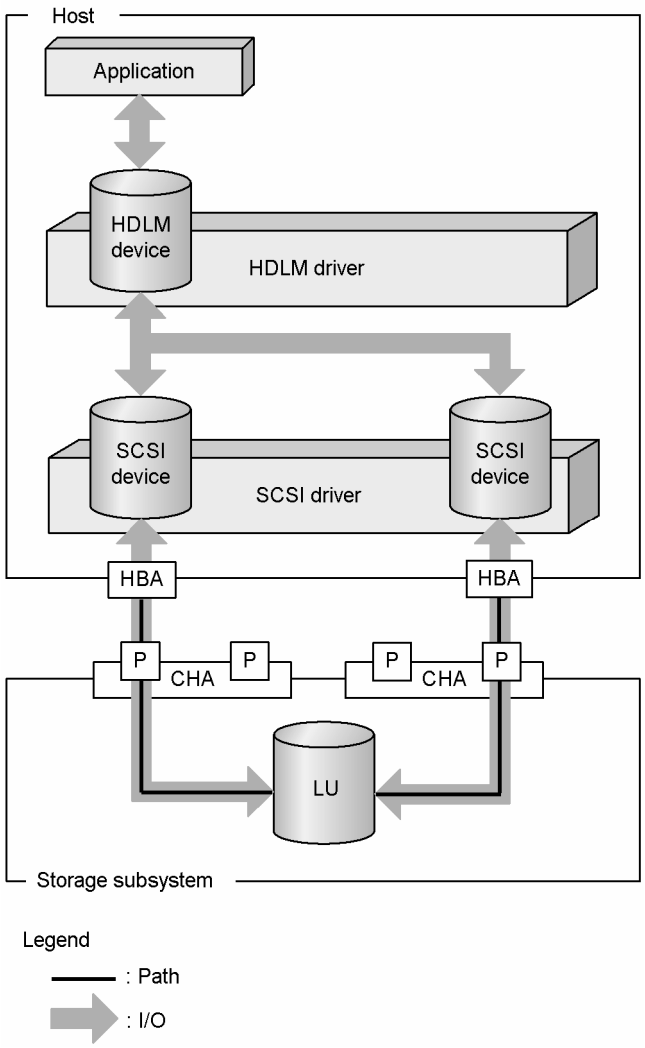

Figure 2.4 Position of the HDLM Driver and HDLM Devices

2.6 Logical Device Files for HDLM Devices

When you install HDLM, a logical device file to be used by HDLM will be created for each LU on per-Dev (slice) basis. Setting this logical device file name in an application, such as volume management software, enables the application to access an LU by using the HDLM function.

The logical device files existing before HDLM installation (the logical device files of an sd or ssd) will be deleted.

The following explains the names and locations of the logical device files for HDLM devices

Logical device file names for HDLM devices

The logical device file name of an HDLM device is a changed version of the controller number of the logical device file name of the sd or ssd device. For example, let us assume that an LU has two physical paths, and for one of the Dev (slices) in that LU, the corresponding logical device file names of the sd or ssd devices are `c2t1d1s0` and `c3t2d1s0`. In this case, when you install HDLM, these logical device files will be deleted. Then, a logical device file that has a different controller number, such as `c4t1d1s0`, is created for the HDLM device.

The following explains each part of the logical device file name format `c`*U*`t`*X*`d`*Y*`s`*Z*:

U

The controller number reserved by HDLM using a nexus driver

X

The target ID or WWN (World Wide Name) of the sd or ssd device that corresponds to the HDLM device

Y

The LUN of the sd or ssd device that corresponds to the HDLM device

Z

The device slice number of the sd or ssd device that corresponds to the HDLM device

Locations of logical device files for HDLM devices

Block logical device files for HDLM devices are created in `/dev/dsk`. Character logical device files for HDLM devices are created in `/dev/rdsk`.

2.7 Load Distribution Using Load Balancing

When the system contains multiple paths to an LU, HDLM can distribute the load across the paths by using multiple paths for I/O. This function is called *load balancing*, and it prevents a heavily loaded path from affecting the performance of the entire system.

Note that some I/O operations managed by HDLM can be distributed to each path, while others cannot. Therefore, even though load balancing function is used, I/O operations may not be equally allocated to each path.

Figure 2.5 shows the I/O flow when the load balancing function is not being used. Figure 2.6 shows the I/O flow when the load balancing function is being used. Both figures show an example of an I/O being issued for the same LU from multiple applications.

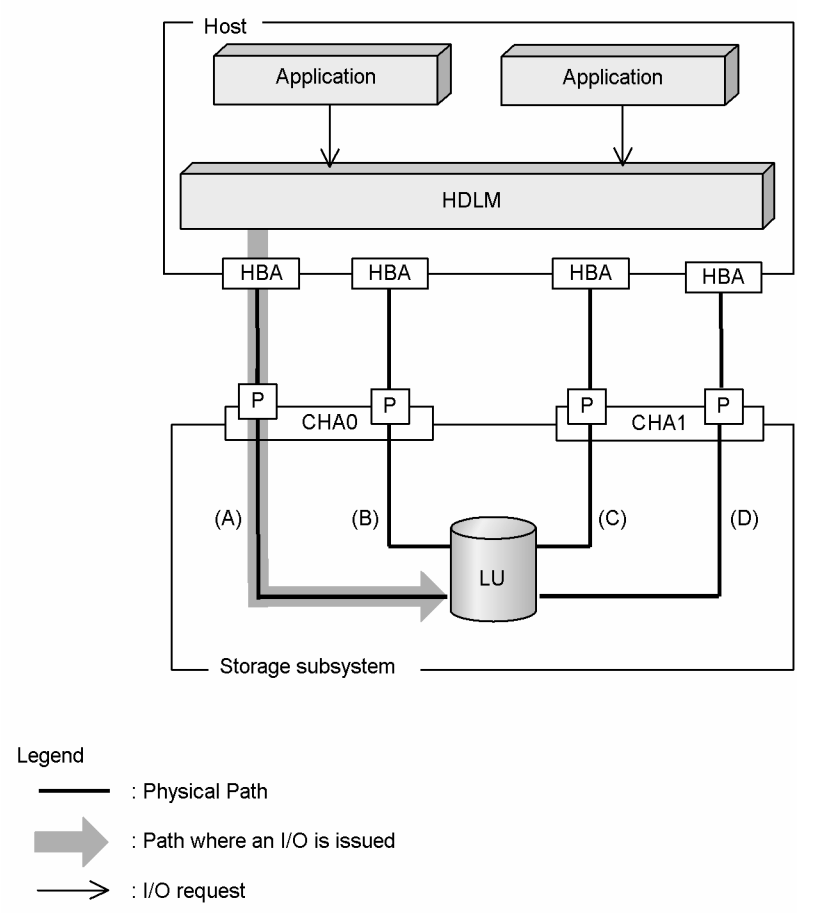

Figure 2.5 I/O Flow When the Load Balancing Function is Not Being Used

When the load balancing function is not being used, I/O operations converge on the physical path (A). The load on the physical path (A) will cause a bottleneck, which might cause deterioration of the whole system's performance.

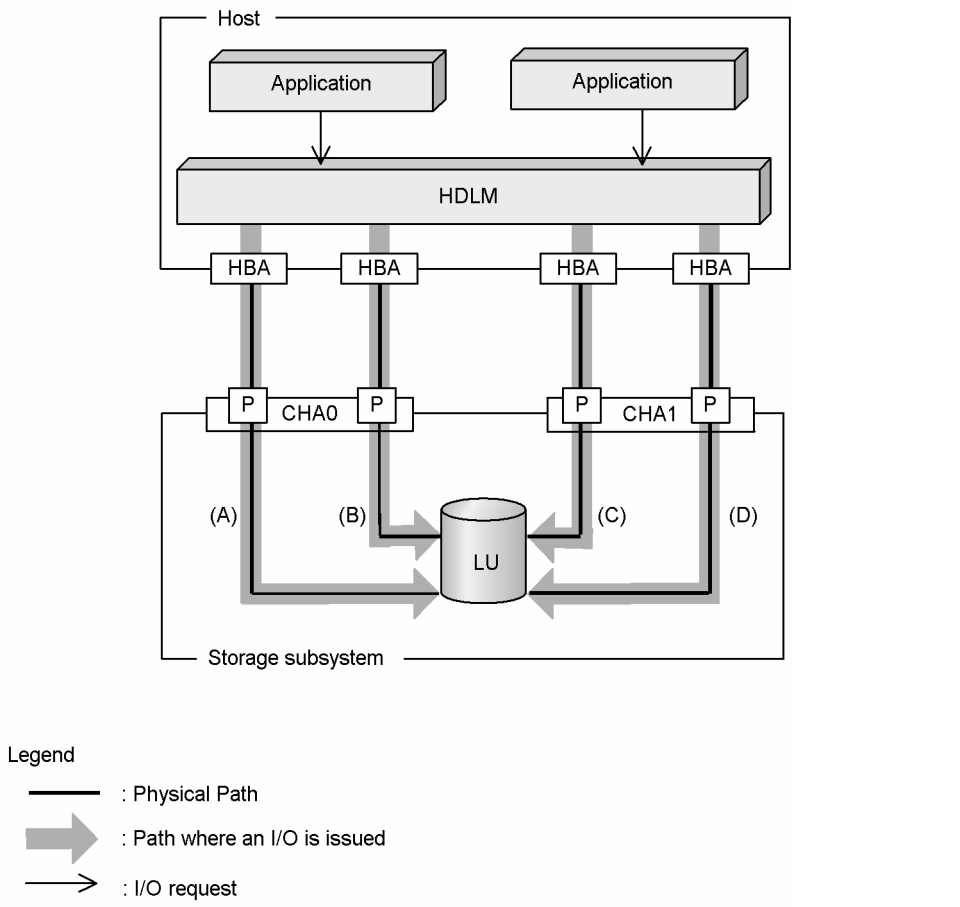

Figure 2.6 I/O Flow When the Load Balancing Function Is Being Used

When the load balancing function is being used, I/O operations are distributed to physical paths (A), (B), (C), and (D). This prevents deterioration of the whole system's performance from a bottleneck on one path.

2.7.1 appliedappliedPaths to Which Load Balancing is Applied

This section describes, for each storage subsystem, the paths to which the load balancing function is applied.

2.7.1.1 When Using the Freedom Storage 5700E, Freedom Storage 5800, Thunder 9200, Thunder 9500V Series, or TagmaStore AMS/WMS series

HDLM performs load balancing between owner paths or between non-owner paths. An *owner path* is a path that passes through the channel adapter. This path is set on the *owner controller* of the storage subsystem LU. Since the owner controller varies depending on the LU, the owner path also varies depending on the LU. A *non-owner path* is a path that uses a channel adapter other than the owner controller (a *non-owner controller*). The paths to be used are selected, in the order of owner paths and non-owner paths. To prevent performance in the entire system from deteriorating, HDLM does not perform load balancing between owner paths and non-owner paths. When some owner paths cannot be used due to a problem such as a failure, load balancing is performed among the remaining usable owner paths. When all owner paths cannot be used, load balancing is performed among the non-owner paths.

For the example in Figure 2.7, suppose that the owner controller of the LU is CHA0. When the LU is accessed, the load is balanced between the physical paths (A) and (B) (that is, between owner paths). When the physical path (A) cannot be used due to a problem such as a failure, the LU can only be accessed via the physical path (B). When the physical paths (A) and (B) cannot be used, the load is balanced between the physical paths (C) and (D) (that is, between non-owner paths).

The following explains the mechanism of load balancing. The explanation uses the example shown in Figure 2.7.

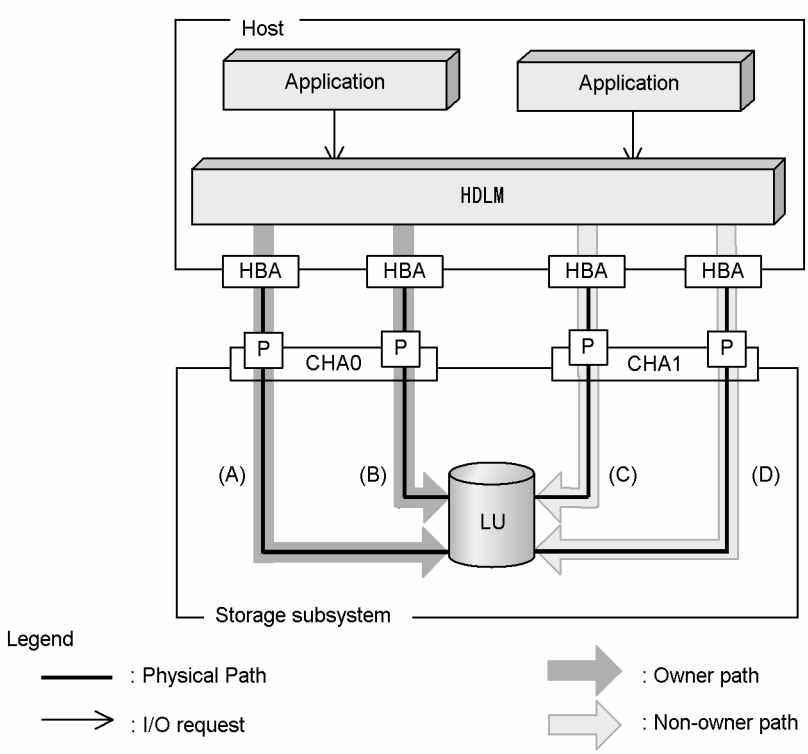

Figure 2.7 Overview of Load Balancing

2.7.1.2 When Using the Freedom Storage 7700E, Lightning 9900 Series, Lightning 9900V Series, TagmaStore USP, or Universal Storage Platform V

All the online paths are owner paths. Therefore, for the example in Figure 2.6, when the LU is accessed, the load is balanced among the physical paths (A), (B), (C), and (D). When one of the physical paths cannot be used due to a problem such as a failure, the load is balanced among the remaining physical paths.

2.7.2 Algorithms for Load Balancing

HDLM has the following two types of load balancing:

- Round robin
- Extended round robin

Round robin distributes all I/Os among multiple paths.

Extended round robin distributes I/Os to paths depending on the type of the I/O, which can be either sequential access or random access. For sequential access, I/Os are issued to a single path. For random access, I/Os will be distributed to multiple paths.

Table 2.4 describes the type of load balancing (round robin and extended round robin) for each I/O operation type.

Table 2.4 Algorithm for Load Balancing

Algorithm	For sequential access	For random access
Round robin	▪ After an I/O is issued to a path once or a certain number of times, the path is switched to the next path. ▪ The storage subsystem cache might not be fully usable. ▪ When multiple applications that request sequential access run concurrently, we recommend that you use the round robin algorithm in order to distribute I/Os across multiple paths.#	After an I/O is issued to a path once or a certain number of times, the path is switched to the next path.
Extended round robin	▪ After an I/O is issued to a path a certain number of times in succession, the path is switched to the next path. If sequential access is switched to random access before an I/O is issued to a path a certain number of times, the path is switched to the next path. ▪ The storage subsystem cache can be used. ▪ When you execute only a single application that requests sequential access, such as a batch job running at night, we recommend that you use the extended round robin algorithm.#	

#

The recommended algorithm depends on the type of applications, and the operations policy.

You can specify the load balancing function by the `dlnkmgr` command's `set` operation. For details on the `set` operation, see 6.6.

Note

Some I/O operations managed by HDLM can be distributed across all paths, and some cannot. Thus, you should be aware that even when you use the load balancing function, I/O operations cannot always be allocated uniformly across all paths.

2.8 Failover and Failback Using Path Switching

When the system contains multiple paths to an LU and an error occurs in the path being used, HDLM can switch to another normal path to allow the system to continue to operate. This function is called *failover*.

When the path in which an error occurred recovers from the error, HDLM can switch the paths so that the recovered path is used. This function is called *failback*.

Two types of failover and failback are available:

- Automatic path switching
- Manual path switching

Failover and failback change the path statuses and switch the paths. Path statuses are classified into *online* statuses and *offline* statuses. Online statuses allow the path to normally receive I/Os. Offline statuses prevent the path from receiving I/Os for the following reasons:

- An error occurred in the path.
- A user executed the HDLM command's `offline` operation.

  For details on the `offline` operation, see 6.4.

For details on path statuses and the status transitions, see 2.8.3.

Note

– When the Freedom Storage 5700E, Freedom Storage 5800, Thunder 9200, Thunder 9500V Series, or TagmaStore AMS/WMS series is being used, make sure that you set the data share mode to **ON** (**Used**).

  For details on the data share mode, see the manual for Thunder 9200, Thunder 9500V Series, or TagmaStore AMS/WMS series.

– Switching a reserved path may take several seconds. A reserved path is switched in the following cases:

  - The reserved path is placed offline.

  - An owner path is placed online when a path has been reserved while only non-owner paths are online.

2.8.1 Automatic Path Switching

This section describes the automatic failover and automatic failback functions that automatically switch a paths.

2.8.1.1 Automatic Failover

If you detect an error in a path being used, you can keep operating the system by changing the path state to offline, and using other online paths. This function is called *automatic failover*. Automatic failover is applicable to the following levels of errors that occur on a path.

Critical

A fatal error that may stop the system.

Error

A high-risk error that can be avoided by using failover or other countermeasures.

For details on error levels, see 2.11.2.

When the Freedom Storage 5700E, Freedom Storage 5800, Thunder 9200, Thunder 9500V Series, or TagmaStore AMS/WMS series is being used, HDLM selects the switching destination from the paths that access the same LU, in the order of owner paths and non-owner paths. For example, in Figure 2.8, the owner controller of LU is CHA0, and access to the LU is made only via the physical path (A). After the access path is placed offline, the first candidate for the switching destination is the physical path (B) and the second candidate is the physical path (C) or (D).

When the Freedom Storage 7700E, Lightning 9900 Series, Lightning 9900V Series, TagmaStore USP, or Universal Storage Platform V is being used, all the paths are owner paths. This means all the paths accessing the same LU can be possible switching destinations. For example, in Figure 2.8, the LU is accessed using only the physical path (A). After the access path is placed offline, the switching destination is one of the physical paths (B), (C), and (D).

Paths are switched in units of physical paths. Therefore, if an error occurs in a path, HDLM switches all the other paths that run through the same physical path.

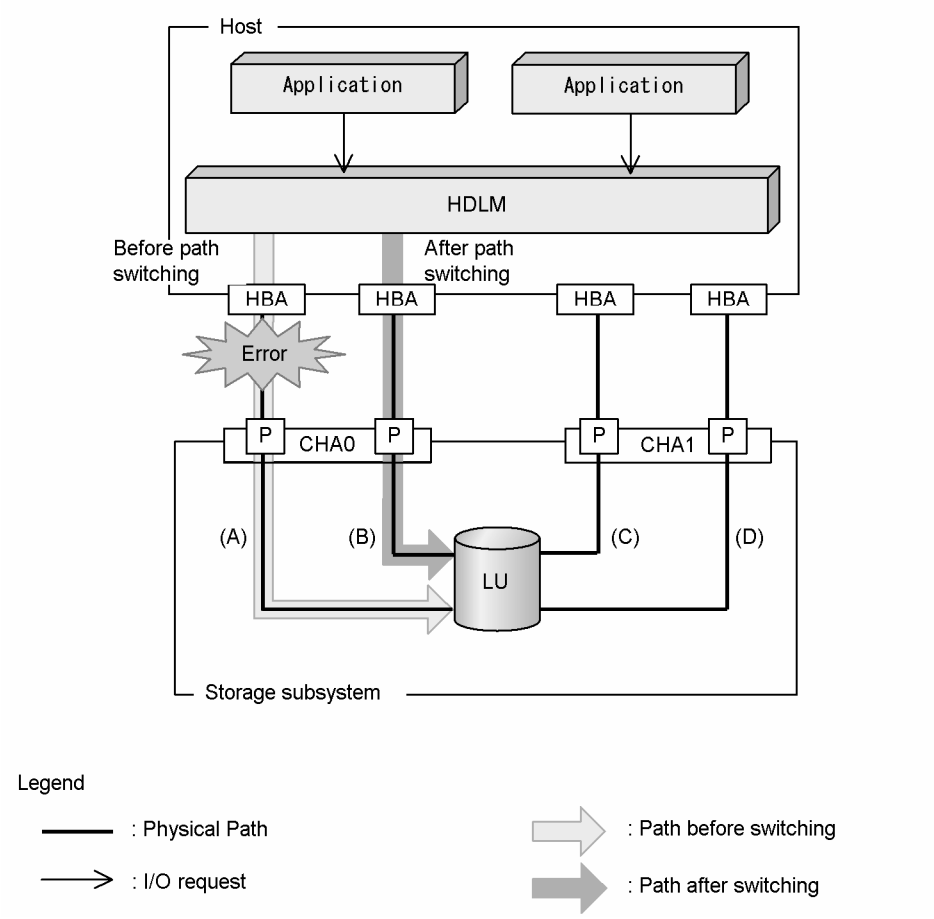

Figure 2.8 Path Switching

2.8.1.2 Automatic Failback

After a path recovers from an error, HDLM can automatically place the recovered path online. This function is called *automatic failback*.

When using this function, HDLM monitors error recovery on a regular basis.

When the Freedom Storage 5700E, Freedom Storage 5800, Thunder 9200, Thunder 9500V Series, or TagmaStore AMS/WMS series is being used, HDLM selects the path to use from online owner paths, and then from online non-owner paths. Therefore, if an owner path recovers from an error and HDLM automatically places the recovered path online while any non-owner path is in use, the path in use will be switched to the recovered owner path.

When the Freedom Storage 7700E, Lightning 9900 Series, Lightning 9900V Series, TagmaStore USP, or Universal Storage Platform V is being used, all the paths are owner paths. Therefore, if an owner path recovers from an error and HDLM automatically places the recovered path online, the path in use will not be switched to the recovered owner path.

When intermittent errors# occur in paths, the path status may frequently alternate between the online and offline status if you are using automatic failback, so the I/O performance might deteriorate. In such a case, if there is a path in which an intermittent error might be occurring, we recommend that you set up intermittent error monitoring to remove that path from those subject to automatic failback.

You can specify the automatic failback function or intermittent error monitoring by the `dlnkmgr` command's `set` operation. For details on the `set` operation, see 6.6.

#

An *intermittent error* means an error that occurs irregularly because of some reason such as a loose cable connection.

2.8.2 Manual Path Switching

You can switch a paths by manually placing a path online or offline. Manually switching a paths temporarily is useful for maintenance of the system.

You can manually place a path online or offline in the following ways:

- Execute the `dlnkmgr` command's `online` or `offline` operation.

  For details on the `online` operation, see 6.5. For details on the `offline` operation, see 6.4.

However, the status of the last path for a specific LU in offline status cannot be manually switched to offline. Also, the status of a path whose error has not been recovered cannot be switched to online.

HDLM selects the switching destination path the same way as for automatic path switching.

When the Freedom Storage 5700E, Freedom Storage 5800, Thunder 9200, Thunder 9500V Series, or TagmaStore AMS/WMS series is being used, HDLM selects the switching destination path from owner paths and then from non-owner paths. When the Freedom Storage 7700E, Lightning 9900 Series, Lightning 9900V Series, TagmaStore USP, or Universal Storage Platform V is being used, all paths that access the same LU are candidates for the switching destination path.

Paths are switched in units of physical paths. Therefore, if an error occurs in a path, all the other paths that run through the same physical path are switched.

Executing the `online` operation places the offline path online. For details on the `online` operation, see 6.5. After the path status is changed to online, HDLM selects the path to use in the same way as for automatic path switching. When the Freedom Storage 5700E, Freedom Storage 5800, Thunder 9200, Thunder 9500V Series, or TagmaStore AMS/WMS series is being used, HDLM selects the path to use from online owner paths, and then from online non-owner paths. When the Freedom Storage 7700E, Lightning 9900 Series, Lightning 9900V Series, TagmaStore USP, or Universal Storage Platform V is being used, since all the paths are owner paths, the path to use is not switched even if you change the path status to online.

2.8.3 Path Status Transition

Each of the online and offline statuses described in 2.8 is further subdivided into two statuses. The following explains the two online path statuses and the two offline path statuses.

2.8.3.1 The Online Path Status

The online path statuses are as follows:

- `Online`

  I/O can be issued normally.

- `Online(E)`

  An error has occurred in the path. Also, none of the paths that access the same LU are in the `Online` status.

  If none of the paths accessing a single LU are in the `Online` status, one of the paths will be changed to the `Online(E)` status. All the paths that access the same LU will not be in the offline statuses. This ensures access to the LU.

  The text (E) of Online(E) indicates the error attribute, which indicates that an error occurred in the path.

2.8.3.2 The Offline Path Status

The offline path statuses are as follows:

- `Offline(C)`

  The status in which I/O cannot be issued because the `offline` operation was executed. For details on the `offline` operation, see 6.4.

  The (C) indicates the command attribute, which indicates that the path was placed offline by using the command.

- `Offline(E)`

  The status in which I/O cannot be issued because an error occurred in the path.

  The (E) in Offline(E) indicates the error attribute, which indicates that an error occurred in the path.

2.8.3.3 Status Transitions of a Path

Figure 2.9 shows the status transitions of a path.

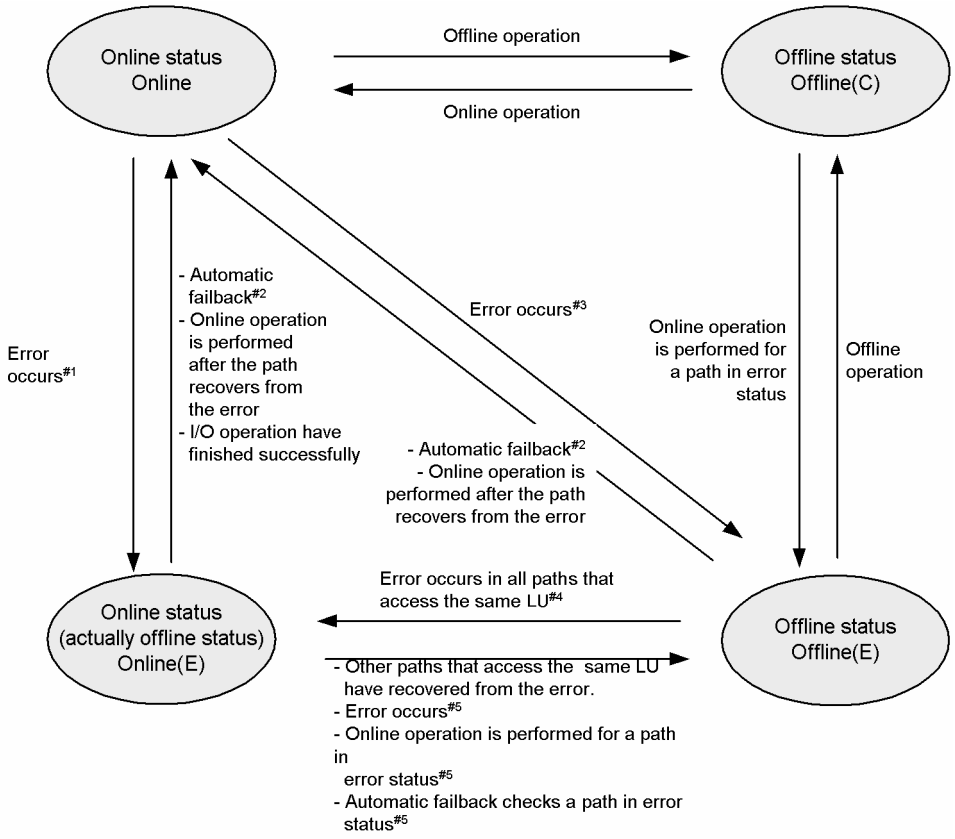

Figure 2.9 Path Status Transitions

Legend:

Online operation: Online operation performed by executing the `dlnkmgr` command's `online` operation.

Offline operation: Offline operation performed by executing the `dlnkmgr` command's `offline` operation.

#1

When no Online or Offline(E) paths exist among the paths that access the same LU.

#2

When the following conditions are satisfied, a path that has been determined to have an intermittent error also becomes subject to automatic failback:

- All the paths connected to an LU are Online(E), Offline(E), or Offline(C).
- All the paths connected to an LU have been determined to have an intermittent error.
- The processing of continuous I/O issued to an LU is successful.

#3

When an Online or Offline(E) path exists among the paths that access the same LU.

#4

One of the Offline(E) paths is changed to the Online(E) path.

#5

When an Offline(E) path exists among the paths that access the same LU.

The last available online path for each LU cannot be placed offline by executing the `offline` operation. This ensures access to the LU. For details on the `offline` operation, see 6.4.

If none of the paths accessing a single LU are in the `Online` status, one of the paths will be changed to the `Online(E)` status.

If you are using automatic failback, when the path recovers from an error, HDLM automatically places the path online.

When you are using intermittent error monitoring, the path in which the intermittent error occurred is not automatically placed even when the path recovers from the error. In such a case, place the path online manually.

Note

If there is a path failure immediately after a path is placed offline by using the `dlnkmgr` command, Offline (C) might change to Offline(E). If an offline operation was performed, wait for a fixed period of time (about 2 minute), check the path status by using the `dlnkmgr` command, and make sure that the status has changed to Offline(C). If it is Offline(E), retry the offline operation.

2.9 Monitoring Intermittent Errors (Functionality When Automatic Failback Is Used)

An intermittent error means an error that occurs irregularly because of some reason such as a loose cable connection. I/O performance might decrease when an intermittent error occurs while automatic failback is used, because automatic failback is performed repeatedly. To prevent this phenomenon, HDLM can automatically remove the path where an intermittent error is occurring from those paths subject to automatic failback. This process is called *intermittent error monitoring*.

We recommend that intermittent error monitoring be used along with automatic failback.

With intermittent error monitoring, a path in which an error occurs a specified number of times within a specified interval is determined to have an intermittent error. The path where an intermittent error occurs has an error status until the user places the path online. Automatic failback is not performed for the path. This status is called *not subject to auto failback*.

2.9.1 Checking Intermittent Errors

You can check the path in which an intermittent error occurs by using the execution result of the HDLM command's `view` operation.

For details on the `view` operation, see 6.7.

2.9.2 Setting Up Intermittent Error Monitoring

When you use the intermittent error function, you can enable or disable the function. If you enable the function, specify the monitoring conditions: the error monitoring interval, and the number of times that the error is to occur. If an error occurs in a path the specified number of times within the specified error monitoring interval, the system determines that the path has an intermittent error. For example, if you specify 30 for the error monitoring interval and 3 for the number of times that the error is to occur, the path is determined to have an intermittent error if an error occurs 3 or more times in 30 minutes.

You can set up intermittent error monitoring by executing the `dlnkmgr` command's `set` operation.

Intermittent error monitoring can be used only when automatic failback is enabled. The setting value depends on the setting value for automatic failback. For details on how to set up the setting values, see 6.6.

2.9.3 Actions for Intermittent Error Monitoring

Intermittent error monitoring is performed for each path, and it starts when a path is recovered from an error by using automatic failback.

This subsection describes the actions for intermittent error monitoring in the following cases:

- When an intermittent error occurs
- When an intermittent error does not occurs
- When the conditions for the intermittent error are changed during intermittent error monitoring

2.9.3.1 When an Intermittent Error Occurs

When an error occurs in a path the specified number of times within the specified interval, the error monitoring finishes, the path is determined to have an intermittent error, and then the path is removed from those items subject to automatic failback. The path that is removed from the paths that are subject to automatic failback has an error status until the `online` operation is performed properly. However, if the path satisfies certain conditions, it will be subject to automatic failover and the operation status will automatically change to `online`. For details on the conditions, see Figure 2.9.

Figure 2.10 shows the action for intermittent error monitoring when an intermittent error occurs. In this example, the path is determined to have an intermittent error when the error occurs 3 or more times in 30 minutes. The events occurring in one path are described on the time arrow.

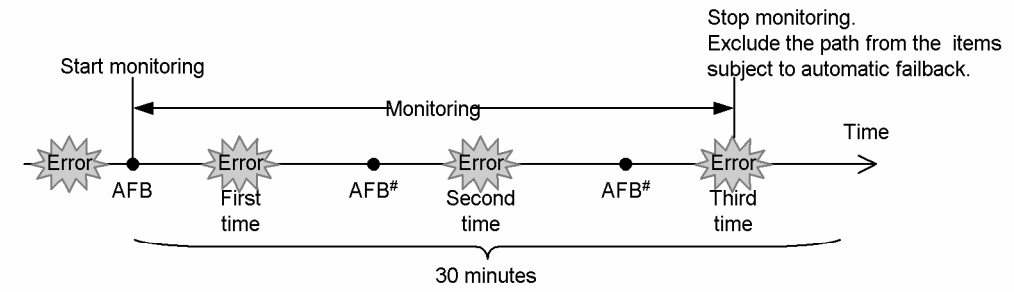

(Legend)
AFB: Indicates where the path was changed from error status to online status by automatic failback.

#
This includes online operation performed by a user.

Figure 2.10 Action When an Intermittent Error Occurs in the Path

2.9.3.2 When an Intermittent Error Does Not Occur

If an error does not occur in the path the specified number of times within the specified interval, the system determines that the path does not have an intermittent error. In this case, the error monitoring finishes when the specified error monitoring interval finishes and the number of errors is reset to 0. If an error occurs in the path later, the error monitoring resumes at the time the path is recovered from the error by using automatic failback.

If errors occur after a long interval, an intermittent error can be detected by increasing the error monitoring interval or by decreasing the number of times that the error is to occur (in order for the system to determine that an intermittent error is occurring).

Figure 2.11 shows the action for intermittent error monitoring when an intermittent error does not occur. In this example, the path is determined to have an intermittent error if the error occurs three or more times in 30 minutes. The events occurring in one path are described on the time arrow.

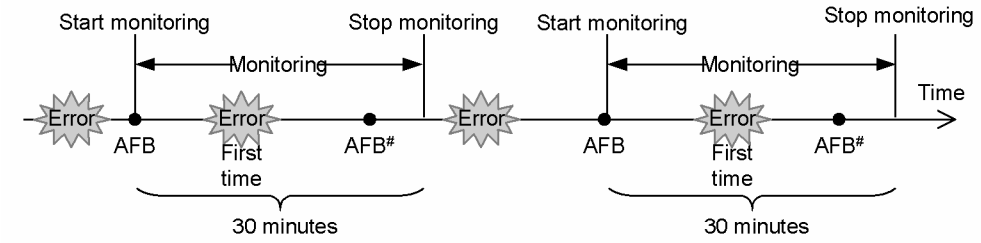

(Legend)
AFB: Indicates where the path was changed from error status to online status by AFB.

#
This includes online operation performed by a user.

Figure 2.11 Action When an Intermittent Error Does Not Occur in the Path

The number of times that the error is to occur is counted when an error occurs. As shown in Figure 2.11, normally the number of times that the error is to occur is counted each time an error occurs after the path is recovered online by automatic failback during intermittent error monitoring. However, if all the paths connected to the LU are in Offline(E), Online(E), or Offline(C) status due to disconnection of the paths or other reasons, the path is not recovered online by automatic failback. If I/O are continuously issued to such a LU, the number of times that the error is to occur might be counted even if the path is not placed online. If the number of times that the error is to occur reaches the specified value, the path is determined to have an intermittent error. In such a case, remove the cause of the error, and then manually place the path online.

2.9.3.3 When the Conditions for an Intermittent Error Are Changed During Error Monitoring

When the conditions (error monitoring interval and the number of times that the error is to occur) for an intermittent error are changed during error monitoring, the number of errors and the time that has passed since the error monitoring started are reset to 0. The error monitoring does not finish and resumes from the time the conditions are changed.

When you change the conditions outside the error monitoring time, at the time the path is recovered from the error by using automatic failback, the error monitoring starts with the changed conditions.

Figure 2.12 shows the action for intermittent error monitoring when the conditions for an intermittent error are changed during error monitoring. In this example, the conditions have been changed from 3 or more errors in 30 minutes, to 3 or more errors in 40 minutes. The events occurring in one path are described on the time arrow.

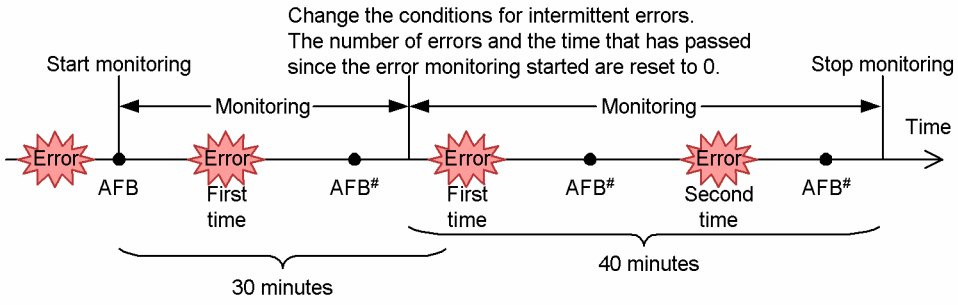

(Legend)
AFB: Indicates where the path was changed from error status to online status by automatic failback.

#
This includes online operations performed by a user.

Figure 2.12 Action When the Conditions for the Intermittent Error Are Changed During Error Monitoring

2.9.4 When User Operations Change the Intermittent Error Information

The following might be reset when the user changes the values set for an intermittent error or the path status: the number of errors that are counted during error monitoring, the time that has passed since error monitoring started, and the information about whether an intermittent error occurs (the path has been removed from those paths subject to automatic failback). Table 2.5 lists whether the above items are reset.

If you want to check whether intermittent error monitoring is being performed for the path, check the IEP item displayed when the `dlnkmgr` command's `view -path` operation is executed with the `-iem` parameter. If a numerical value of 0 or greater is displayed in the above items, then intermittent error monitoring is being performed.

Table 2.5 When User Operations Change the Intermittent Error Information

User operation		Number of errors and time passed since error monitoring started	Information about paths not subject to automatic failback
Changing the setting for intermittent error monitoring	Setting `off`	Reset	Reset#1
	Changing the conditions for an intermittent error during intermittent error monitoring	Reset#2	Inherit
	Setting `on` during intermittent error monitoring by executing the `set` operation (the conditions for the intermittent error monitoring are not changed)		
	Changing the conditions for an intermittent error outside the intermittent error monitoring.	(Not applicable) (Not counted.)	Inherit
Changing the automatic failback settings	Setting `off`	Reset	Reset
Changing the path status	Placing the path Offline(C)	Reset	Reset
	Placing the path Online outside the intermittent error monitoring	(Not applicable) (Not counted.)	Reset
	Placing the path Online during intermittent error monitoring	Inherit	(Not applicable)(If a path has been removed from the paths subject to automatic monitoring, that path is not monitored.)
Restarting the HDLM manager		Reset#3	Inherit
Restarting the host		Reset	Reset

#1

When you disable the intermittent error monitoring function, information about paths not subject to automatic failback will be reset. When you change the intermittent error monitoring functionality to off and you do not want to reset information about paths not subject to automatic failback, place the target paths Offline(C).

#2

The number of errors and the time since monitoring started is reset to `0`, and then monitoring restarts in accordance with the changed monitoring conditions.

#3

The number of errors and the time since monitoring started is reset to `0`, and then monitoring restarts at the time the HDLM manager starts.

2.10 Detecting Errors by Using Path Health Checking

HDLM can check the status of paths to which I/Os are not issued, at regular intervals, and detect errors. This function is called *path health checking*.

Without path health checking, an error is not detected unless I/O is issued because the system only checks the path status when I/O is issued. With path health checking, however, the system checks the status of online paths at regular intervals regardless of whether I/O is issued. If an error is detected in a path, path health checking function switches the status of that path to `Offline(E)` or `Online(E)`, so you can use the `dlnkmgr` command's `view` operation to check the path error.

For example, in a normal state, I/O is not issued on the paths of the standby host in the cluster configuration or on the non-owner paths (that is, some of the paths that access the Freedom Storage 5700E, Freedom Storage 5800, Thunder 9200, Thunder 9500V Series, and TagmaStore AMS/WMS series storage subsystem). Because of this, for the standby host or a host connected to non-owner paths, we recommend that you use path health checking to detect errors. This enables the system to use the most recent path-status information when selecting the switching destination.

You can configure path health checking by executing the `dlnkmgr` command's `set` operation. For details on the `set` operation, see 6.6.

2.11 Error Management

For troubleshooting, HDLM collects information into log files. The error information can be filtered according to the error level, and collected into the log files.

Figure 2.13 shows the data flow when collecting error information on a host where HDLM is running.

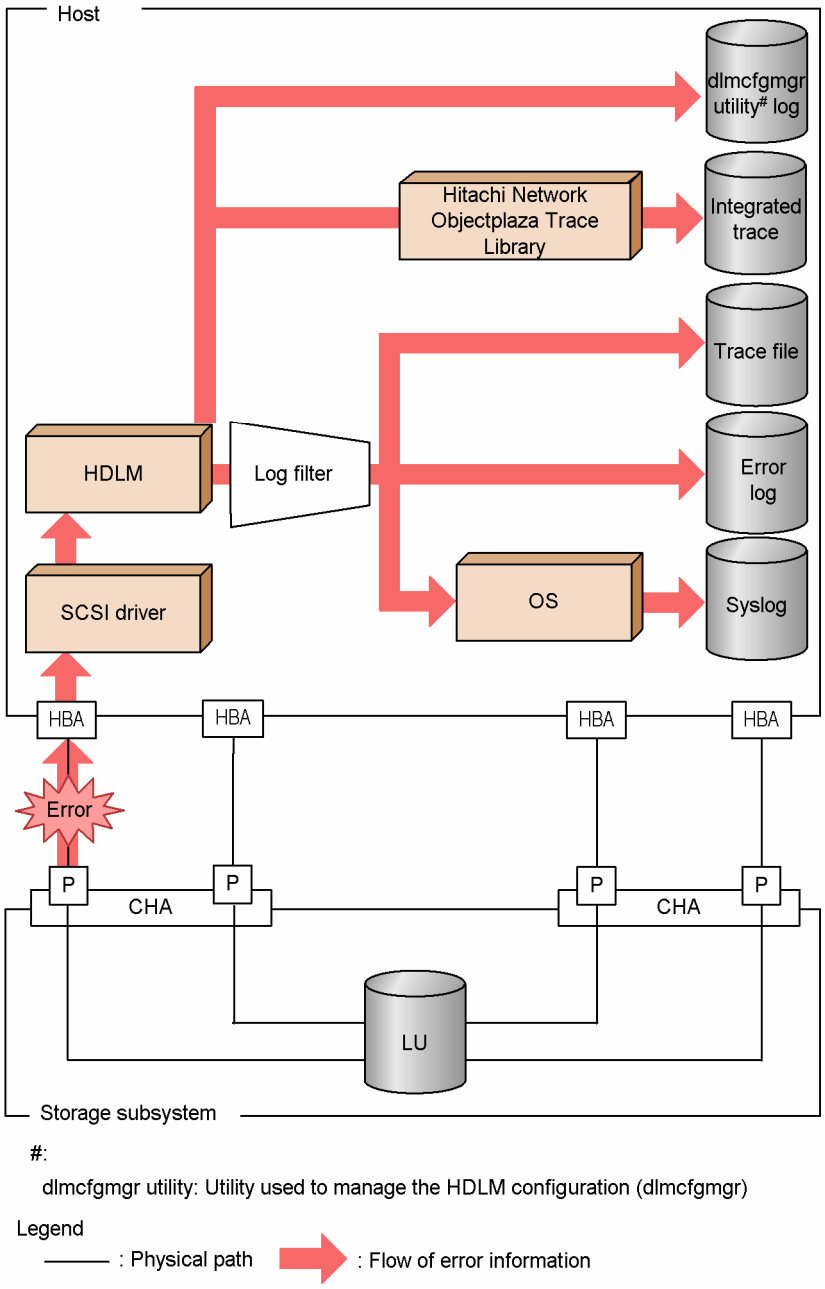

Figure 2.13 Data Flow When Collecting Error Information

Logs might be collected in layers lower than HDLM, such as for the SCSI driver. For details on the logs, refer to the Solaris documentation.

2.11.1 Types of Logs Collected

HDLM collects information on detected errors and trace information in the *integrated trace file*, *trace file*, *error logs*, *log for the dlmcfgmgr utility for managing the HDLM configuration*, and *syslog*. You can use the error information to examine the status of an error and analyze the cause of the error.

Table 2.6 describes the types of error information.

Table 2.6 Types of Error Information

Log name	Description	Output destination
Integrated trace file	Operation logs of the HDLM command is collected.	The default file path is `/var/opt/hitachi/HNTRLib2/spool/hntr2[1-16].log`. To specify the output destination directory and the file prefix for the integrated trace file, use a utility of Hitachi Network Objectplaza Trace Library (`HNTRLib2`).
Trace file	Trace information on the HDLM manager is collected for the level set by the user. If an error occurs, you might need to change the settings to collect trace information.	The trace file name is `/var/opt/DynamicLinkManager/log/hdlmtr[1-64].log`

Error log	Error information for the user-defined level is collected from detected errors. By default, HDLM collects all detected error information.	HDLM Manager logs:
		`/var/opt/DynamicLinkManager/log/dlmmgr[1-16].log`
		HDLM remote access interface log:
		`/var/opt/DynamicLinkManager/log/dlmwebagent[1-`*n*`].log`
		The value *n* depends on the setting in the file `dlmwebagent.properties`.

log for the dlmcfgmgr utility	Logs are collected when the `dlmcfgmgr` utility is executed.	The log file name is `/var/opt/DynamicLinkManager/log/dlmcfgmgr1[1-2].log`
Syslog	The HDLM messages on and above the level set by the user with `/etc/syslog.conf` are collected. # We recommend that you configure the system so that information at the Information level and higher is output. Syslogs can be checked using a text editor.	The default file path is `/var/adm/messages`. The syslog file path is specified in the file `/etc/syslog.conf`. For details, refer to the Solaris documentation.

#

When you want to configure the system so that HDLM messages are output to syslog, specify `user` for the facility in the `/etc/syslog.conf` file. The following shows an example where the system function name is `user`, and messages at the `info` level or higher are output to the `/tmp/syslog.user.log` file:

```
user.info          /tmp/syslog.user.log
```

For details on error levels, see 2.11.2.

2.11.2 Filtering of Error Information

Errors that HDLM detects are classified into error levels. Table 2.7 lists the error levels, in order of adverse effect on the system.

Table 2.7 Error Levels

Error level	Meaning	Level output in syslog
Critical	Fatal errors that may stop the system.	error
Error	Errors that crucially affect the system. This type of error can be avoided by using failover or other countermeasures.	error
Warning	Errors that enable the system to continue but, if left, might cause the system to operate improperly.	warning
Information	Information that indicates the operating history when the system operates normally.	info

Error information is filtered according to the error level, and then collected.

The error level is equivalent to the level of the messages output by HDLM. For details on the level of the message, see 8.1.1.

In syslog, the HDLM messages on and above the level set by the user configured in `/etc/syslog.conf` are collected. It is recommended that you set the information to be output at the info level and higher.

Note that the facility and level (in *facility* format) when HDLM outputs messages to syslog are all `user`.

The error information in error logs and trace files are collected based on a user-defined collection level. The collection levels are as follows:

Collection levels for error logs

- Collects no error information.
- Collects error information at the Error level and higher.
- Collects error information at the Warning level and higher.
- Collects error information at the Information level and higher.
- Collects error information at the Information level and higher (including maintenance information).

Collection levels for log information in trace files:

- Outputs no trace information
- Outputs error information only
- Outputs trace information on program operation summaries
- Outputs trace information on program operation details
- Outputs all trace information

For details on how to set the collection level, see 3.9.

2.11.3 Collecting Error Information Using the Utility for Collecting HDLM Error Information (DLMgetras)

HDLM provides a utility for collecting HDLM error information (`DLMgetras`).

By executing this utility, you can simultaneously collect all the information required for analyzing errors: information such as error logs, integrated trace files, trace files, definition files, core files, system crash dump files, and libraries You can use the collected information when you contact your HDLM vendor or maintenance company (if there is a maintenance contract for HDLM).

For details on the `DLMgetras` utility, see 7.2.

2.12 Cluster Support

HDLM is also available for cluster configurations.

For details on cluster software supported by HDLM, the supported Solaris version, and usable volume management software, see in Table 3.3 in section 3.1.2.4.

HDLM uses a path of the *active host* to access an LU.

Paths are switched in units of physical paths. Therefore, if an error occurs in a path, all the other paths that run through the same physical path are switched. The switching destination is a physical path of the active host. The details of host switching depend on the application.

Note

- When you use HDLM in a cluster configuration, you must install the same version of HDLM on all the nodes that comprise the cluster. If different versions of HDLM are installed, the cluster system may not operate correctly. If the `HDLM Version` and `Service Pack Version`, which are displayed by executing the following command, are the same, the versions of HDLM are the same:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -sys
```

- The DiskReservation agent of the VERITAS Cluster Server is not supported.
- I/O fencing of VCS 4.0 or later versions is not supported.

Chapter 3 Creating an HDLM Environment

This chapter describes the procedures for setting up an HDLM environment and the procedure for canceling those settings.

Make sure that HDLM is installed and the functions have been set up. Set up volume groups and cluster software to suit your operating environment.

- 3.1 System Requirements
- 3.2 Flow for Creating an HDLM Environment
- 3.3 HDLM installation types
- 3.4 Notes on Creating an HDLM Environment
- 3.5 Installing HDLM
- 3.6 Configuring a Boot Disk Environment
- 3.7 Migrating from a Boot Disk Environment to the Local Boot Disk Environment
- 3.8 Checking the Path Configuration
- 3.9 Setting Up HDLM Functions
- 3.10 Setting up Integrated traces
- 3.11 Creating File Systems for HDLM (when volume management software is not used)
- 3.12 Setting Up VxVM
- 3.13 Setting Up SDS
- 3.14 Setting Up SVM
- 3.15 Setting Up VCS
- 3.16 Uninstalling HDLM

3.1 System Requirements

Check the following before installing HDLM:

3.1.1 Applicable Hardware Models

3.1.1.1 Applicable Host

You can install HDLM on Sun SPARC series computer on which an OS listed in Table3.1 is running.

Table3.1 Applicable OSs for the host

OS	Required patches

Solaris 8	108434-04 or later, 108974-10 or later, and Recommended Patch Cluster Aug/27/02 or later
Solaris 9	Recommended Patch Cluster Nov/12/02 or later
Solaris 10	119685-07 or later is required. Also, other patches are required depending on the host bus adapters being used. For details on the other patches, see the HDLM *Release Notes*.

3.1.1.2 Applicable host bus adapters (HBAs)

For details on the applicable host bus adapters (HBAs), see HDLM *Release Notes*.

3.1.1.3 Applicable Storage subsystems

Applicable storage subsystems for use as data disks:

Freedom Storage 5700E#/Freedom Storage 5800#/Thunder 9200/Thunder 9500V Series#, Freedom Storage 7700E/Lightning 9900 Series/Lightning 9900V Series, TagmaStore AMS, TagmaStore USP, TagmaStore WMS, XP48/XP128/XP512/XP1024/XP10000/XP12000/XP24000, and Hitachi Universal Storage Platform V

#

Supports the Fibre Channel interface only.

The applicable storage subsystems require a dual controller configuration. If you use the system in a hub-connected environment, you must set unique loop IDs for all connected hosts and storage subsystems.

For details on the micro-programs and settings information for storage subsystems, which are required to use HDLM, see *Release Notes*.

Applicable storage subsystems for use as boot disks:

TagmaStore AMS/WMS series, TagmaStore USP, XP10000/XP12000/XP24000, and Hitachi Universal Storage Platform V

3.1.2 Related programs

3.1.2.1 When combining cluster configurations

Table 3.2 lists the cluster software versions supported by HDLM when building a cluster configuration.

Table 3.2 Supported cluster software versions

Cluster software	Supported versions	
	Local boot disk environment#1	Boot disk environment#6
VCS	VCS 1.3.0#2, 2.0#3, 3.5, 4.0#4, 4.1#4, or 5.0#4	VCS 4.1#4 or 5.0#4

Sun Cluster	Sun Cluster 3.0 (5/02 release)#5 or Sun Cluster 3.1#5	Sun Cluster 3.1 8/05 (Update 4)#5
Oracle RAC	Oracle9i RAC or Oracle RAC 10g	--

#1

An environment with a boot disk located on the host.

#2

Can be used when VERITAS Cluster Server 1.3.0 Patch 03 or later is applied.

#3

Can be used when VERITAS Cluster Server 2.0 Patch 02 or later is applied.

#4

Does not support the I/O fencing function of VERITAS Cluster Server 4.0 or later.

#5

Can be used when the prerequisite Sun Cluster patches are applied.

#6

An environment with a boot disk located in a storage subsystem instead of in the host.

3.1.2.2 When combining volume manager configurations

- When combining configurations by using VxVM
  - Local boot disk environment

    VERITAS Volume Manager 3.1.1#1, 3.2#2, #3, 3.5#2, 4.0#2, 4.1#2, or 5.0#2
  - Boot disk environment

    VERITAS Volume Manager 4.1#2 or 5.0#2
- When combining configurations by using SDS or SVM

  Solstice DiskSuite 4.2.1#4, or Solaris Volume Manager 1.0

HDLM-managed boot disks do not support a mirrored boot disk configuration incorporating a volume manager such as SDS, SVM, or VxVM. For this reason, you cannot register an HDLM-managed boot disk in `bootdg` when using VxVM.

#1

Can be used when VERITAS Volume Manager 3.1.1 Patch 04 or later is applied.

#2

When used with Thunder 9500V Series, TagmaStore AMS, TagmaStore WMS, Lightning 9900V Series, TagmaStore USP, XP512, XP48, Universal Storage Platform V, install the Array Support Library of VxVM before installing HDLM. For details on how to install the Array Support Library, see the storage subsystem documentation.

#3

Can be used when VERITAS Volume Manager 3.2 Patch 02 or later is applied.

#4

Can be used when Patch 108693-07 or later is applied for Solaris 8.

3.1.2.3 When you exchange data by using intermediate volumes managed by Hitachi RapidXchange

When you exchange data by using intermediate volumes managed by Hitachi RapidXchange, the following version of File Access Library and File Conversion Utility (FAL/FCU) is required:

- For Lightning 9900V Series

  01-03-56/20 or later

- TagmaStore USP

  01-04-64/21 or later

For details on Hitachi RapidXchange, see the manual *Reference Manual / File Conversion Utility & File Access Library*.

3.1.2.4 Combinations of related programs supported by HDLM

Table 3.3 lists the combinations of related programs, other than Oracle RAC, that are supported by HDLM.

Table 3.3 Combinations of related programs supported by HDLM

OS	Cluster	Volume manager
Solaris 8	None	None
		SDS 4.2.1
		VxVM 3.1.1
		VxVM 3.2
		VxVM 3.5
		VxVM 4.0
		VxVM 4.1
		VxVM 5.0
	Sun Cluster 3.0[#1]	None
		SDS 4.2.1
		VxVM 3.2
	Sun Cluster 3.1[#1]	None

OS	Cluster	Volume manager
		SDS 4.2.1
		VxVM 3.2
		VxVM 3.5
	Sun Cluster 3.1 (9/04)#1	None
		SDS 4.2.1
		VxVM 4.0
	Sun Cluster 3.1 (8/05)#1	None
		SDS 4.2.1
		VxVM 4.1
	VCS 1.3.0	None
		VxVM 3.1.1
	VCS 2.0	None
		VxVM 3.2
		VxVM 3.5
	VCS 3.5	None
		VxVM 3.2
		VxVM 3.5
	VCS 4.0#2	None
		VxVM 4.0
	VCS 5.0#2	None
		VxVM 5.0
Solaris 9	None	None
		SVM 1.0#3
		VxVM 3.5
		VxVM 4.0
		VxVM 4.1
		VxVM 5.0
	Sun Cluster 3.0#1	None
		SVM 1.0#3
	Sun Cluster 3.1#1	None
		SVM 1.0#3
		VxVM 3.5
	Sun Cluster 3.1 (9/04)#1	None
		SVM 1.0#3

OS	Cluster	Volume manager
		VxVM 4.0
	Sun Cluster 3.1 (8/05)#1	None
		SVM 1.0#3
		VxVM 4.1
	VCS 3.5	None
		VxVM 3.5
	VCS 4.0#2	None
		VxVM 4.0
	VCS 4.1#2	None
		VxVM 4.1
	VCS 5.0#2	None
		VxVM 5.0
Solaris 10	None	None
		SVM 1.0#3 #4
		VxVM 4.1
		VxVM 5.0
	Sun Cluster 3.1(8/05) #1	None
		SVM 1.0#3 #5
		VxVM 4.1
	VCS 4.1#2	None
		VxVM 4.1
	VCS 5.0#2	None
		VxVM 5.0

#1

In either of the following cases, the load balancing function is disabled because a reservation is issued to one of the LUs that is in use:

- When a failure occurs on one of the nodes in a two-node configuration running Sun Cluster and the LU cannot be accessed

- When the SDS 4.2.1 or SVM 1.0 shared diskset is being used in an environment without Sun Cluster

#2

Does not support the I/O fencing function for VCS 4.0 or later. Does not support linkage with SFVS (Storage Foundation Volume Server) for VCS 5.0.

#3

Does not support the following SVM functions:

- Multi-terabyte volumes

- Multi-owner disksets

- Diskset import

- Automatic (top down) volume creation

#4

In a configuration that uses a driver other than the Sun Microsystems HBA driver (other than the qlc or emlxs driver), the SVM shared diskset cannot use disks managed by HDLM.

#5

When the SVM shared diskset uses disks managed by HDLM in a configuration that uses a driver other than the Sun Microsystems HBA driver (other than the qlc or emlxs driver), use Sun Cluster device IDs (logical device files under /dev/did/dsk). The SVM shared diskset cannot use HDLM logical device names.

3.1.2.5 When you create an Oracle9i RAC environment

Required programs

Table 3.4 and Table 3.5 list the programs required to create an Oracle9i RAC environment.

Table 3.4 Programs required to create an Oracle9i RAC environment (for Solaris 8)

Program		Remarks
OS	Solaris 8	--
	Sun StorageTek SAN Foundation Software 4.2	Required package: ▪ SUNWsan
Cluster	Sun Cluster 3.0(5/02 release)	HDLM supports the two-node configuration only.
	Sun Cluster Support for Oracle Parallel Server/Real Application Clusters 3.0	Required packages: ▪ SUNWscucm ▪ SUNWudlm ▪ SUNWudlmr ▪ SUNWcvmr ▪ SUNWcvm
	Oracle9i 9.2.0.1	RAC is bundled with Oracle9i.
	Oracle UNIX Distributed Lock Manager 3.3.4.5	Required package: ▪ ORCLudlm
	Oracle9i Database Patch Set Release 9.2.0.3	--
Volume Manager	VERITAS Volume Manager 3.2	Patch 02 or later has to be applied.

Table 3.5 Programs required to create an Oracle9i RAC environment (for Solaris 9)

Program		Remarks
OS	Solaris 9	--
Cluster	Sun Cluster 3.1	HDLM supports the two-node configuration only.
	Sun Cluster Support for Oracle Parallel Server/Real Application Clusters 3.1	Required packages: ▪ SUNWscucm ▪ SUNWudlm ▪ SUNWudlmr ▪ SUNWcvmr ▪ SUNWcvm
	Oracle9i 9.2.0.4	RAC is bundled with Oracle9i. For Oracle9i versions prior to 9.2.0.4.0, Oracle9i Database Patch Set Release 9.2.0.4.0 must be applied.
	Oracle9i Database Patch Set Release 9.2.0.4.0	--
Volume Manager	VERITAS Volume Manager 3.5	--

Required patches (for Solaris 8)

Table 3.6 shows patches that are provided by Sun Microsystems, Inc. and are required to create an Oracle9i RAC environment.

Table 3.6 Patches required to create an Oracle9i RAC environment

Target program	Patch ID	Timing for applying
Sun Cluster	110648-28 or later	Apply the patch after installing Sun Cluster and before restarting the host.
	111554-09 or later	
	112108-07 or later	
VxVM	111095-13 or later	Apply the patch before installing VxVM.
	111096-06 or later	
	111097-12 or later	
	111412-11 or later	
	111413-10 or later	
	113201-03 or later	Apply the patch after installing VxVM and before restarting the host.
Oracle9i RAC	111857-07 or later	Apply the patch after installing Oracle9i.

3.1.2.6 When you create an Oracle RAC 10g environment

Required programs

Table 3.7 and Table 3.8 show programs required to create an Oracle RAC 10g environment.

Table 3.7 Programs required to create an Oracle RAC 10g environment (For Solaris 8 or 9)

Program		Remarks
OS	Solaris 8 or 9	When using Oracle 10g Database and ASM: ▪ In Solaris 8, use Update 7 or later. ▪ In Solaris 9, use Update 6 or later.
Cluster	Oracle 10g Database 10.1.0.2.0	--
	Oracle Cluster Ready Services (CRS) 10.1.0.2.0	--
	Oracle 10g Database 10.2.0.1.0	--
	Oracle Clusterware 10.2.0.1.0	--

Program		Remarks
Volume Manager	ASM	▪ ASM is bundled with Oracle RAC 10g. ▪ ASM is used for the disk memory area for files and recovery files of the Oracle database. In Oracle RAC 10g, HDLM devices can be used in the same way as procedures for disk devices. For details on how to install ASM, refer to the documentation for Oracle RAC 10g.

Table 3.8 Programs required to create an Oracle RAC 10g environment (For Solaris 10)

Program		Remarks
OS	Solaris 10	--
Cluster	Oracle 10g Database 10.1.0.4.0	--
	Oracle Cluster Ready Services (CRS) 10.1.0.4.0	--
	Oracle 10g Database 10.2.0.1.0	--
	Oracle Clusterware 10.2.0.1.0	--
	Oracle 10g Database 10.2.0.2.0	--
	Oracle Clusterware 10.2.0.2.0	--
Volume Manager	ASM	▪ ASM is bundled with Oracle RAC 10g. ▪ ASM is used for the disk memory area for files and recovery files of the Oracle database. In Oracle RAC 10g, HDLM devices can be used in the same way as procedures for disk devices. For details on how to install ASM, refer to the documentation for Oracle RAC 10g.

Required patches

Table 3.9, Table 3.10 show patches that are provided by Sun Microsystems Inc. and are required to create an Oracle RAC 10g environment. Table 3.11 show patches that are provided by Oracle Corporation and are required to create an Oracle RAC 10g environment.

Table 3.9 Patches required to create an Oracle RAC 10g environment (For Solaris 8)

Target program	Patch ID	Timing for applying
Oracle RAC 10g	108528-23 or later	Apply the patch before installing Oracle RAC 10g.
	108652-66 or later	
	108773-18 or later	
	108921-16 or later	
	108940-53 or later	
	108987-13 or later	
	108989-02 or later	
	108993-19 or later #	
	109147-24 or later	
	110386-03 or later	
	111023-02 or later	
	111111-03 or later	
	111308-03 or later	
	111310-01 or later	
	112396-02 or later	
	111721-04 or later	
	112003-03 or later	
	112138-01 or later	

#

When using Oracle RAC 10g 10.2.0.1.0, apply 108993-45, not 108993-19.

Table 3.10 Patches required to create an Oracle RAC 10g environment (For Solaris 9)

Target program	Patch ID	Timing for applying
Oracle RAC 10g	112233-11 or later	Apply the patch before installing Oracle RAC 10g.
	111722-04 or later	

Table 3.11 Patches required to create an Oracle RAC 10g environment (For Solaris 10)

Target program	Patch ID	Timing for applying
Oracle RAC 10g	P4332242#	Apply the patch after installing Oracle RAC 10g.

#

When using Oracle RAC 10g 10.2.0.1.0, this patch is not required.

Note

- When the Oracle RAC 10g Oracle Cluster Registry and the Oracle Cluster Ready Services voting disk are assigned to an HDLM device, the logical device file name of that HDLM device must be the same among the hosts. Therefore, when creating an environment, set the HDLM configuration after all the following conditions are satisfied:
  - HBA configuration conditions

    The HBA configuration (including HBAs not connecting to the SAN) of each host must match with respect to the following:

    - The HBA model

    - The slot where an HBA is attached, and the number of such slots

    - The HBA driver version

    - Whether or not FCode is installed in an HBA, and (if FCode is installed) the FCode version

    - The port locations of the HBA ports connected to the storage subsystem, and the number of such ports
  - Shared-disk configuration conditions

    The configuration of a shared disk with a Fibre Channel connection recognized by each host must match with respect to the following:

    - The path configuration (topology and number of paths)

    - Target IDs and LUNs at the time each host recognizes the disk
- When a host and an Oracle RAC 10g voting disk are connected by multiple paths, HDLM performs failover processing for those paths (in the same way as for normal paths) when an I/O timeout occurs for one of the paths.

  Note that, depending on the settings of Oracle RAC 10g, Oracle RAC 10g might determine that a node error has occurred before the failover processing performed by HDLM is completed, and then re-configure the cluster.

Therefore, when HDLM manages the paths that are connected to an Oracle RAC 10g voting disk, change the following settings according to your version of Oracle RAC 10g:

When using Oracle RAC 10g version 10.1.0.3.0 or later

Change the value of `MISSCOUNT` to the following value or greater: (*number-of-paths-connected-to-the-voting-disk* x 60 seconds).

When using Oracle RAC 10g version 10.2.0.2.0 or later

Change the value of `MISSCOUNT` to the following value or greater: (*number-of-paths-connected-to-the-voting-disk* x 60 seconds).

If four or more paths are connected to a voting disk, change `DISKTIMEOUT`, which is the I/O timeout threshold value for a voting disk, to the following value or greater: (*number-of-paths-connected-to-the-voting-disk* x 60 seconds).

For details on how to change `MISSCOUNT` and `DISKTIMEOUT`, contact the company with which you have a contract for Oracle Support Services.

Note that when you uninstall HDLM from the above configuration, you must reset the values of `MISSCOUNT` and `DISKTIMEOUT` to their original values. Therefore, make a note of the original values of `MISSCOUNT` and `DISKTIMEOUT` before changing them.

3.1.3 Memory and Disk Space Requirements

This subsection explains the memory and disk space requirements.

3.1.3.1 Memory requirements

The host requires 11 MB of memory.

3.1.3.2 Disk space requirements

Table 3.12 lists the disk space requirements for a host.

Table 3.12 Disk space requirements for a host

Directory	Disk space requirement
/etc	3MB
/kernel/drv	2MB + *x* bytes (*x* = *the-number-of-LUs* × *the-number-of-physical-paths* × 1.5KB)
/opt	264MB
/ `var`[#1]	*p*MB[#2] + *q*MB[#3] + 11MB

#1

100 MB of free space is required in `/var/tmp` during installation.

#2

This value depends on the error log file settings. The maximum is 30000MB.

$p = (s \times m) \div 1024$ (in MB),

where:

s is the size of the error log file (units: KB, default is 9,900), and

m is the number of error log files (default is 2). The resulting fraction will be rounded up.

#3

This value depends on the trace file settings. The maximum is 1000 MB.

$q = (t \times n) \div 1024$ (in MB),

where:

t is the size of the trace file (units: KB, default is 1,000), and

n is the number of trace files (default is 4). The resulting fraction will be rounded up.

3.1.4 The number of paths guaranteed by HDLM

Table 3.13 lists the number of paths per LU and the number of total physical paths that are guaranteed by HDLM.

Table 3.13 Maximum number of LUs and paths

Item	Maximum value
The number of LUs	256
The number of physical paths per LU	64
The number of total physical paths	4096

When Sun Cluster 3.0 or 3.1 is used, configure a server so that the sum of the number of physical paths per an LU for all nodes will be less than or equal to 15 as well as satisfying the conditions shown in Table 3.13.

3.2 Flow for Creating an HDLM Environment

Set up the environment to use HDLM as follows.

If you are using volume management software, you will not need to set up a file system.

- The installation procedure includes the Sun Cluster setup procedure.
- Execute the procedures described in *Preparations for an Upgrade Installation or Re-installation of HDLM* and *Performing an Upgrade Installation or Re-installation of HDLM* only when upgrading or re-installing HDLM.
- Execute the procedures described in *Creating File Systems for HDLM (when volume management software is not used)* and subsequent sections as required.

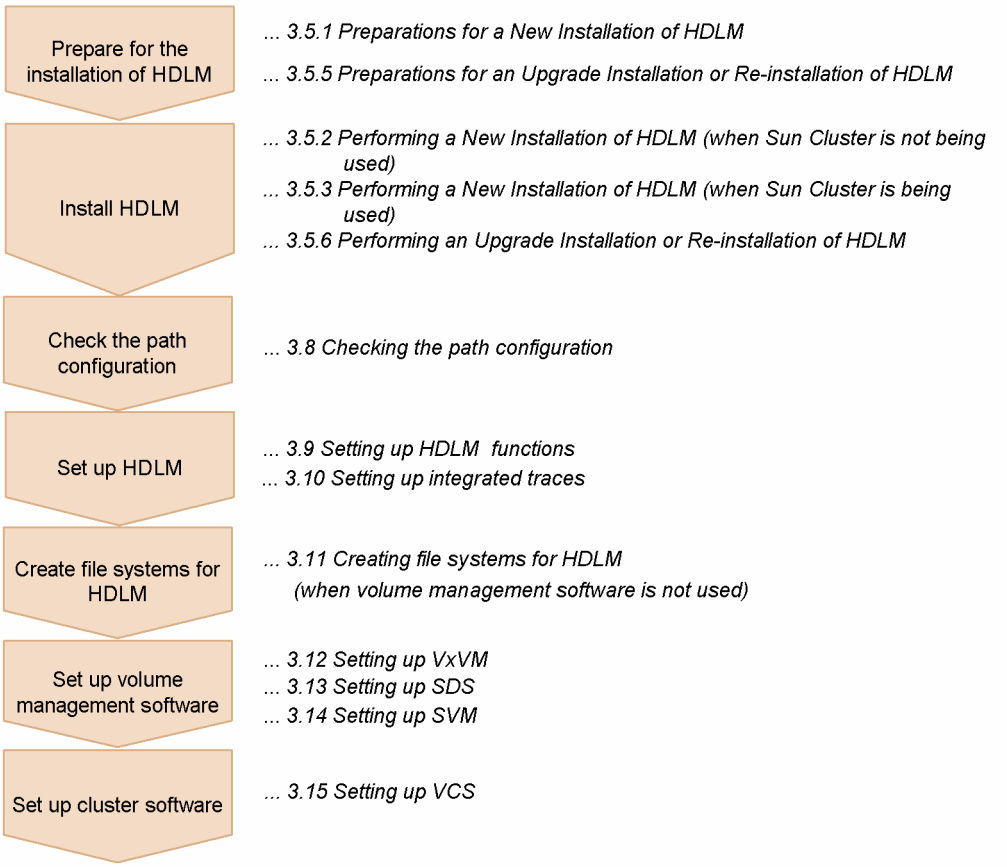

Figure 3.1 Flow of HDLM Environment Setup

3.3 HDLM installation types

This subsection describes the following types of HDLM installation: new installation, upgrade installation, and re-installation.

New installation of HDLM:

Installing HDLM in a server on which HDLM has not been installed is called a *new installation of HDLM*.

Upgrade installation of HDLM:

Installing a newer version than the existing version without uninstalling the existing version is called an *upgrade installation of HDLM*.

An HDLM 5.9 upgrade installation, can only be performed from HDLM 05-00 and later versions. If the current version is earlier than HDLM 05-00.

Re-installation of HDLM:

Installing the same version of HDLM, to restore the existing version, without first uninstalling that version is called a *re-installation of HDLM*.

3.4 Notes on Creating an HDLM Environment

This section gives points to be noted when creating an HDLM environment.

For points to be noted during operation of HDLM, see 4.1.

3.4.1 Notes on Hardware Settings

For details on the settings of the storage subsystem, see HDLM *ReleaseNotes*.

- When using HBAs that can set up a binding between target IDs and storage-subsystem ports, you must set up the binding before installing HDLM. For details, see the *Release Notes* and the HBA documentation.
- Do not change the vendor ID or product ID of a storage subsystem. If you do, HDLM will not be able to recognize the storage subsystem.
- The I/O timeout value reported to the sd or ssd driver depends on the timeout values for the switch and HBA.
- Depending on the switch and the HBA timeout value that is set, a timeout for I/O might take considerable time. Accordingly, switching a path or detecting a path error may take some time. Even the path health checking function may take longer than the checking interval set for detecting path errors.
- Make sure that all HBAs that are on a single host and connected to HDLM-managed devices are of the same type, and have the same HBA driver version, the same firmware, and the same FCode. If you using more than one type of HBA, paths will not be able to be switched when an error occurs.
- When you change the status of FCode on the HBA or FCode version, the physical path name of the sd or ssd device that HDLM manages might be changed. When you change the status of FCode after HDLM is installed, make sure that the physical path names of HDLM-managed devices have not changed. If the physical path name has changed, execute the `dlmsetconf` utility for creating the HDLM driver configuration definition file, and then restart the host with the reconfiguration option specified.

  In addition, when you use VxVM and remove the devices from VxVM control, make sure that the physical path names configured in the settings file (`/etc/vx/vxvm.exclude`) that is used to remove from VxVM control have not been changed. If a physical path name has been changed, re-set the physical path name. For details on the procedure for removing the devices from VxVM control, see 3.12.3.
- When using Fibre Channel, make sure that the target IDs of the SCSI comply with the HBA settings. To determine the target IDs, refer to the HBA documentation.

3.4.2 Notes on Installation

- You cannot edit the HDLM driver configuration definition file (`/kernel/drv/dlmfdrv.conf`). To change an HDLM-managed device to a non-HDLM-managed device, use the HDLM device unconfiguration function of the HDLM device. For details on the HDLM device unconfiguration function, see 3.5.4 or 4.4.5

- When you install HDLM, a device file dedicated to HDLM is created in the `/dev/dlm` directory. Do not use this device file.
- Install any other applications after completing the HDLM environment setup.
- If Hitachi Path Manager is installed on the host, uninstall it and restart the host before installing HDLM. For details on uninstalling Hitachi Path Manager, see the Hitachi Path Manager manual.
- If you are using HDLM 05-00 or earlier, uninstall the program while referring to 3.16.3, and then perform a new installation of HDLM 5.9.
- A license key is required to install HDLM.
- To update the HDLM license, execute the `dlnkmgr` command's `set -lic` operation. The license key specified in the license key file or the type of input license key determines the expiration of the license key. For details on license key types and the `set` operation,see 6.6.
- An upgrade installation of HDLM does not update the following files:
  - The HDLM driver configuration definition file (`/kernel/drv/dlmfdrv.conf`)
  - The HDLM nexus driver configuration definition file (`/kernel/drv/dlmndrv.conf`)
  - The HDLM function configuration file (`/opt/DynamicLinkManager/config/dlmmgr.xml`)
  - The HDLM driver unconfiguration definition file (`/opt/DynamicLinkManager/config/dlmfdrv.unconf`)

  The information in these files will be inherited when the update installation is complete.

  However, if both the conditions listed below are satisfied, the HDLM driver configuration definition file and the HDLM driver unconfiguration definition file are updated:

  - SVS is an HDLM management-target device.
  - You are upgrading from HDLM 5.6.1 or earlier.

  When both these conditions are satisfied, the previous generation of the HDLM driver configuration definition file (`/kernel/drv/dlmfdrv.conf.bak`) and HDLM driver unconfiguration definition file (`/opt/DynamicLinkManager/config/dlmfdrv.unconf.bak`) are also updated. For details about the `dlmfdrv.conf.bak` and `dlmfdrv.unconf.bak` files, see Note in section 7.5.2.

  If both the conditions above are satisfied, the following files are also created:

  - `/kernel/drv/dlmfdrv.conf_oldbackup`

    Backup of the `dlmfdrv.conf` file before the upgrade installation
  - `/kernel/drv/dlmfdrv.conf.bak_oldbackup`

    Backup of the `dlmfdrv.conf.bak` file before the upgrade installation
  - `/opt/DynamicLinkManager/config/dlmfdrv.unconf_oldbackup`

    Backup of the `dlmfdrv.unconf` file before the upgrade installation
  - `/opt/DynamicLinkManager/config/dlmfdrv.unconf.bak_oldbackup`

    Backup of the `dlmfdrv.unconf.bak` file before the upgrade installation

- If you install HDLM 5.9 after uninstalling HDLM 03-02, the format of the logical device file name of the HDLM device will change from `ch`*`W`*`t`*`X`*`d`*`Y`*`s`*`Z`* to `c`*`U`*`t`*`X`*`d`*`Y`*`s`*`Z`*. For details on the elements of `c`*`U`*`t`*`X`*`d`*`Y`*`s`*`Z`*, see in 2.6.

- In HDLM 5.6 or later, trace files for versions of HDLM earlier than 5.6 are divided into integrated trace files and trace files. The logs for the HDLM command (`dlnkmgr`)operation is output to integrated trace files. Trace information for an HDLM manager is output to trace files. The output destinations for the files are changed as follows:

  When you install HDLM 5.6 or later after uninstalling HDLM 04-00 or earlier

  Trace files before uninstallation:

  `/opt/hitachi/HNTRLib/spool/hntr`*`n`*`.log` (*n* indicates a file number)

  Integrated trace files after re-installation:

  `/var/opt/hitachi/HNTRLib2/spool/hntr2`*`n`*`.log` (*n* indicates a file number)

  Trace files after re-installation:

  `/var/opt/DynamicLinkManager/log/hdlmtr`*`n`*`.log` (*n* indicates a file number)

  When you re-install HDLM 5.6 or later after uninstalling HDLM 04-01, or when you perform an upgrade installation from HDLM 05-00 or later

  Trace files before uninstallation:

  `/var/opt/hitachi/HNTRLib2/spool/hntr2`*`n`*`.log` (*n* indicates a file number)

  Integrated trace files after upgrade installation or re-installation:
  `/var/opt/hitachi/HNTRLib2/spool/hntr2`*`n`*`.log` (*n* indicates a file number)

  Trace files after upgrade installation or re-installation:
  `/var/opt/DynamicLinkManager/log/hdlmtr`*`n`*`.log` (*n* indicates a file number)

- If you interrupted a new installation, check whether the `/var/opt/DynamicLinkManager` directory exists. If the `/var/opt/DynamicLinkManager` directory exists, execute the following command to delete it:

```
# rm -rf /var/opt/DynamicLinkManager
```

  If you interrupted an upgrade installation or re-installation, do not delete this directory.

- When you install a license by specifying the `set -lic` operation, you must execute one command at a time. If you execute multiple commands at the same time, a core file might be created and a message might be displayed.

  If a message is output, execute the `view -sys` operation to make sure that the license has been installed correctly.

- HDLM can be installed in a global zone only. If you install HDLM in a non-global zone, the following message will be output and the installation will be interrupted:

```
KAPL09140-E HDLM cannot be installed in this zone. Install it in the global zone.
```

- Paths can be managed in a global zone only. To use HDLM devices in a non-global zone, allocate HDLM devices from a global zone to the non-global zone for each LU or for each controller. The following shows an example of allocating the raw device of an HDLM device for `LU0` that is in a global zone to a non-global zone, and then setting the configuration of the non-global zone from the global zone.

  Example:

```
# zonecfg -z zone1
zonecfg:zone1> add device
```

  - To allocate the HDLM device for each LU, execute the following command:

```
zonecfg:zone1:device> set match=/dev/rdsk/c5t50060E800436E211d0*
```

  - To allocate the HDLM device for each controller, execute the following commands, and then restart the non-global zone:

```
zonecfg:zone1:device> set match=/dev/rdsk/c5t*
zonecfg:zone1:device> end
zonecfg:zone1> verify
zonecfg:zone1> commit
zonecfg:zone1> exit
#
```

  For details on the commands, refer to the Solaris 10 documentation. Table 3.14 shows the HDLM support for the Solaris zones.

Table 3.14 Support for the Solaris zones

Item	Global zone	Non-global zone	Remarks
Installing HDLM	Supported	Not supported	--
Path operations	Supported	Not supported	--
Displaying path statuses	Supported	Not supported	--
Accessing HDLM devices	Supported	Supported	The devices must be allocated to the relevant zone.
Uninstalling HDLM	Supported	Not supported	--

- When performing a new installation or upgrade installation of HDLM using the `pkgadd` command, respond as follows when you are prompted to enter information other than the license key:
  - When you want to continue the installation, respond to all prompts by entering `y`.
  - When you want to interrupt or stop the installation, respond by entering `q`.
- Before you perform the following operations in an environment whose run level is less than 2, execute the procedure described in *(1) Initializing HDLM when the host is started in single-user mode* in 4.1.9 *Operation in single-user mode*:

- Installing HDLM
- Configuring a boot disk environment
- Setting up HDLM functions
- Setting up an integrated trace#
- Creating a file system
- Executing an HDLM command
- Using an HDLM device
- Changing the configuration of the HDLM operating environment
- Applying an HDLM service pack

#

In single-user mode, HDLM command activity is not logged to integrated trace files.

- When you install HDLM while the 32-bit kernel mode is being used, only the 32-bit version of HDLM will be installed. When you install HDLM while the 64-bit kernel mode is being used, both the 32-bit and 64-bit versions of HDLM will be installed. Before installing HDLM, decide which kernel mode you will use to operate Solaris and, if needed, change to the desired kernel mode. If you switch the kernel mode from 32 bits to 64 bits when only the 32-bit version of HDLM is installed, you must reinstall HDLM after the switch is completed. For details on how to switch the kernel mode, see 3.5.1 in *Preparations for a New Installation of HDLM*.
- When installing HDLM on a host where a Device Manager Agent 5.0 or later is installed, do not execute any of the following commands of Device Manager Agent during installation:

```
hbsasrv, HiScan, hdvmagt_account, hdvmagt_schedule, hldutil, TIC
```

3.4.3 Notes on related software

- When using VxVM 3.1.1, VxVM 3.2, or VxVM 3.5 on the same host as HDLM, create rootdg on a device not managed by HDLM.
- When using VxVM 4.0 or later on the same host as HDLM, do not register any disk groups containing HDLM devices in `bootdg`.
- When configuring a boot disk environment on a host that is using VxVM 4.1, and then executing a VxVM command such as the `vxdisk list`, or `vxdmpadm listenclosure all` command, the enclosure type `FAKE_ARRAY` is displayed for the boot disk but there is no problem. For details on this phenomenon, see the Symantec documentation *VERITAS Storage Foundation(TM) 4.1 Release Notes Solaris Maintenance Pack 1*.
- If you use the `vxinstall` program or the `vxdiskadm` utility to perform an encapsulation while creating a disk group with an HDLM device or adding an HDLM device to a disk group, the KAPL08019-E or KAPL08022-E message might be output; however, this will not affect operations.
- Make sure that the version of HDLM is the same for all nodes on a cluster.
- HDLM and Sun StorageTek RAID Manager cannot be both present. Before installing HDLM, uninstall Sun StorageTek RAID Manager. If either of the following is applied, delete the `/devices/devices` link file by using the following command, and then install HDLM:

- When uninstalling Sun StorageTek RAID Manager
- When Sun StorageTek RAID Manager was previously installed and the `/devices/devices` link file remains

  ```
  # rm /devices/devices
  ```

For details on the `/devices/devices` link file, refer to the bug report for Bug ID 4673499 on the SUNSOLVE ONLINE website. You need a SUNSOLVE ONLINE account to view the bug report.

- Sun StorageTek Traffic Manager Software cannot be used for devices manageable by HDLM. If you need to use Sun StorageTek Traffic Manager Software for devices that cannot be managed by HDLM, use HBA ports that can differentiate between devices that can be managed by HDLM and devices that cannot be managed by HDLM separately.# In addition, for HBA ports that can differentiate between devices that can be managed by HDLM, the Sun StorageTek Traffic Manager Software function must be disabled.

  #

  You can use zoning of Fibre Channel switches or storage-subsystem's WWN security function to separately use HBA ports. You can also use point-to-point connections to physically separate the HBA ports.

3.4.4 Notes on command execution

When HDLM is used, you cannot specify the `-r` option in the `drvconfig`, `devlinks`, `disks`, or `devfsadm` command to reconfigure the HDLM devices on another boot disk. To reconfigure HDLM devices, you must boot them from the applicable boot disk.

3.4.5 Notes on the disk label

Make sure that you execute the `format` command to add the VTOC label to a disk to be managed by HDLM before changing that disk into an HDLM-managed disk. If you do not, the path status will not change to `Online`.

Note that you cannot change a disk that has the EFI (Extensible Firmware Interface) label into an HDLM-managed disk.

3.5 Installing HDLM

When you install HDLM, Hitachi Network Objectplaza Trace Library (HNTRLib2) will also be installed.

The file path of the integrated trace information file of Hitachi Network Objectplaza Trace Library (HNTRLib2) is `/var/opt/hitachi/HNTRLib2/spool/hntr2`*n*`.log` (*n* indicates a file number).

3.5.1 Preparations for a New Installation of HDLM

This section explains how to apply patches to Solaris, back up HDLM-managed devices, set up hardware, and set up volume management software.

3.5.1.1 Performing operations on devices to be managed by HDLM

If any devices to be managed by HDLM are operating, you need to perform the following procedure.

To prepare a device that is already operating:

1. Terminate the processes of all applications that access the device to be managed by HDLM.
2. If necessary, back up the contents of the device to tape.
3. Collect all relevant information about the logical device file names of the device, and internal LUNs of the LUs in the storage subsystem. This is so that you can keep the correspondence with the logical device file names once installation is complete.
4. When you execute the `dlmsetconf` utility for creating the HDLM driver configuration definition file, and restart the host with the reconfiguration option specified, the logical device files of the sd or ssd devices to be managed by HDLM will be deleted. Back up all pertaining to the logical device files under the `/dev/dsk` and `/dev/rdsk` directories to another directory.

   The following shows an example of executing the commands:

```
# mkdir /var/tmp/bkfile1
# ls -l /dev/dsk > /var/tmp/bkfile1/devices1
# ls -l /dev/rdsk > /var/tmp/bkfile1/devices2
```

5. Unregister the device.

   If the device is registered with any cluster software program applications (other than Volume Manager), unregister it. Since HDLM devices are accessed with the logical device file names assigned by HDLM, these applications will no longer be able to access such devices using the logical device file names currently used.

6. Unmount the device.

   If the device is mounted, unmount it. Also, if the device is set to be mounted automatically when the host starts, delete this setting from the `/etc/vfstab` file. The following is an example of this command.

   – Unmounting a device.

     First, check the current settings.

     ```
     # mount -p
     ```

     The current settings will be output as shown in Figure 3.2.

```
# mount -p
/dev/dsk/c0t8d0s0 - / ufs - no rw,intr,largefiles,onerror=panic,suid,dev=800038
/proc - /proc proc - no dev=4080000
fd - /dev/fd fd - no rw,suid,dev=4140000
mnttab - /etc/mnttab mntfs - no dev=4240000
swap - /var/run tmpfs - no dev=1
swap - /tmp tmpfs - no dev=2
/dev/dsk/c0t0d0s7 - /Sol7_export ufs - no rw,intr,largefiles,onerror=panic,suid,dev=800007
/dev/dsk/c0t0d0s0 - /Sol7_sys ufs - no rw,intr,largefiles,onerror=panic,suid,dev=800000
/dev/dsk/c0t8d0s7 - /export/home ufs - no rw,intr,largefiles,onerror=panic,suid,dev=80003f
/dev/dsk/c9t1d3s0 - /mntpt  ufs - no rw,intr,largefiles,onerror=panic,suid,dev=37c0000
#
```

Figure 3.2 Execution Results of the mount -p Command

The shaded portion shows the device to be managed by HDLM. Execute the following command to unmount it:

```
# umount /mntpt
```

– Clear the automatic mount setting.

An example of how to edit the `/etc/vfstab` file is shown in Figure 3.3.

```
#device         device          mount      FS   fsck mount   mount
#to mount       to fsck         point      type pass at boot options
#
#/dev/dsk/c1d0s2  /dev/rdsk/c1d0s2  /usr       ufs  1   yes     -
fd               -               /dev/fd    fd   -   no      -
/proc            -               /proc      proc -   no      -
/dev/dsk/c0t8d0s1 -              -          swap -   no      -
/dev/dsk/c0t8d0s0 /dev/rdsk/c0t8d0s0/          ufs  1   no      -
/dev/dsk/c0t8d0s7 /dev/rdsk/c0t8d0s7/export/home ufs  2   yes     -
swap             -               /tmp       tmpfs -  yes     -
/dev/dsk/c0t0d0s0 -              /Sol7_sys  ufs  -   yes     -
/dev/dsk/c0t0d0s7 -              /Sol7_export ufs  -   yes     -
#/dev/dsk/c9t1d3s0/dev/rdsk/c9t1d3s0/mntpt     ufs  2   yes     -
```

Figure 3.3 An example of how to edit the /etc/vfstab file

Comment the shaded portions out by placing a hash mark (#) at the beginning of the line, or delete the line.

7. Unregister the device from the volume management software.

   If the device is registered in VxVM, SDS, or SVM, cancel the registration.

3.5.1.2 Apply Solaris patches

See 3.1.1.1 Applicable Hardware Models, and then apply any necessary patches to the Solaris OS.

3.5.1.3 Set up the hardware

For a Fibre Channel connection, check the topology (Fabric, AL, etc.) and perform setup as appropriate.

To set up the hardware:

1. Set up the storage subsystem.

   For details on how to set up the storage subsystem, see the *HDLM Release Notes* that come with HDLM and the maintenance documentation for the storage subsystem.

   HDLM cannot manage command devices used by Hitachi CCI. To configure HDLM so that it does not manage a command device used by Hitachi CCI, on the Thunder 9200, enable the Standard Inquiry Data extended mode of the storage subsystem. For details on how to perform this setting, see the Storage Subsystem Maintenance documentation.

   Use the Hitachi CCI functionality to configure any command device used by Hitachi CCI so that it is in a redundant configuration.

2. Set up the switches.

   For details on how to set up a switch, refer to the switch documentation. This setup is not necessary if no switch is being used.

3. Set up the host bus adapter (HBA).

   For details on how to set up an HBA, see the *HDLM Release Notes*, and the HBA documentation.

3.5.1.4 Set up the /kernel/drv/sd.conf File

When using an sd driver, depending on the HBA driver, you need to edit the `/kernel/drv/sd.conf` file to recognize the disks. For details on how to edit this file, refer to the HBA documentation.

When using an ssd driver, you will not need to edit this file, since the connected disks will be automatically recognized.

After editing the `/kernel/drv/sd.conf` file, restart the host by executing either of the following commands with the reconfiguration option (`-r`) specified:

- `# reboot -- -r`
- `boot -r` (at the `ok` prompt)

After restarting the host, use the `format` command to confirm that the sd or ssd device that is to be managed by HDLM is displayed.

```
# format
```

For details on the settings, refer to the documentation for the HBA or Solaris.

3.5.1.5 Switch the kernel mode

When you install HDLM while the 32-bit kernel mode is being used, only the 32-bit version of HDLM will be installed. When you install HDLM while the 64-bit kernel mode is being used, both the 32-bit and 64-bit versions of HDLM will be installed. Before installing HDLM, decide which kernel mode you will use to operate Solaris, and if needed, change to the desired kernel mode. If you switch the kernel mode from 32 bits to 64 bits when only the 32-bit version of HDLM is installed, you must re-install HDLM after the switch is completed.

To switch the kernel mode currently being used:

1. Check the kernel mode currently being used.

   Execute the following command:

   ```
   # isainfo -b
   ```

   `32` will be displayed when a 32-bit kernel is being used, and `64` will be displayed when a 64-bit kernel is being used.

2. Switch the kernel mode.

   Execute the following command to switch to either the 32-bit or 64-bit kernel mode:

   - To switch to the 32-bit kernel mode:

     If Solaris is running, execute the following command:

     `# reboot -- "`*`boot-device`* `kernel/unix` *`other-boot-options`*`"`

     To boot at the `ok` prompt, execute the following command:

     `boot` *`boot-device`* `kernel/unix` *`other-boot-options`*

   - To switch to the 64-bit kernel mode:

     If Solaris is running, execute the following command:

     `# reboot -- "`*`boot-device`* `kernel/sparcv9/unix` *`other-boot-options`*`"`

     To boot at the `ok` prompt, execute the following command:

     `boot` *`boot-device`* `kernel/sparcv9/unix` *`other-boot-options`*

3.5.1.6 Set up the /etc/system File

HDLM uses a maximum of 8 KB of the stack for the kernel thread. Add the usage of HDLM to this maximum value, specify that total value as the stack size for the kernel thread, and then install HDLM. In Solaris 10, `default_stksize` has priority over `lwp_default_stksize`. Therefore, if `default_stksize` is specified in `/etc/system`, add 8 KB to the value specified in `default_stksize`. If the documentation for the volume manager or file systems that use HDLM devices specify that the stack size for the kernel thread is to be expanded, the value mentioned in the documentation must be also added to the stack size.

The procedure for increasing the stack size of the kernel thread from the default size in a system with a 64-bit kernel is as follows:

In Solaris 8 or 9

Add `set lwp_default_stksize=0x6000` to `/etc/system`, and then restart the host.

In Solaris 10

Add one of the following parameters to `/etc/system`, and then restart the host:

- `set lwp_default_stksize=0x6000`
- `set default_stksize=0x6000`

The default_stksize parameter takes precedence when both parameters are specified.

3.5.1.7 Set up the /etc/syslog.conf File

The dlmsetconf utility and the DLMgetras utility for collecting HDLM error information output messages to syslog. See the /etc/syslog.conf file (system log configuration definition file), and then check if it is set so that messages with a level of "Error" or higher are to be output.

Make sure that the /etc/syslog.conf file contains the following configuration entry. We recommend setting the output destination for *.err in the /etc/syslog.conf file to the default /var/adm/messages directory.

```
*.err;kern.debug;daemon.notice;mail.crit;    /var/adm/messages
```

3.5.1.8 Set up VxVM

To use HDLM and VxVM on the same host, set up VxVM as follows:

1. Make sure that the logical device file under the /dev/dsk and /dev/rdsk directories is backed up.

   If the information is not backed up, perform step 4 in 3.5.1.1 in 3.5.1.

2. If a created disk group already contains a device that is to be managed by HDLM, deport the disk group.

   For details, refer to the VxVM documentation. You can use the disk group again by importing the deported disk group after installing HDLM.

3. When used with Thunder 9500V Series, Lightning 9900V Series, TagmaStore USP, TagmaStore AMS/WMS, XP512, XP48, Universal Storage Platform V install the Array Support Library of VxVM before installing HDLM.

   For details, see the storage subsystem manual.

4. If you are using VxVM 3.1.1, VxVM 3.2, or VxVM 3.5, after installing VxVM, confirm that rootdg has been created on an sd or ssd device that will not be managed by HDLM.

   If any HDLM-managed sd or ssd devices have been added to rootdg, use a VxVM command to remove those devices from rootdg.

5. When you are using VxVM 4.0 or later, after installing VxVM, confirm that the disk group containing the sd or ssd device to be managed by HDLM has not been registered in bootdg.

   If the disk group containing the sd or ssd device to be managed by HDLM has been registered in bootdg, perform one of the following:

   - Use a VxVM command to remove the HDLM-managed sd or ssd device from the disk group.
   - Remove the target disk group from the bootdg registration

   For details, refer to the VxVM documentation.

6. Remove any sd or ssd disk to be managed by HDLM from VxVM control.

For details on how to remove any such sd or ssd device from VxVM control, refer to the VxVM documentation. If HDLM has been installed, remove any such sd or ssd device from VxVM control as directed in see 3.12.3.

3.5.1.9 Set up SDS and SVM

When SDS or SVM uses the sd or ssd device to be managed by HDLM, clear the device settings. For details, refer to the SDS or SVM documentation.

If you plan to use a SDS or SVM shared diskset, make sure that the hardware configurations of all the hosts are identical (host type; HBA type, number, slots; the number of paths and the HUBs and switches through which each path passes; and the storage ports and LUN used by each LU).

3.5.1.10 Set up Sun Cluster

If you are using Sun Cluster, before installing HDLM, perform the following operations to release the device group and Quorum device settings.

- Release the disk device groups that consist of the sd or ssd devices to be managed by HDLM.

  For details on how to unregister an SDS or SVM diskset, see 3.16.2.5 in 3.16.2 or 3.16.2.6 in 3.16.2. Unregistration of a diskset deletes the data in the diskset. Therefore, back up a diskset before unregistering it.

  To remove a VxVM disk group from Sun Cluster's disk device group, perform the following procedures:

  a) Place the target disk device group offline.

  ```
  # /usr/cluster/bin/scswitch -F -D disk-device-group-name
  ```

  b) Start the `scsetup` utility.

  ```
  # /usr/cluster/bin/scsetup
  ```

  c) Select `4` (Device group and volumes) to display `Device Group Menu`.

  d) To unregister a VxVM disk group, select `3` (Unregister a VxVM device group).

  Following the displayed directions, enter the name of the disk group you want to unregister.

- If an sd or ssd device to be managed by HDLM has been set as a Quorum device, perform the following procedures to remove the Quorum device settings.

  a) Use the following commands to check the device ID of the Quorum device:

  ```
  # /usr/cluster/bin/scconf -pv | grep Quorum
  # /usr/cluster/bin/scdidadm -L
  ```

b) If multiple Quorum devices exist, execute the following command until the number of Quorum devices becomes one. When the number of Quorum devices finally becomes one, go to step 3.

```
# /usr/cluster/bin/scconf -r -q globaldev=Quorum-device's-device-ID
```

c) Switch to become the superuser in the node that is to be changed to the maintenance mode in order to release the settings of the last Quorum device.

d) Back up all resource groups and sd or ssd devices from the node that is to be changed to maintenance mode.

The following shows an example of executing the command:

```
host1# /usr/cluster/bin/scswitch -S -h host1
```

e) Shut down the node to be changed to the maintenance mode so that the `ok` prompt appears.

The following shows an example of executing the command:

```
host1# shutdown -g0 -y -i0
```

f) At another node in the same cluster, switch to become the superuser, and then place the node that was stopped in step 5 into the maintenance mode.

The following shows an example of executing the command:

```
host2# /usr/cluster/bin/scconf -c -q node=host1,maintstate
```

g) Confirm that the cluster node is in the maintenance mode.

Make sure that the node that is in the maintenance node has the following status:

- The node is in the Offline status.

- Both the current `Quorum votes present` and `Quorum votes possible` are 0.

The following shows an example of executing the command:

```
host2# /usr/cluster/bin/scstat -q
```

h) If the cluster is a two -node configuration, enter the install mode.

If the cluster consists of three or more nodes, this step is not necessary so go to step 9.
The following shows an example of executing the command:

```
host2# /usr/cluster/bin/scconf -c -q installmode
```

i) Release the settings of the last Quorum device.

The following shows an example of executing the command:

```
host2# /usr/cluster/bin/scconf -r -q
globaldev=Quorum-device's-device-ID
```

j) At a node that is not in the maintenance mode, switch to become the superuser.

If the cluster consists of two nodes, go to step 12. If the cluster consists of three or more nodes, go to step 11.

k) If the cluster consists of three or more nodes, reset the cluster Quorum votes.

The following shows an example of executing the command:

```
host2# /usr/cluster/bin/scconf -c -q node=host1,reset
```

l) Change the mode of the node, which you changed into the maintenance mode, back to the previous mode.

At the `ok` prompt, execute `boot`.

m) If the cluster consists of two nodes, make sure that the node (which was changed into the maintenance mode) that you want to change back to the previous mode is active, and then reset the cluster Quorum votes. If the cluster consists of three or more nodes, this step is not necesssary.

The following shows an example of executing the command:

```
host2# /usr/cluster/bin/scconf -c -q node=host1,reset
```

3.5.2 Performing a New Installation of HDLM (when Sun Cluster is not being used)

To perform a new installation of HDLM:

1. When the license key file has been provided, Store the license key file directly under the `/var/tmp/` directory by using the `hdlm_license` name.

```
/var/tmp/hdlm_license
```

Note that the `hdlm_license` files is deleted after installation finishes successfully.

If the license key file does not exist, a message prompts you to enter a license key during installation. Enter the license key at this time.

2. Log in to Solaris as a user with root privileges.

3. Execute the `format` command to make sure that the sd or ssd device to be managed by HDLM is displayed.

```
# format
```

4. If you want to remove a device from HDLM management, execute the `format` command to check information for each path of that device.

The information you are checking here will be used in step 11.The following shows an example of executing the `format` command:

```
# format
Searching for disks...done

AVAILABLE DISK SELECTIONS:
           0. c1t0d0 <SUN36G cyl 24620 alt 2 hd 27 sec 107>
              /pci@8,600000/SUNW,qlc@4/fp@0,0/ssd@w21000004cf991b9b,0
           1. c1t1d0 <SUN36G cyl 24620 alt 2 hd 27 sec 107>
              /pci@8,600000/SUNW,qlc@4/fp@0,0/ssd@w21000004cf99152c,0
           2. c5t0d0 <HITACHI-OPEN-3-2103 cyl 3336 alt 2 hd 15 sec 96>
              /pci@8,700000/JNI,FCR@2/sd@0,0
           3. c5t1d0 <HITACHI-OPEN-3-2103 cyl 3336 alt 2 hd 15 sec 96>
              /pci@8,700000/JNI,FCR@2/sd@1,0
           4. c5t2d0 <HITACHI-OPEN-3-2103 cyl 3336 alt 2 hd 15 sec 96>
              /pci@8,700000/JNI,FCR@2/sd@2,0
Specify disk (enter its number):
```

In this example, the shaded portion shows information for each path of the device you are removing from HDLM management. Type in the appropriate device number, and then type `quit` to end the `format` command.

5. Insert the CD-ROM.

   After a few seconds, the CD-ROM will be automatically mounted, and the mount point of the CD-ROM will be displayed in a list of files.

   If the CD-ROM cannot be automatically mounted, execute the following command to mount the CD-ROM:

   ```
   # mount -F hsfs -o ro CD-ROM-device-name CD-ROM-mount-point
   ```

   The following shows an example of executing the commands for mounting the CD-ROM:

   ```
   # mkdir /cdrom/dlmanager
   # mount -F hsfs -o ro /vol/dev/dsk/c0t6d0/dlmanager /cdrom/dlmanager
   ```

6. Execute the following installation command:

   ```
   # pkgadd -d CD-ROM-mount-point DLManager
   ```

7. Continue by following the instructions displayed on the screen.

   If the license key file does not exist, a message prompts you to enter a license key. Enter the license key.

   - When a correct license key is entered, the installation program displays a message that indicates the license was installed.
   - Take care that if you enter an incorrect license key three times, installation stops.

   If you interrupt an installation, check whether the `/var/opt/DynamicLinkManager` directory exists.

   If the `/var/opt/DynamicLinkManager` directory exists, execute the following command to delete it:

   ```
   # rm -rf /var/opt/DynamicLinkManager
   ```

8. Execute the following command to check whether HDLM 5.9 has been installed successfully:

   Make sure that VERSION appears as 05-90.

```
# pkginfo -l DLManager
   PKGINST:  DLManager
      NAME:  Dynamic Link Manager
  CATEGORY:  system
      ARCH:  sparc
   VERSION:  xx-xx
   BASEDIR:  /
    VENDOR:
    ...
#
```

9. Make sure that the logical device file of the sd or ssd device is backed up.

   After you execute the dlmsetconf utility for creating the HDLM driver configuration definition file in the next step(s) (either step 10, or steps 10 and 11) and restart the host with the reconfiguration option specified, the logical device files of the sd or ssd devices to be managed by HDLM will be erased. Before you execute the dlmsetconf utility and restart the host, confirm that these files have been backed up. For details on how to back up files, see 3.5.1.1 in 3.5.1.

10. Execute the dlmsetconf utility to create the HDLM driver configuration definition file (/kernel/drv/dlmfdrv.conf) and the HDLM driver unconfiguration definition file (/opt/DynamicLinkManager/config/dlmfdrv.unconf).

```
# /opt/DynamicLinkManager/bin/ dlmsetconf
```

    Type y in response to the KAPL10242-I message displayed when you execute the dlmsetconf utility.

    Information pertaining to LUs to which the physical path accesses will be defined in the HDLM driver configuration definition file (/kernel/drv/dlmfdrv.conf). Do not edit this file.

11. Specify the devices that are to be excluded from management by HDLM, if there are any.

    This step is unnecessary if you want HDLM to manage all applicable devices. Devices that cannot be managed by HDLM are automatically removed from the HDLM driver configuration definition file. For details on devices that cannot be managed by HDLM, see 2.1.

    To remove devices from HDLM management, edit the HDLM driver unconfiguration definition file (/opt/DynamicLinkManager/config/dlmfdrv.unconf).

    For details on how to perform this setting, see 3.5.4.

    Note

    Set the boot disk in the post-migration environment as an HDLM-managed device if, before this step, you executed the procedure described in *Migration by installing HDLM in the existing local boot disk environment Procedure for Configuring a Boot Disk Environment*.

Note the following if, before this step, you executed the procedure described in *Migration by installing HDLM in the existing boot disk environment* or *Migration by building a new pre-migration environment* in *Procedure for Configuring a Boot Disk Environment*:

- If the boot disk in the pre-migration environment is an HDLM-managed device, change the settings to exclude it from HDLM management.

- Set the post-migration boot disk as an HDLM-managed device.

12. When you use VxVM, from VxVM, disable the devices that you do not need to access.

    Specify settings as shown in 3.12.3.2 or 3.12.3.3 in 3.12.3.

13. Restart the host with the reconfiguration option specified to make the settings valid and the HDLM device recognizable.

    Execute one of the following commands:

    - `# reboot -- -r`
    - `boot -r` (at the `ok` prompt)

14. Execute the `dlnkmgr` command's `view` operation to display the status of the programs.

    The following shows an example of executing the commands:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -sys
HDLM Version                : xx-xx
Service Pack Version        :
Load Balance                : on(rr)
Support Cluster             :
Elog Level                  : 3
Elog File Size (KB)         : 9900
Number Of Elog Files        : 2
Trace Level                 : 0
Trace File Size(KB)         : 1000
Number Of Trace Files       : 4
Path Health Checking        : on(30)
Auto Failback               : off
Intermittent Error Monitor  : off
HDLM Manager Ver      WakeupTime
Alive        xx-xx    yyyy/mm/dd hh:mm:ss
HDLM Alert Driver Ver      WakeupTime          ElogMem Size
Alive             xx-xx    yyyy/mm/dd hh:mm:ss
HDLM Driver Ver      WakeupTime
Alive       xx-xx    yyyy/mm/dd hh:mm:ss
License Type Expiration
Permanent    -
KAPL01001-I The HDLM command completed normally. Operation name = view, completion time =
yyyy/mm/dd hh:mm:ss
#
```

    Nothing is displayed for `Support Cluster` even if you are using cluster software. However, the cluster-support functions are, in fact, operating without problem.

15. By checking to the execution results of the `view` operation, verify the version of the installed HDLM.

    If `HDLM Version` is `05-90`, the installed version of HDLM is correct.

16. By checking to the execution results of the `view` operation, verify that the programs are running normally.

If HDLM Manager, HDLM Alert Driver, and HDLM Driver are all Alive, the programs are running normally.

17. In the environment settings files (such as .profile) for root privileges, add /opt/DynamicLinkManager/bin to the PATH environment variable.

    When the Bourne shell or Korn shell is used, add the path as shown below:

    ```
    PATH=$PATH:/opt/DynamicLinkManager/bin ; export PATH
    ```

    When the C shell is used, add the path as shown below:

    ```
    set path=($path /opt/DynamicLinkManager/bin)
    ```

18. Execute the following command to view the HDevName and iLU columns:

    ```
    # /opt/DynamicLinkManager/bin/dlnkmgr view -path
    ```

    Check the execution results of this command to determine the correspondence between each HDevName (the logical device file name of the HDLM device) and iLU (the internal LU number). Then, compare the correspondence of the logical device file names and internal LUs that took place before HDLM was installed, and then re-register the correct HDLM devices in those programs that will use these devices.

19. If you uninstalled HDLM 04-00 or HDLM 04-01, and then installed HDLM 5.9, create a table that shows the correspondence between the HDLM logical device file names before and after installation of HDLM 5.9.

    Specify the file (the file you created before you uninstalled HDLM 04-00 or HDLM 04-01) in which you saved the output results of the dlnkmgr command's view -path operation, as the input file for the dlmlisthdev utility for assisting HDLM transitions. The following shows an example of executing the dlmlisthdev utility when upgrading from HDLM 04-00 to HDLM 5.9.

```
# /opt/DynamicLinkManager/bin/dlmlisthdev  /var/tmp/old_view_path_list.txt
OldHDevName NewHDevName DskName                              iLU
c3t0d10     c5t0d10     HITACHI .DF500F      .0207     0010
c3t0d11     -           HITACHI .DF500F      .0207     0011
-           c5t0d12     HITACHI .DF500F      .0207     0012
...
#
```

Table 3.15 describes the items shown above.

Table 3.15 Information displayed by the dlmlisthdev utility for assisting HDLM transitions

Item	Description
OldHDevName	Indicates the HDLM 04-00 logical device file name (HDevName). If a hyphen (–) is displayed, the corresponding HDev did not exist before the HDLM 04-00 installation.
NewHDevName	Indicates the HDLM 5.9 logical device file name (HDevName). If a hyphen (–) is displayed, the corresponding HDev does not exist after the HDLM 5.9 installation.
DskName	Indicates the storage subsystem name.
iLU	Indicates the LU number in the storage subsystem.

20. Based on the results of step 19, compare the logical device file names from before and after the HDLM 5.9 installation, and re-register the correct HDLM devices in programs that use HDLM devices.

Once installation is complete, check the path information according to the procedures listed in 3.8 to check whether HDLM has been installed correctly or not.

3.5.3 Performing a New Installation of HDLM (when Sun Cluster is being used)

Note

Do not restart any nodes until you have finished installing HDLM on all nodes.

To perform a new installation of HDLM:

1. When the license key file has been provided, Store the license key file directly under the `/var/tmp/` directory by using the `hdlm_license` name.

```
/var/tmp/hdlm_license
```

   Note that the `hdlm_license` files is deleted after installation `finishes successfully`.

   If the license key file does not exist, a message prompts you to enter a license key during installation. Enter the license key at this time.

2. Log in to Solaris as a user with root privileges.
3. Execute the `format` command to make sure that the sd or ssd device to be managed by HDLM is displayed.

```
# format
```

4. If you want to remove a device from HDLM management, execute the `format` command to check information for each path of that device.

   The information you are checking here will be used in step 17.The following shows an example of executing the `format` command:

```
# format
Searching for disks...done

AVAILABLE DISK SELECTIONS:
           0. c1t0d0 <SUN36G cyl 24620 alt 2 hd 27 sec 107>
              /pci@8,600000/SUNW,qlc@4/fp@0,0/ssd@w21000004cf991b9b,0
           1. c1t1d0 <SUN36G cyl 24620 alt 2 hd 27 sec 107>
              /pci@8,600000/SUNW,qlc@4/fp@0,0/ssd@w21000004cf99152c,0
           2. c5t0d0 <HITACHI-OPEN-3-2103 cyl 3336 alt 2 hd 15 sec 96>
              /pci@8,700000/JNI,FCR@2/sd@0,0
           3. c5t1d0 <HITACHI-OPEN-3-2103 cyl 3336 alt 2 hd 15 sec 96>
              /pci@8,700000/JNI,FCR@2/sd@1,0
           4. c5t2d0 <HITACHI-OPEN-3-2103 cyl 3336 alt 2 hd 15 sec 96>
              /pci@8,700000/JNI,FCR@2/sd@2,0
Specify disk (enter its number):
```

In this example, the shaded portion shows information for each path of the device you are removing from HDLM management. Type in the appropriate device number, and then type `quit` to end the `format` command.

5. If the settings of a Quorum device or disk device group are specified for the device to be managed by HDLM, remove those settings.

   For details on how to remove the Quorum device and the VxVM disk group settings, see 3.5.1.10 in 3.5.1.

   For details on how to unregister an SDS or SVM diskset, see 3.16.2.5 or 3.16.2.6 in 3.16.2. Unregistering a diskset deletes the data in the diskset. Therefore, back up a diskset before unregistering it.

6. Make sure that the key to be used for disk reservation has been registered with the device to be managed by HDLM.

   On any one node, execute the following command:

```
# /usr/cluster/lib/sc/scsi -c inkeys -d
slice-2-logical-device-file-name-for-the-device-to-be-HDLM-managed
```

   The following shows an example of executing the command:

```
# /usr/cluster/lib/sc/scsi -c inkeys -d /dev/rdsk/c5t0d0s2
Reservation keys(3):
0x40235c5200000002
0x40235c5200000001
0x40235c5200000003
#
```

   In this example, keys are displayed.

   If keys are displayed, perform the steps beginning with step 7. Even if the Quorum device or disk device group settings are removed, you need to restart the node in non-cluster mode and cancel the key registration because the keys to be used for disk reservation are left registered.

   If no key is displayed, perform the steps beginning with step 11.

7. Stop all nodes that comprise the cluster. On any one node, execute the following command:

```
# /usr/cluster/bin/scshutdown -g0 -y
```

8. Start all nodes comprising the cluster in non-cluster mode.

   Execute `boot -x` at the `ok` prompt.

9. Cancel the key registration for the device to be managed by HDLM.

   On any one node, execute the following command:

```
# /usr/cluster/lib/sc/scsi -c scrub -d
slice-2-logical-device-file-name-for-the-device-to-be-HDLM-managed
```

   The following shows an example of executing the command:

```
# /usr/cluster/lib/sc/scsi -c scrub -d /dev/rdsk/c5t0d0s2
Reservation keys currently on disk:
0x40235c5200000002
0x40235c5200000001
0x40235c5200000003
Attempting to remove all keys from the disk...
Scrubbing complete, use '/usr/cluster/lib/sc/scsi -c inkeys -d /dev/rdsk/c5t0d0s2' to
verify success
#
```

10. Make sure that the key registration was canceled for the device to be managed by HDLM.

    On any one node, perform step 6 again and confirm that no key is displayed.

    The following shows an example of executing the command:

```
# /usr/cluster/lib/sc/scsi -c inkeys -d /dev/rdsk/c5t0d0s2
Reservation keys(0):
#
```

    In this example, key registration is canceled.

11. Insert the CD-ROM and install HDLM on all nodes that comprise the cluster.

    After a few seconds, the CD-ROM will be automatically mounted, and the mount point of the CD-ROM will be displayed in a list of files.

    If the CD-ROM cannot be automatically mounted, execute the following command to mount the CD-ROM:

```
# mount -F hsfs -o ro CD-ROM-device-name CD-ROM-mount-point
```

    The following shows an example of executing the command for mounting the CD-ROM:

```
# mkdir /cdrom/dlmanager
# mount -F hsfs -o ro /vol/dev/dsk/c0t6d0/dlmanager /cdrom/dlmanager
```

12. Execute the following installation command:

```
# pkgadd -d CD-ROM-mount-point DLManager
```

13. Continue by following the instructions displayed on the screen:

    If the license key file does not exist, a message prompts you to enter a license key. Enter the license key.

- When a correct license key is entered, the installation program displays a message that indicates the license was installed.
- Take care that if you enter an incorrect license key three times, installation stops.

If you interrupt an installation, check whether the /var/opt/DynamicLinkManager directory exists. If the /var/opt/DynamicLinkManager directory exists, execute the following command to delete it:

```
# rm -rf /var/opt/DynamicLinkManager
```

14. Execute the following command to check whether that HDLM 5.9 has been installed successfully.

    Make sure that VERSION appears as 05-90.

```
# pkginfo -l DLManager
   PKGINST:  DLManager
      NAME:  Dynamic Link Manager
  CATEGORY:  system
      ARCH:  sparc
   VERSION:  xx-xx
   BASEDIR:  /
    VENDOR:
    ...
#
```

15. Make sure that the logical device file of the sd or ssd device is backed up

    After you execute the dlmsetconf utility for creating the HDLM driver configuration definition file in the next step(s) (either step 16, or steps 16 and 17) and restart the node with the reconfiguration option specified, the logical device files for the sd or ssd devices to be managed by HDLM will be deleted. Before you execute the dlmsetconf utility and restart the node, confirm that these files have been backed up. For details on how to back up files, see 3.5.1.1 in 3.5.1.

16. Execute the following dlmsetconf utility to create the HDLM driver configuration definition file (kernel/drv/dlmfdrv.conf) and the HDLM driver unconfiguration definition file (/opt/DynamicLinkManager/config/dlmfdrv.unconf).

```
# /opt/DynamicLinkManager/bin/dlmsetconf
```

    Type y in response to the KAPL10242-I message displayed when you execute the dlmsetconf utility.

    Information associated with LUs that access a physical path is defined in the HDLM driver configuration definition file (/kernel/drv/dlmfdrv.conf). Do not edit this file.

17. Specify the devices that are to be excluded from management by HDLM, if there are any.

    This step is unnecessary if you want HDLM to manage all applicable devices. Devices that cannot be managed by HDLM are automatically removed from the HDLM driver configuration definition file. For details on devices that cannot be managed by HDLM, see 2.1.

    To remove devices from HDLM management, edit the HDLM driver unconfiguration definition file (/opt/DynamicLinkManager/config/dlmfdrv.unconf).

    For details on how to perform this setting, see 3.5.4.

    Note

Set the boot disk in the post-migration environment as an HDLM-managed device if, before this step, you executed the procedure described in *Migration by installing HDLM in the existing local boot disk environment* in *Procedure for Configuring a Boot Disk Environment*.

Note the following if, before this step, you executed the procedure described in *Migration by installing HDLM in the existing boot disk environment* or *Migration by building a new pre-migration environment* in *Procedure for Configuring a Boot Disk Environment*:

- If the boot disk in the pre-migration environment is an HDLM-managed device, change the setting to exclude it from HDLM management.

- Set the post-migration boot disk as an HDLM-managed device.

18. When you use VxVM, from VxVM disable the devices that you do not need to access.

    Specify settings as shown in 3.12.3.2 or 3.12.3.3 in 3.12.3.

19. Stop all nodes that comprise the cluster.

    If one or more keys were displayed at step 6 and the node was started in non-cluster mode at step 8, execute the following command on all nodes:

```
# shutdown -g0 -y -i0
```

    If the nodes were started in cluster mode, execute the following command on any one node:

```
# /usr/cluster/bin/scshutdown -g0 -y
```

20. With the reconfiguration option specified, start all nodes comprising the cluster in cluster mode, at the same time.

    Execute `boot -r` at the `ok` prompt.

    Note

    While a node is starting, the message below may be output to the console or syslog:

```
Could not read symbolic link for: /dev/rdsk/c5t0d0s2 path not loaded
```

    This message is output because HDLM deleted the logical device file of an sd or ssd device managed by HDLM. Perform the procedure for updating the device ID information (steps 26 to 29) so that this message is not output the next time the node starts. For details, refer to the Sun Cluster documentation.

21. Execute the `dlnkmgr` command's `view` operation on all nodes comprising the cluster to display the status of the programs.

    The following shows an example of executing the command:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -sys
HDLM Version               : xx-xx
Service Pack Version       :
Load Balance               : on(rr)
Support Cluster            :
Elog Level                 : 3
Elog File Size (KB)        : 9900
Number Of Elog Files       : 2
```

```
Trace Level                    : 0
Trace File Size(KB)            : 1000
Number Of Trace Files          : 4
Path Health Checking           : on(30)
Auto Failback                  : off
Intermittent Error Monitor     : off
HDLM Manager Ver      WakeupTime
Alive        xx-xx    yyyy/mm/dd hh:mm:ss
HDLM Alert Driver Ver      WakeupTime          ElogMem Size
Alive             xx-xx    yyyy/mm/dd hh:mm:ss
HDLM Driver Ver      WakeupTime
Alive       xx-xx    yyyy/mm/dd hh:mm:ss
License Type Expiration
Permanent    -
KAPL01001-I The HDLM command completed normally. Operation name = view, completion time =
yyyy/mm/dd hh:mm:ss
#
```

Nothing is displayed for `Support Cluster` even if you are using cluster software. However, the cluster-support functions are, in fact, operating without problem.

22. By checking the execution results of the `view` operation, verify the version of the installed HDLM.

    If `HDLM Version` is `05-90`, the installed version of HDLM is correct.

23. By checking the execution results of the `view` operation, verify that the programs are running normally.

    If `HDLM Manager`, `HDLM Alert Driver`, and `HDLM Driver` are all `Alive`, the programs are running normally.

24. If you uninstalled HDLM 04-00 or HDLM 04-01 and then installed HDLM 5. 9, create a table that shows the correspondence between the HDLM logical device file names before and after installation of HDLM 5.9.

    Specify the file (the file you created before you uninstalled HDLM 04-00 or HDLM 04-01) in which you saved the output results of the `dlnkmgr` command's `view -path` operation, as the input file for the `dlmlisthdev` utility for assisting HDLM transitions. The following shows an example of executing the `dlmlisthdev` utility when upgrading from HDLM 04-00 to HDLM 5.9:

```
# /opt/DynamicLinkManager/bin/dlmlisthdev  /var/tmp/old_view_path_list.txt
OldHDevName NewHDevName DskName                                  iLU
c3t0d10     c5t0d10     HITACHI .DF500F      .0207     0010
c3t0d11     -           HITACHI .DF500F      .0207     0011
-           c5t0d12     HITACHI .DF500F      .0207     0012
...
#
```

Table 3.16 describes the items shown above.

Table 3.16 Information displayed by the dlmlisthdev utility for assisting HDLM transitions

Item	Description
`OldHDevName`	Indicates the HDLM 04-00 logical device file name (`HDevName`). If a hyphen (-) is displayed, the corresponding `HDev` did not exist before the HDLM 5.9 installation.

NewHDevName	Indicates the logical device file name (HDevName) after the HDLM 5.9 upgrade installation. If a hyphen (–) is displayed, the corresponding HDev does not exist after the HDLM 5.9 installation.
DskName	Indicates the storage subsystem name.
iLU	Indicates the LU number in the storage subsystem.

25. Based on the results of step 24, compare the logical device file names from before and after the HDLM 5.9 installation, and re-register the correct HDLM devices in programs that use HDLM devices.

    After you have checked the HDLM operation, repeat steps 26 to 29 below on all nodes.

26. Remove the device ID information of devices that do not exist. Execute the following command:

```
# /usr/cluster/bin/scdidadm -C
```

27. Assign device IDs. Execute the following command:

```
# /usr/cluster/bin/scdidadm -r
```

28. Notify the server of the new device ID information. Execute the following command:

```
# /usr/cluster/bin/scdidadm -ui
```

29. Update the Sun Cluster global device namespace. Execute the following command:

```
# /usr/cluster/bin/scgdevs
```

30. To use an HDLM device for a Quorum device, specify the Quorum device settings at any of the nodes that constitute the cluster.

    The following shows examples of executing the commands:

    – Check the device ID of the HDLM device.

      ```
      # /opt/DynamicLinkManager/bin/dlnkmgr view -path
      # /usr/cluster/bin/scdidadm -L
      ```

    – Set up the Quorum device.

      ```
      # /usr/cluster/bin/scconf -a -q globaldev=d4
      ```

    – In a two-node configuration, execute the following command to release the cluster installation mode:

      ```
      # /usr/cluster/bin/scconf -c -q reset
      ```

31. On all nodes in the cluster, add /opt/DynamicLinkManager/bin to the PATH environment variable in the environment settings files (such as .profile) for root privileges.

    When the Bourne or Korn shells is used, add the path as shown below:

    ```
    PATH=$PATH:/opt/DynamicLinkManager/bin ; export PATH
    ```

When the C shell is used, add the path as shown below:

```
set path=( $path /opt/DynamicLinkManager/bin )
```

32. Execute the following command to view the HDevName and iLU columns:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -path
```

Check the execution results of this command to determine the correspondence between each HDevName (logical device file name of the HDLM device) and iLU (the internal LU number). Then, compare the correspondence of the logical device file names and internal LUs that took place before HDLM was installed, and then re-register the correct HDLM devices in those programs that will use these devices.

To use an HDLM-managed device with a file system created by using volume management software, perform step 33.

To use an HDLM-managed device with a file system created without using volume management software, perform steps 34 to 36.

33. Register the disk device group.

If you are using VxVM, see 3.12.5.

If you are using SDS, see 3.13.2.2 in 3.13.2.

If you are using SVM, see 3.14.2.2 in 3.14.2.

34. In Sun Cluster, to create a new shared file system, create it in any HDLM-managed device from a desired node.

The following shows an example of executing the command:

```
host1# newfs /dev/global/rdsk/d5s0
```

35. Execute the following command to create a directory for mounting on each node:

The following shows an example of executing the command:

```
host1# mkdir -p /global/mntpt
host2# mkdir -p /global/mntpt
```

36. On the desired node, mount the file system. The following shows an example of executing the command:

```
host1# mount -g /dev/global/dsk/d5s0 /global/mntpt
```

3.5.4 Using the HDLM device unconfiguration function when performing a new installation of HDLM

You can use the HDLM device unconfiguration function to remove Hitachi storage subsystem devices from HDLM management on a per-LU basis. This section explains how to use the HDLM device unconfiguration function.

1. To remove an LU from HDLM management, use a text editor to edit the HDLM driver unconfiguration definition file (/opt/DynamicLinkManager/config/dlmfdrv.unconf).

In the lines that define the LU that is to be removed from HDLM management, delete the hash mark (#) in front of the LDEV property line. Do not change anything else. Figure 3.4 and Figure 3.5 show an example of editing the HDLM driver unconfiguration definition file (/opt/DynamicLinkManager/config/dlmfdrv.unconf).

```
####### T0 D0(9970/9980.15001.066A) #######
### Device=sd46 target=0 lun=0
### Path=/pci@8,700000/JNI,FCR@2/sd@0,0
### Device=sd47 target=1 lun=0
### Path=/pci@8,700000/JNI,FCR@2/sd@1,0
### Device=sd48 target=2 lun=0
### Path=/pci@8,700000/JNI,FCR@2/sd@2,0
#LDEV="9970/9980.15001.066A";
```

Figure 3.4 Example of an /opt/DynamicLinkManager/config/dlmfdrv.unconf file (before editing)

```
####### T0 D0(9970/9980.15001.066A) #######
### Device=sd46 target=0 lun=0
### Path=/pci@8,700000/JNI,FCR@2/sd@0,0
### Device=sd47 target=1 lun=0
### Path=/pci@8,700000/JNI,FCR@2/sd@1,0
### Device=sd48 target=2 lun=0
### Path=/pci@8,700000/JNI,FCR@2/sd@2,0
LDEV="9970/9980.15001.066A";
```

Figure 3.5 Example of an /opt/DynamicLinkManager/config/dlmfdrv.unconf file (after editing)

The shaded portion indicates the text that was edited.

2. Apply the contents of the HDLM driver unconfiguration definition file to the HDLM driver configuration definition file.

   Use the following command to apply the settings in the HDLM driver unconfiguration definition file (/opt/DynamicLinkManager/config/dlmfdrv.unconf) to the HDLM driver configuration definition file (/kernel/drv/dlmfdrv.conf):

   Execute the dlmsetconf utility for creating the HDLM driver configuration definition file.

   ```
   # /opt/DynamicLinkManager/bin/dlmsetconf -u
   ```

3. Restart the host with the reconfiguration option specified.

   Execute one of the following commands:

   - # reboot -- -r
   - boot -r (at the ok prompt)

4. Execute the following command to verify that the settings were applied correctly:

   ```
   # /opt/DynamicLinkManager/bin/dlnkmgr view -drv
   ```

   Confirm that the LDEV that was removed from HDLM management is not shown.

3.5.5 Preparations for an Upgrade Installation or Re-installation of HDLM

Before performing an upgrade installation or re-installation of HDLM, back up the HDLM-managed devices.

Back up the HDLM driver configuration definition file (/kernel/drv/dlmfdrv.conf)

Back up the HDLM driver configuration definition file (`/kernel/drv/dlmfdrv.conf`).

Stopping applications

HDLM manager is inactive during an upgrade installation or re-installation. For this reason, quit all applications that require log output and other HDLM manager operations.

3.5.6 Performing an update Performing an Upgrade Installation or Re-installation of HDLM

You can perform an upgrade installation on a host on which HDLM 05-00 or later is installed. If you are using an earlier version, uninstall it while referring to 3.16.3 *Uninstalling HDLM*, and then perform a new installation of HDLM 5.9.

When performing an update installation, you can use the existing HDLM driver configuration definition file (`/kernel/drv/dlmfdrv.conf`). If the path configuration has not changed, you will not need to execute the `dlmsetconf` utility for creating the HDLM driver configuration definition file.

When upgrading to HDLM 5.9 from a version earlier than 5.4, or when the license has expired, you must specify a license key file or enter an appropriate license key.

Note

- Because HDLM manager is inactive during an upgrade installation or re-installation, no error logs are output. Also, you cannot set HDLM functions during these types of installation. Restart the host immediately after the upgrade installation or re-installation has finished.
- Do not interrupt the upgrade installation or re-installation.
- When you install HDLM while the 32-bit kernel mode is being used, only the 32-bit version of HDLM will be installed. When you install HDLM while the 64-bit kernel mode is being used, both the 32-bit and 64-bit versions of HDLM will be installed. Before installing HDLM, decide which kernel mode you will use to operate Solaris and, if needed, change to the desired kernel mode. If you switch the kernel mode from 32 bits to 64 bits when only the 32-bit version of HDLM is installed, you must reinstall HDLM after the switch is completed. For details on how to switch the kernel mode, see *(5) Switch the kernel mode* in *Preparations for a New Installation of HDLM*.

1. When the license key file has been provided, Store the license key file directly under the `/var/tmp/` directory by using the `hdlm_license` name.

```
/var/tmp/hdlm_license
```

Note that the `hdlm_license` files is deleted after installation `finishes successfully`.

If the license key file does not exist, a message prompts you to enter a license key during installation. Enter the license key at this time.

2. Log in to Solaris as a user with root privileges.
3. Insert the HDLM CD-ROM into the CD-ROM drive.

   After a few seconds, the CD-ROM will be automatically mounted, and the mount point of the CD-ROM will be displayed in a list of files.

   If the CD-ROM cannot be automatically mounted, execute the following command to mount the CD-ROM:

```
# mount -F hsfs -o ro CD-ROM-device-name CD-ROM-mount-point
```

   The following shows an example of executing the command for mounting the CD-ROM:

```
# mkdir /cdrom/dlmanager
# mount -F hsfs -o ro /vol/dev/dsk/c0t6d0/dlmanager /cdrom/dlmanager
```

4. Execute the following install command:

```
# pkgadd -a CD-ROM-mount-point/DLManager/install/update -d CD-ROM-mount-point DLManager
```

   If the license key file does not exist, a message prompts you to enter a license key. Enter the license key.

5. Execute the following command to check that the package is installed, by the version number in the output results:

   The following shows an example of executing the command:

```
# pkginfo -l DLManager

   PKGINST:  DLManager
      NAME:  Dynamic Link Manager
  CATEGORY:  system
      ARCH:  sparc
   VERSION:  xx-xx
   BASEDIR:  /
    VENDOR:
    ...
#
```

6. Restart the host.

   When Sun Cluster is not being used:

   Execute one of the following commands:

   - `# reboot`
   - `boot` (at the `ok` prompt)

   When Sun Cluster is being used:

   Execute the following command at one of the nodes:

```
# /usr/cluster/bin/scshutdown -g0 -y
```

   Then, execute the following command at the `ok` prompt in all nodes:

```
boot
```

Note

You must restart the host immediately after a successful upgrade installation or re-installation has been confirmed.

7. Execute the `dlnkmgr` command's `view` operation to display the status of the programs.

   The following shows an example of executing the command:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -sys
HDLM Version                : xx-xx
Service Pack Version        :
Load Balance                : on(rr)
Support Cluster             :
Elog Level                  : 3
Elog File Size (KB)         : 9900
Number Of Elog Files        : 2
Trace Level                 : 0
Trace File Size(KB)         : 1000
Number Of Trace Files       : 4
Path Health Checking        : on(30)
Auto Failback               : off
Intermittent Error Monitor  : off
HDLM Manager Ver      WakeupTime
Alive        xx-xx    yyyy/mm/dd hh:mm:ss
HDLM Alert Driver Ver      WakeupTime          ElogMem Size
Alive             xx-xx    yyyy/mm/dd hh:mm:ss
HDLM Driver Ver      WakeupTime
Alive       xx-xx    yyyy/mm/dd hh:mm:ss
License Type Expiration
Permanent    -
KAPL01001-I The HDLM command completed normally. Operation name = view, completion time =
yyyy/mm/dd hh:mm:ss
#
```

   - The set values that were effective for each function before the update are inherited as they are.

     If you performed an upgrade installation from HDLM 05-00 or 05-01, the load-balancing setting will remain `on` if it was previously set `on`, and the algorithm will be `rr` (round robin). If you performed an upgrade installation from HDLM 05-02 or later, the setting and algorithm will be inherited if the load-balancing setting was previously `on`. The load-balancing algorithm options are round robin and extended round robin. See 3.9.2, and select the appropriate load balancing algorithm as required.

   - Nothing is displayed for `Support Cluster` even if you are using cluster software. However, the cluster-support functions are, in fact, operating without problem.

8. By checking the execution results of the `view` operation, verify the version of the installed HDLM.

   If `HDLM Version` is `05-90`, the installed version of HDLM is correct.

9. By checking the execution results of the `view` operation, verify that the programs are running normally.

   If `HDLM Manager`, `HDLM Alert Driver`, and `HDLM Driver` are all `Alive`, the programs are running normally.

3.6 Configuring a Boot Disk Environment

This section describes how to configure a boot disk environment in which HDLM manages the boot disk. An environment with a boot disk located on the host is called a *local boot disk environment*.

3.6.1 Overview of Configuring a Boot Disk Environment

To configure a boot disk environment in which HDLM manages the boot disk, prepare a separate disk that is not managed by HDLM, and set up an environment including Solaris, HDLM, and the applications to be used installed on that disk. If you already have such an environment, you do not need to set it up again. In this section, such an environment is called the *pre-migration environment* or *pre-migration*. Configure the boot disk environment using this prepared, or existing, pre-migration environment. The boot disk environment in which HDLM manages the boot disk is called the *post-migration environment* or *post-migration* in this section.

The procedures for using these three types of existing environments as pre-migration environments, and for building a new pre-migration environment, are summarized as follows:

- To use an existing local boot disk environment in which HDLM is installed as the pre-migration environment:

  Perform the procedure described in 3.6.2.1, and then execute the procedure described in 3.6.2.5.

- To use an existing local boot disk environment without HDLM installed as the pre-migration environment:

  Perform the procedure described in 3.6.2.2, and then execute the procedure described in 3.6.2.5.

- To use an existing boot disk environment without HDLM installed as the pre-migration environment:

  Perform the procedure described in 3.6.2.3, and then execute the procedure described in 3.6.2.5.

- Building a new pre-migration environment:

  Perform the procedure described in 3.6.2.4, and then execute the procedure described in 3.6.2.5.

3.6.2 Procedure for Configuring a Boot Disk Environment

The command examples and setting examples in this section assume the following slice allocations:

- Slice 0: / (root) file system
- Slice 1: Swap device and dump device
- Slice 3: Sun Cluster global device file system

- Slice 6: `/usr` file system
- Slice 7: Replica of the SDS or SVM state database

For details on procedures involving Solaris functionality, see your Solaris documentation.

3.6.2.1 Migration from an existing HDLM environment

Perform the steps below in the pre-migration environment. The workflow is shown in Figure 3.6.

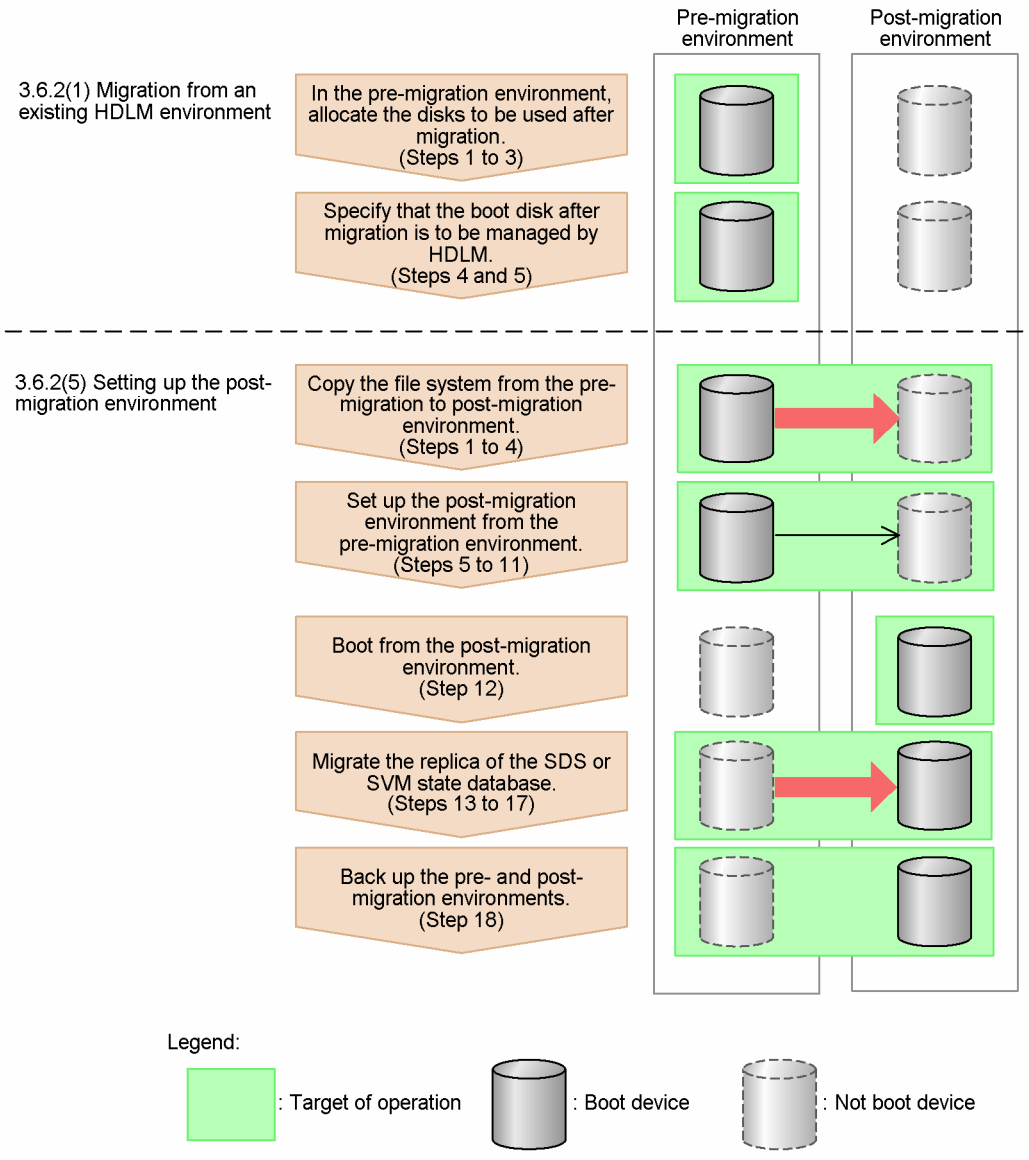

Figure 3.6 Flow of setting up the boot disk environment when using an existing HDLM environment as the pre-migration environment

1. Allocate the LUs to be used by the boot disk in the post-migration environment.
2. Get Solaris to recognize the LUs to be used by the post-migration boot disk.
3. Set the LU slices to be used by the post-migration boot disk.
4. Execute the `dlmsetconf` utility as follows to create the HDLM driver configuration definition file (`/kernel/drv/dlmfdrv.conf`) and the HDLM driver unconfiguration definition file (`/opt/DynamicLinkManager/config/dlmfdrv.unconf`):

   ```
   # /opt/DynamicLinkManager/bin/dlmsetconf
   ```

   This step and the next step set the LU to be used for the post-migration boot disk as an HDLM-managed LU.
5. Restart the host with the reconfiguration option specified to make the settings valid and the HDLM device recognizable.

   Execute one of the following commands:

   - `# reboot -- -r`
   - `boot -r` (at the `ok` prompt)

After you finish these steps, perform the procedure described in *Setting up the post-migration environment*.

3.6.2.2 Migration by installing HDLM in the existing local boot disk environment

Perform the steps below in the pre-migration environment. The workflow is shown in Figure 3.7.

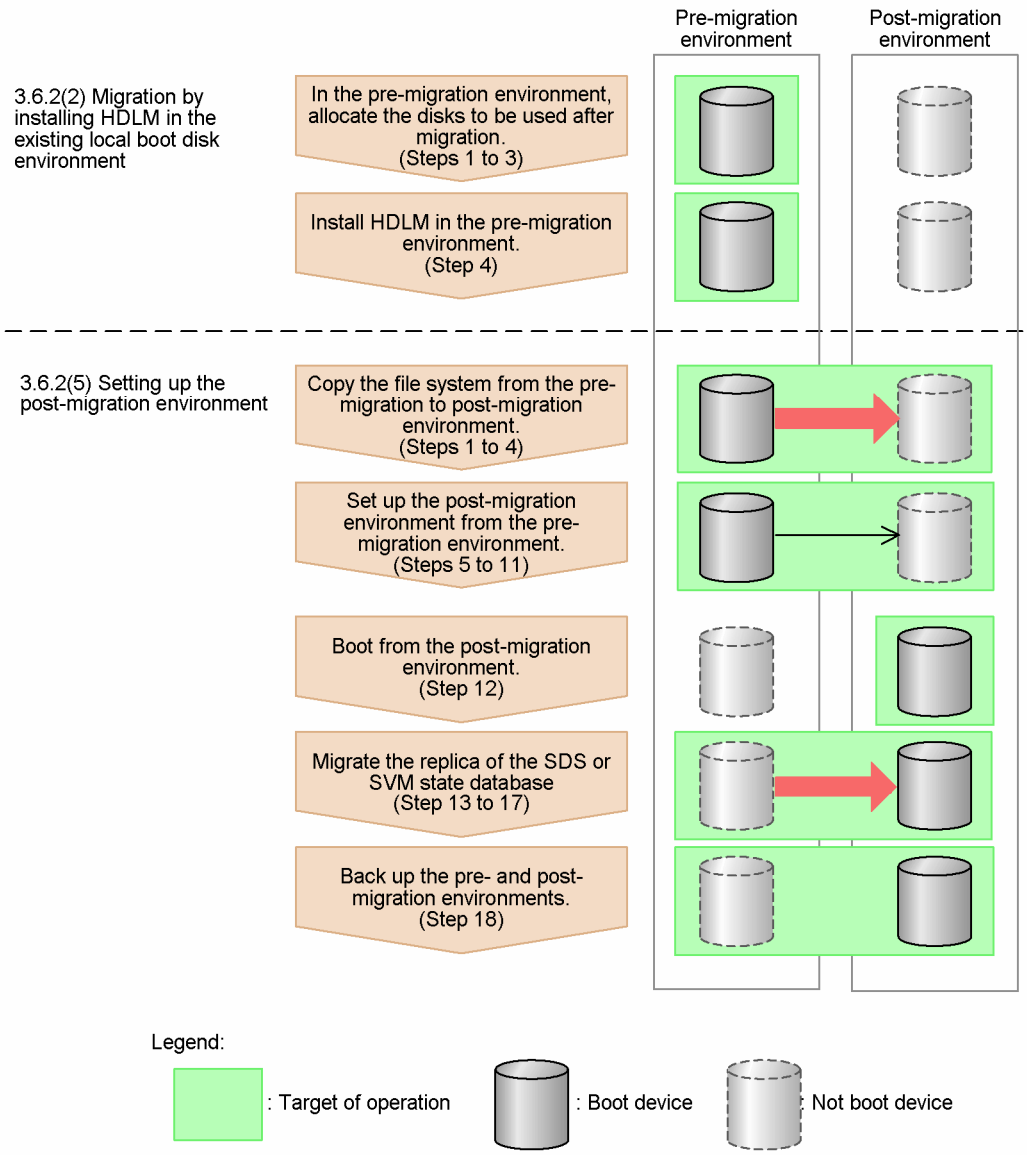

Figure 3.7 Flow of setting up the boot disk environment when installing HDLM in the existing local boot disk environment

1. Allocate the LUs to be used by the boot disk in the post-migration environment.

   Configure the LU or LUs to be used by the boot disk as follows:

   - To the slices of a single LU, allocate the Solaris file system, swap device, dump device, Sun Cluster global device file system, and replica of the SDS or SVM state database.
   - To the slices of multiple LUs, allocate the Solaris file system, swap device, dump device, Sun Cluster global device file system, and replica of the SDS or SVM state database.

     In this case, make sure that the physical path of each LU is the same.

2. Get Solaris to recognize the LUs to be used by the post-migration boot disk.
3. Set the LU slices to be used by the post-migration boot disk.

4. Complete the preparations for installing HDLM in the pre-migration environment and then install HDLM, referring to the description in *Installing HDLM*.

   Perform the necessary steps.

   Note

   Do not exclude the post-migration boot disk from HDLM management when you set the devices to be excluded in step 12 of *Performing a New Installation of HDLM (when Sun Cluster is not being used)* or in step 18 of *Performing a New Installation of HDLM (when Sun Cluster is being used)*.

After you finish these steps, perform the procedure described in 3.6.2.5.

3.6.2.3 Migration by installing HDLM in the existing boot disk environment

Perform the following steps in the pre-migration environment. The workflow is shown in Figure 3.8.

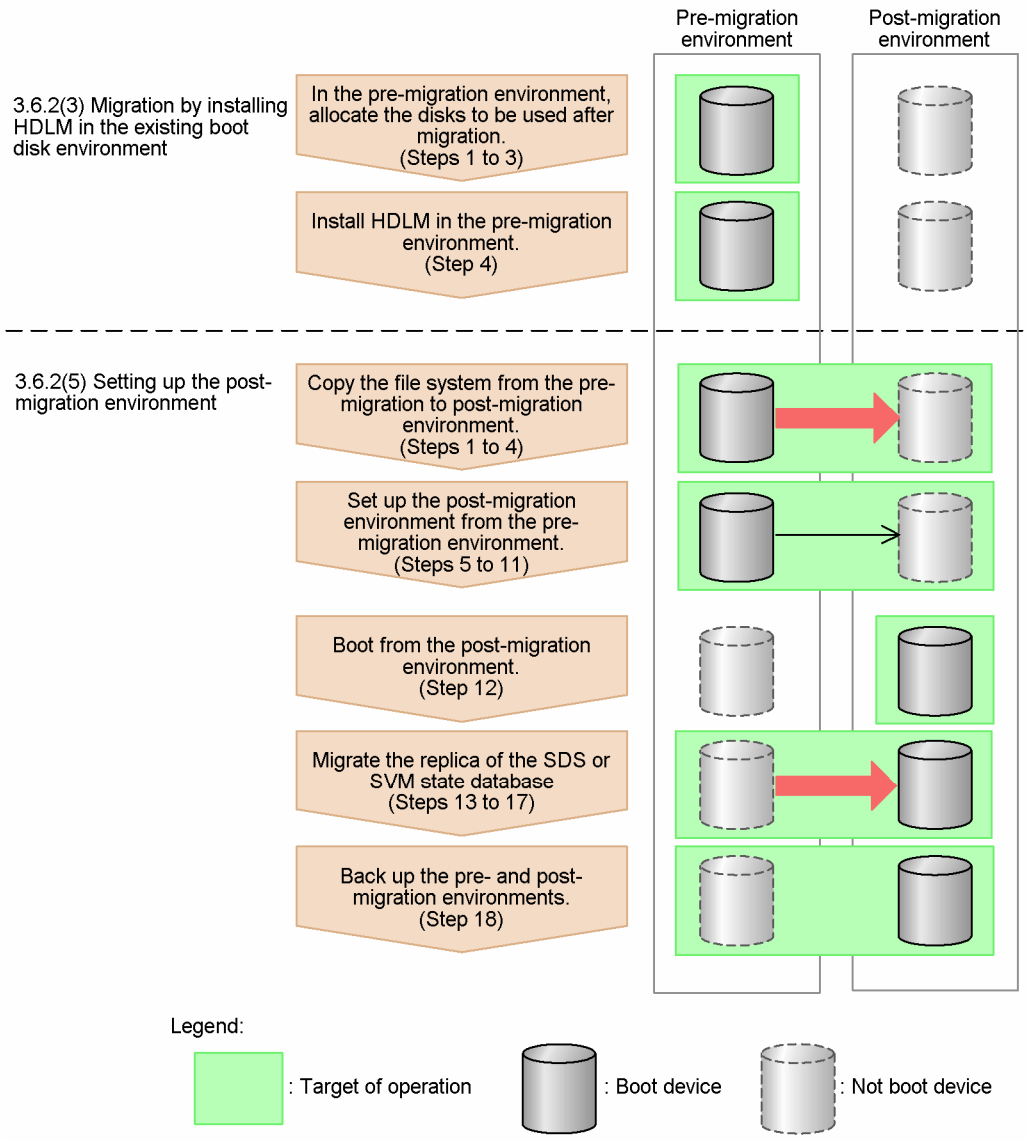

Figure 3.8 Flow of setting up the boot disk environment when using an existing boot disk environment without HDLM installed as the pre-migration environment

1. Allocate the LUs to be used by the boot disk in the post-migration environment.

   Configure the LU or LUs to be used by the boot disk as follows:

   - To the slices of a single LU, allocate the Solaris file system, swap device, dump device, Sun Cluster global device file system, and replica of the SDS or SVM state database.
   - To the slices of multiple LUs, allocate the Solaris file system, swap device, dump device, Sun Cluster global device file system, and replica of the SDS or SVM state database.

     In this case, make sure that the physical path of each LU is the same.

2. Get Solaris to recognize the LUs to be used by the post-migration boot disk.
3. Set the LU slices to be used by the post-migration boot disk.

4. Complete the preparations for installing HDLM in the pre-migration environment and then install HDLM, referring to the description in 3.5.

   Perform the necessary steps.

   Note

   Note the following when you set the devices to be excluded at step 12 of 3.5.2 or in step 18 of 3.5.3.

After you finish these steps, perform the procedure described in *Setting up the post-migration environment*.

3.6.2.4 Migration by building a new pre-migration environment

Perform the following steps in the pre-migration environment. The workflow is shown in Figure 3.9.

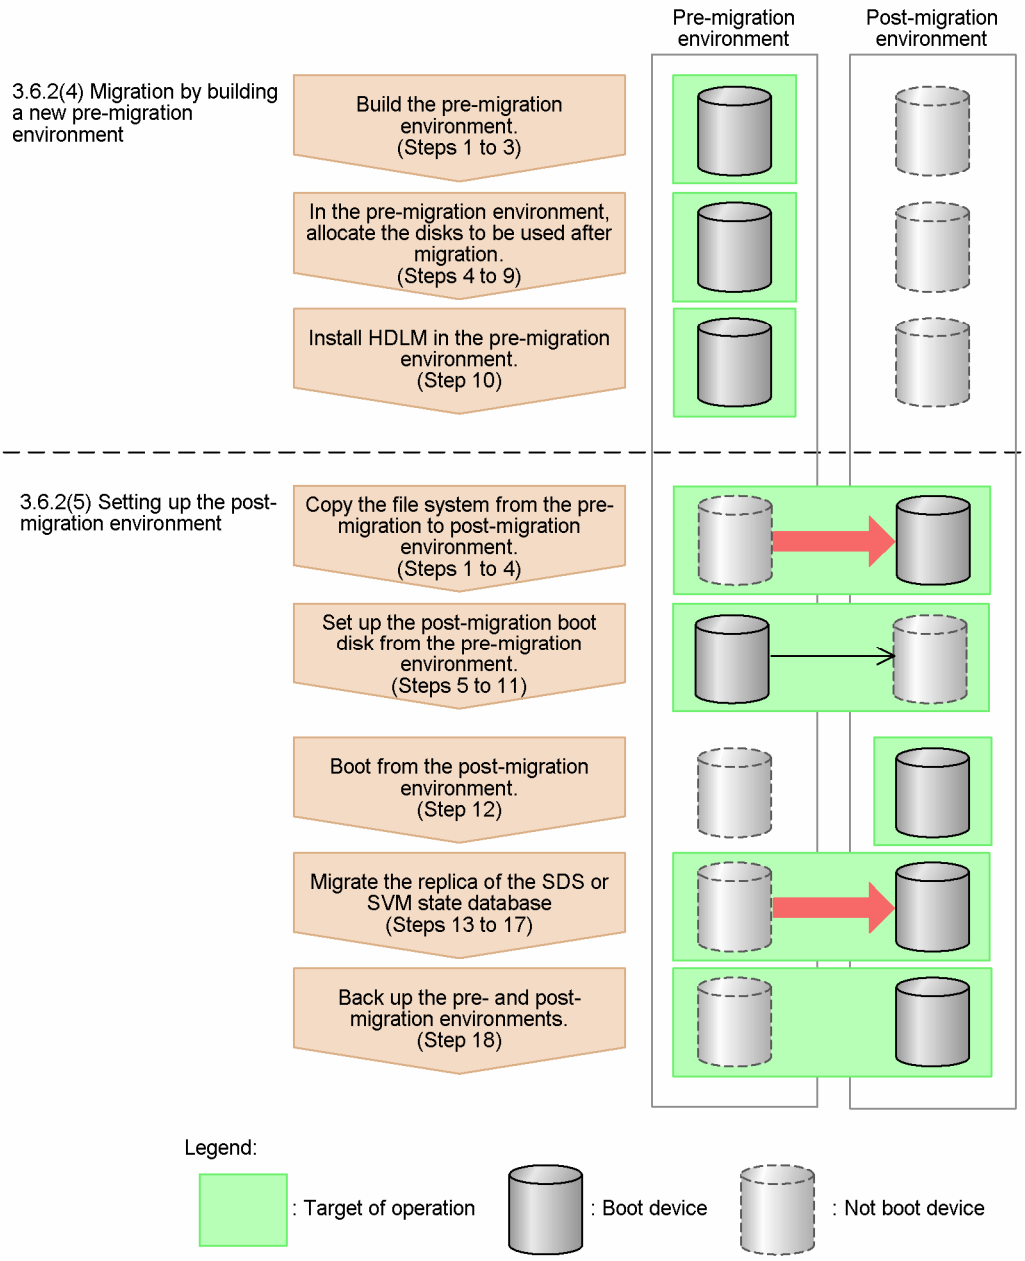

Figure 3.9 Flow of setting up the boot disk environment when building a new pre-migration environment

1. Install Solaris in the pre-migration environment.
2. Install Sun Cluster if required.

   For details, see your Sun Cluster documentation.
3. Install VxVM if required.

   For details, see your VxVM documentation.
4. Allocate the LUs to be used by the boot disk in the post-migration environment.

   Configure the LU or LUs to be used by the boot disk as follows:

- To the slices of a single LU, allocate the Solaris file system, swap devices, dump devices, Sun Cluster global device file system, and replica of the SDS or SVM state database.
- To the slices of multiple LUs, allocate the Solaris file system, swap devices, dump devices, Sun Cluster global device file system, and replica of the SDS or SVM state database.

  In this case, make sure that the physical path of each LU is the same.

5. Get Solaris to recognize the LUs to be used by the post-migration boot disk.
6. Set the LU slices to be used by the post-migration boot disk.
7. Allocate the LUs to be used in the post-migration environment by applications that require storage subsystem LUs other than the boot disk.
8. Get Solaris to recognize the disk used by applications if not already recognized.
9. Set the disk slice to be used by applications if not already set.
10. Complete the preparations for installing HDLM in the pre-migration environment and then install HDLM, referring to the description in 3.5.

    Perform the necessary steps.

After you finish these steps, perform the procedure described in 3.6.2.5.

3.6.2.5 Setting up the post-migration environment

1. Set the boot block for the post-migration environment, specifying the post-migration slice 0 (/ (root) file system).

   The following shows an example of executing the command:

```
# installboot /usr/platform/`uname -i`/lib/fs/ufs/bootblk
/dev/rdsk/c8t50060E80004361D1d1s0
```

2. Initialize the Solaris file system in the post-migration environment. If using Sun Cluster, also initialize the Sun Cluster global device file system.

   The following shows an example of executing the command:

```
# newfs /dev/rdsk/c8t50060E80004361D1d1s0
# newfs /dev/rdsk/c8t50060E80004361D1d1s6
# newfs /dev/rdsk/c8t50060E80004361D1d1s3
```

3. Mount the Solaris file system in the post-migration environment. If using Sun Cluster, also mount the Sun Cluster global device file system.

   The following shows an example of executing the command:

```
# mount /dev/dsk/c8t50060E80004361D1d1s0 /mnt0
# mount /dev/dsk/c8t50060E80004361D1d1s6 /mnt1
# mount /dev/dsk/c8t50060E80004361D1d1s3 /mnt2
```

4. Copy the Solaris file system from the pre-migration environment to the post-migration environment. If using Sun Cluster, also copy the Sun Cluster global device file system to the post-migration environment.

   The following shows an example of executing the command:

```
# ufsdump 0f - / | (cd /mnt0; ufsrestore rf -)
# ufsdump 0f - /usr | (cd /mnt1; ufsrestore rf -)
# ufsdump 0f - /global/.devices/node@1 | (cd /mnt2; ufsrestore rf -)
```

5. Delete the `rootdev` parameter if set in `/etc/system` in the post-migration environment (if you were using SDS, SVM, VxVM, or Sun StorageTek Traffic Manager Software in the pre-migration environment).

6. Execute the `ls -l` command to display information about the slice containing the `/` (root) file system of the post-migration boot disk.

   This example shows how to execute the command. The underlined HDLM physical device file name will be used in step 7.

```
# ls -l /dev/dsk/c8t50060E80004361D1d1s0
lrwxrwxrwx   1 root     root          60 Nov  6 17:09 /dev/dsk/c8t50060E80004361D1d1s0
-> ../../devices/pseudo/dlmndrv@1/dlmfdrv@w50060e80004361d1,1:a
```

7. Set the `rootdev` parameter for the slice containing the `/` (root) file system of the post-migration boot disk to the `/etc/system` file of the post-migration boot disk.

   In the `rootdev` parameter, specify the physical device file name of the HDLM device. In the following example, the underlined part is the physical device file name displayed in the execution example in step 6.

```
rootdev:/pseudo/dlmndrv@1/dlmfdrv@w50060e80004361d1,1:a
```

8. Edit the post-migration boot disk's `/etc/vfstab` file, changing the Solaris file system to the HDLM logical device file name of the post-migration boot disk. If using Sun Cluster, change the mount device# of the Sun Cluster global device file system in the same way.

   An example is shown below.

   Before editing

```
/dev/dsk/c0t0d0s0  /dev/rdsk/c0t0d0s0  /  ufs  1  no  -
/dev/dsk/c0t0d0s6  /dev/rdsk/c0t0d0s6  /usr  ufs  1  no  -
/dev/did/dsk/d24s3  /dev/did/rdsk/d24s3
/global/.devices/node@1  ufs  2  no  global
```

   After editing

```
/dev/dsk/c8t50060E80004361D1d1s0
/dev/rdsk/c8t50060E80004361D1d1s0 /   ufs   1   no   -
/dev/dsk/c8t50060E80004361D1d1s6
/dev/rdsk/c8t50060E80004361D1d1s6 /usr ufs  1  no  -
/dev/did/dsk/d30s3  /dev/did/rdsk/d30s3
/global/.devices/node@1  ufs  2  no  global
```

#

Execute the `scdidadm -l` command to find out the device ID to specify as the mount device of the Sun Cluster global device file system. In the following example, `d24` is the pre-migration boot disk, and `d30` is the post-migration boot disk.

```
# scdidadm -l
...
24   node1:/dev/rdsk/c0t0d0                  /dev/did/rdsk/d24
...
30   node1:/dev/rdsk/c8t50060E80004361D1d1  /dev/did/rdsk/d30
```

9. Edit the post-migration boot disk's `/etc/vfstab` file, changing the swap device definition from the pre-migration boot disk to the post-migration boot disk.

   An example is shown below.

   Before editing

   ```
   /dev/dsk/c0t0d0s1       -    -    swap    -    no    -
   ```

   After editing

   ```
   /dev/dsk/c8t50060E80004361D1d1s1    -    -    swap    -    no
   -
   ```

10. Edit the post-migration boot disk's `/etc/dumpadm.conf` file, changing the dump device definition in the `DUMPADM_DEVICE` parameter from the pre-migration boot disk to the post-migration boot disk.

    An example is shown below.

    Before editing

    ```
    DUMPADM_DEVICE=/dev/dsk/c0t0d0s1
    ```

    After editing

    ```
    DUMPADM_DEVICE=/dev/dsk/c8t50060E80004361D1d1s1
    ```

11. To enable the system to boot from the post-migration environment, find the boot device name to specify in the OpenBoot PROM (OBP).

    The paths that can be specified as the boot device in the OBP depend on the Solaris version, as follows:

    - In Solaris 8 or 9:

      As the boot device, you can specify any owner path of the boot disk.

    - In Solaris 10:

      As the boot device, you can only specify a boot disk path used by HDLM to determine the logical device file name of the HDLM device. To find this path, perform the procedure described in 4.4.1.3.

    An example in which Sun Microsystems' HBA is used with Solaris 8 or 9 is presented below. For details on finding the boot device name when using an HBA other than from Sun Microsystems, and on setting the HBA in the OBP, see the documentation for the HBA you are using.

    Execute the `dlnkmgr` command's `view` operation with the `-lu`, `-item`, `pd`, `type`, and `-hdev` parameters specified. An execution example is shown below.

```
# dlnkmgr view -lu -item pd type -hdev c8t50060E80004361D1d1
Product       : AMS
SerialNumber  : 73010236
LUs           : 50

iLU  HDevName               Device PathID Status     Type Physical Device
0066 c8t50060E80004361D1d1  ssd23  000001 Online     Own  /devices/ssm@0,0/pci@1c,600000/
pci@1/SUNW,qlc@4/fp@0,0/ssd@w50060e80004361d1,1
                            ssd89  000051 Online     Non  /devices/ssm@0,0/pci@1c,600000/
pci@1/SUNW,qlc@4/fp@0,0/ssd@w50060e80004361d3,1
                            ssd145 000101 Online     Own  /devices/ssm@0,0/pci@1c,600000/
pci@1/SUNW,qlc@5/fp@0,0/ssd@w50060e80004361d1,1
                            ssd235 000151 Online     Non  /devices/ssm@0,0/pci@1c,600000/
pci@1/SUNW,qlc@5/fp@0,0/ssd@w50060e80004361d3,1

KAPL01001-I The HDLM command completed normally. Operation name = view, completion time =
yyyy/mm/dd hh:mm:ss
#
```

Edit the `Physical Device` character string of any owner path in the execution results as follows. The edited path will be the boot device name that you need to specify in the OBP.

- Delete the `/device` part.
- Replace the driver name of the ssd driver (`ssd`) with `disk`.
- Add the slice identifier of the `/` (root) file system (any ID from `:a` to `:h`) to the owner path.

By editing the character strings in the underlined part of the above execution example in this way, you get the following character string. This is the boot device name to specify in the OBP:

```
/ssm@0,0/pci@1c,600000/pci@1/SUNW,qlc@4/fp@0,0/disk@w50060e80004361d1,1:a
```

Note

In Solaris 8 or 9, find and make a note of the boot device names to specify in the OBP for all owner paths in case an error occurs on the path specified for the boot device.

12. Boot the system from the post-migration environment.
    - Perform steps 13 to 17 only if SDS or SVM is being used and you are migrating a replica of the SDS or SVM state database to the post-migration environment. In this case, keep the pre-migration environment recognized by Solaris.
    - Skip to step 18 if SDS or SVM is being used but you are not migrating a replica of the SDS or SVM state database to the post-migration environment.
    - Skip to step 18 if SDS or SVM is not being used.
13. Check the state database settings.

    An example is shown below.

```
# metadb
        flags           first blk       block count
     a        u         16              8192            /dev/dsk/c0t0d0s7
     a        u         16              8192            /dev/dsk/c0t1d0s7
     a        u         16              8192            /dev/dsk/c0t2d0s7
#
```

14. Add the post-migration boot disk to the SDS or SVM state databse.

An example is shown below.

```
# metadb -a /dev/dsk/c8t50060E80004361D1d1s7
# metadb -a /dev/dsk/c8t50060E80004361D1d2s7
# metadb -a /dev/dsk/c8t50060E80004361D1d3s7
```

15. Make sure that the post-migration boot disk has been added to the SDS or SVM state database.

    An example is shown below.

```
# metadb
        flags           first blk       block count
     a        u         16              8192            /dev/dsk/c0t0d0s7
     a        u         16              8192            /dev/dsk/c0t0d1s7
     a        u         16              8192            /dev/dsk/c0t0d2s7
     a        u         16              8192           /dev/dsk/c8t50060E80004361D1d1s7
     a        u         16              8192           /dev/dsk/c8t50060E80004361D1d2s7
     a        u         16              8192           /dev/dsk/c8t50060E80004361D1d3s7
#
```

16. Delete the SDS or SVM state database replica that exists in the pre-migration boot disk.

    The following shows an example of executing the command:

```
# metadb -d /dev/dsk/c0t0d0s7 /dev/dsk/c0t0d1s7 /dev/dsk/c0t0d2s7
```

17. Make sure that the SDS or SVM state database replica that exists in the pre-migration boot disk has been deleted.

    The following shows an example of executing the command:

```
# metadb
        flags           first blk       block count
     a        u         16              8192            /dev/dsk/c8t50060E80004361D1d1s7
     a        u         16              8192            /dev/dsk/c8t50060E80004361D1d2s7
     a        u         16              8192            /dev/dsk/c8t50060E80004361D1d3s7
#
```

18. Back up the contents of both the pre-migration and post-migration environments.

    Perform a backup so that you can take action in the following situations:

    - Recover when an error occurs in the post-migration environment.

      Boot from the pre-migration environment, and repair the post-migration environment. The post-migration backup can be used to restore the post-migration environment.

    - Migrate from a boot disk environment to the local boot disk environment, and return to the state before setting up the boot disk environment.

3.7 Migrating from a Boot Disk Environment to the Local Boot Disk Environment

This section describes how to migrate from a boot disk environment in which HDLM manages the boot disk to the local boot disk environment (boot disk configured on the host). In this section, the former is sometimes called the *pre-migration environment* or *pre-migration*, and the latter is sometimes called the *post-migration environment* or *post-migration*. There are two methods of performing this migration:

- By restoring the file system you backed up when configuring the HDLM boot disk environment

  This method restores the system status before you configured the boot disk environment. Restore the file system according to the Solaris and application specifications. For the procedure, see your Solaris and application documentation.

- By copying the file system of the boot disk environment to the local boot disk environment

  This method migrates to the local boot disk environment with the current operating status intact. The procedure is outlined and described below.

An overview of the procedure is shown in Figure 3.10.

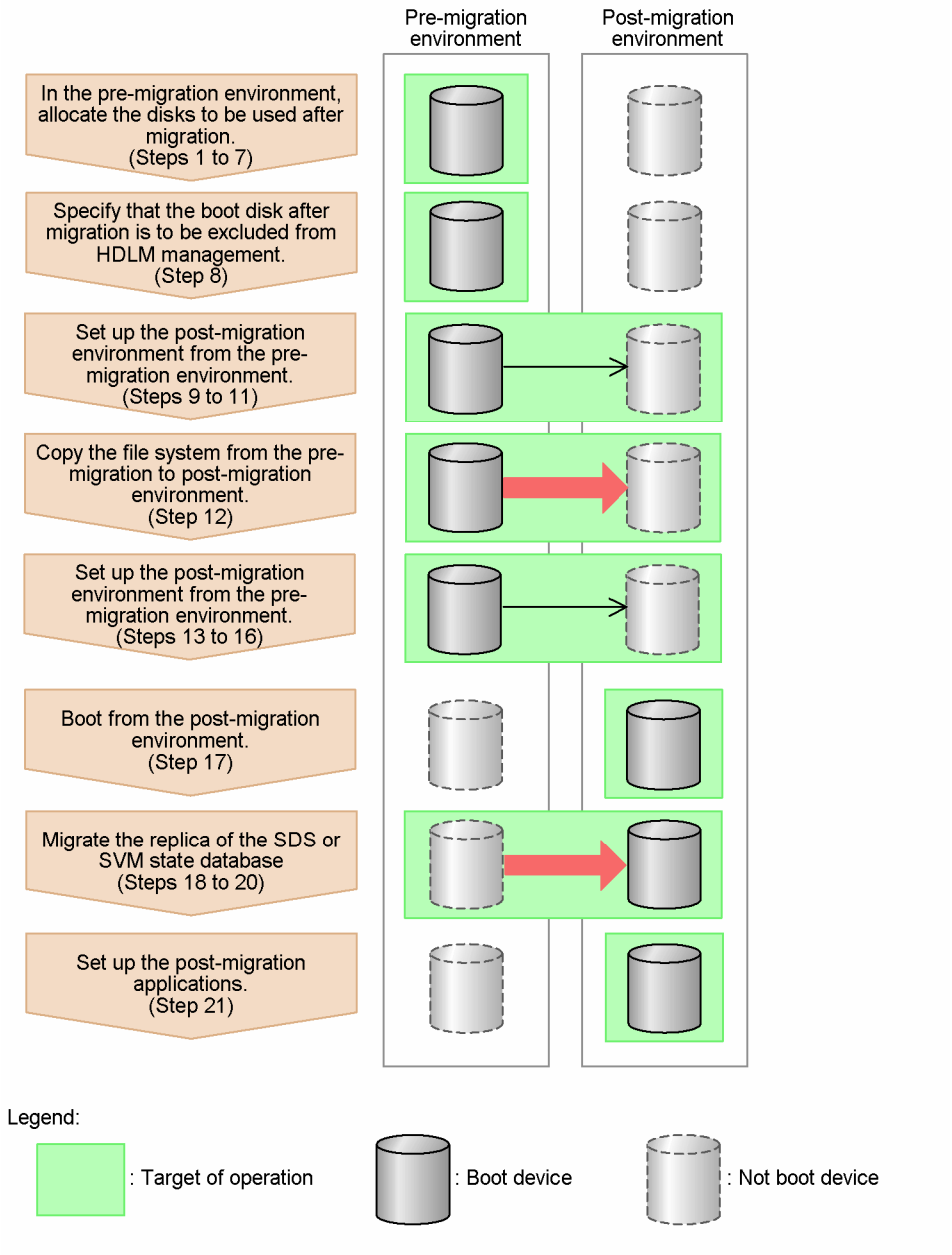

Figure 3.10 Migration flow when coping the file system from a boot disk environment to the local boot disk environment

1. Boot the system from the pre-migration environment.

   Perform steps 1 to 16 in the pre-migration environment.

2. Allocate the LUs to be used by the boot disk in the post-migration environment.
3. Get Solaris to recognize the LUs to be used by the post-migration boot disk.
4. Set the LU slices to be used by the post-migration boot disk.
   - If using Sun Cluster, go to the next step.
   - If not using Sun Cluster, go to step 8.

5. Assign device IDs.

   The following shows an example of executing the command:

```
# /usr/cluster/bin/scdidadm -r
```

6. Notify the server of the new device ID information.

   The following shows an example of executing the command:

```
# /usr/cluster/bin/scdidadm -ui
```

7. Update the Sun Cluster global device namespace.

   The following shows an example of executing the command:

```
# /usr/cluster/bin/scgdevs
```

8. If the post-migration boot disk is an HDLM-managed device, exclude it from HDLM management by executing the procedure described in 3.5.4.

9. Set the boot block for the post-migration environment, specifying the post-migration slice 0 (/ (root) file system).

   The following shows an example of executing the command:

```
# installboot /usr/platform/`uname -i`/lib/fs/ufs/bootblk /dev/rdsk/c0t0d0s0
```

10. Initialize the Solaris file system in the post-migration environment. If using Sun Cluster, also initialize the Sun Cluster global device file system.

    The following shows an example of executing the command:

```
# newfs /dev/rdsk/c0t0d0s0
# newfs /dev/rdsk/c0t0d0s6
# newfs /dev/rdsk/c0t0d0s3
```

11. Mount the Solaris file system in the post-migration environment. If using Sun Cluster, also mount the Sun Cluster global device file system.

    The following shows an example of executing the command:

```
# mount /dev/dsk/c0t0d0s0 /mnt0
# mount /dev/dsk/c0t0d0s6 /mnt1
# mount /dev/dsk/c0t0d0s3 /mnt2
```

12. Copy the Solaris file system from the pre-migration environment to the post-migration environment. If using Sun Cluster, also copy the Sun Cluster global device file system to the post-migration environment.

    The following shows an example of executing the command:

```
# ufsdump 0f - / | (cd /mnt0; ufsrestore rf -)
# ufsdump 0f - /usr | (cd /mnt1; ufsrestore rf -)
# ufsdump 0f - /global/.devices/node@1 | (cd /mnt2; ufsrestore rf -)
```

13. In the post-migration `/etc/system` file, delete the `rootdev` parameter that specifies the HDLM physical device file name.

14. Edit the post-migration boot disk's `/etc/vfstab` file, changing the Solaris file system to the HDLM logical device file name of the post-migration boot disk. If using Sun Cluster, change the mount device# of the Sun Cluster global device file system in the same way.

    An example is shown below.

    Before editing

    ```
    /dev/dsk/c8t50060E80004361D1d1s0
    /dev/rdsk/c8t50060E80004361D1d1s0 /   ufs   1   no   -
    /dev/dsk/c8t50060E80004361D1d1s6
    /dev/rdsk/c8t50060E80004361D1d1s6 /usr ufs  1  no  -
    /dev/did/dsk/d30s3  /dev/did/rdsk/d30s3
    /global/.devices/node@1  ufs  2  no  global
    ```

    After editing

    ```
    /dev/dsk/c0t0d0s0  /dev/rdsk/c0t0d0s0  /  ufs  1  no  -
    /dev/dsk/c0t0d0s6  /dev/rdsk/c0t0d0s6  /usr  ufs  1  no  -
    /dev/did/dsk/d24s3  /dev/did/rdsk/d24s3
    /global/.devices/node@1  ufs  2  no  global
    ```

    #

    Execute the `scdidadm -l` command to find out the device ID to specify as the mount device of the Sun Cluster global device file system. In the following example, `d30` is the pre-migration boot disk, and `d24` is the post-migration boot disk.

    ```
    # scdidadm -l
    ...
    24   node1:/dev/rdsk/c0t0d0                  /dev/did/rdsk/d24
    ...
    30   node1:/dev/rdsk/c8t50060E80004361D1d1  /dev/did/rdsk/d30
    ```

15. Edit the post-migration boot disk's `/etc/vfstab` file, changing the swap device definition from the pre-migration boot disk to the post-migration boot disk.

    An example is shown below.

    Before editing

    ```
    /dev/dsk/c8t50060E80004361D1d1s1    -    -    swap    -    no
    -
    ```

    After editing

    ```
    /dev/dsk/c0t0d0s1       -    -    swap    -    no    -
    ```

16. Edit the post-migration boot disk's `/etc/dumpadm.conf` file, changing the dump device definition in the `DUMPADM_DEVICE` parameter from the pre-migration boot disk to the post-migration boot disk.

    An example is shown below.

    Before editing

```
DUMPADM_DEVICE=/dev/dsk/c8t50060E80004361D1d1s1
```

After editing

```
DUMPADM_DEVICE=/dev/dsk/c0t0d0s1
```

17. Boot the system from the post-migration environment.
    - Skip to step 18 if SDS or SVM is being used and you are migrating a replica of the SDS or SVM state database to the post-migration environment. In this case, keep the pre-migration environment recognized by Solaris.
    - Skip to step 21 if SDS or SVM is being used but you are not migrating a replica of the SDS or SVM state database.
    - Skip to step 21 if SDS or SVM is not being used.
18. Check the state database settings.

    An example is shown below.

```
# metadb
        flags           first blk       block count
     a        u         16              8192            /dev/dsk/c8t50060E80004361D1d1s7
     a        u         16              8192            /dev/dsk/c8t50060E80004361D1d2s7
     a        u         16              8192            /dev/dsk/c8t50060E80004361D1d3s7
#
```

19. Add the post-migration boot disk to the SDS or SVM state database.

    An example is shown below.

```
# metadb -a /dev/dsk/c0t0d0s7
# metadb -a /dev/dsk/c0t0d1s7
# metadb -a /dev/dsk/c0t0d2s7
# metadb
        flags           first blk       block count
     a        u         16              8192            /dev/dsk/c8t50060E80004361D1d1s7
     a        u         16              8192            /dev/dsk/c8t50060E80004361D1d2s7
     a        u         16              8192            /dev/dsk/c8t50060E80004361D1d3s7
     a        u         16              8192            /dev/dsk/c0t0d0s7
     a        u         16              8192            /dev/dsk/c0t0d1s7
     a        u         16              8192            /dev/dsk/c0t0d2s7
#
```

20. Delete the SDS or SVM state database replica that exists in the pre-migration boot disk.

    An example is shown below.

```
# metadb -d /dev/dsk/c8t50060E80004361D1d1s7
# metadb -d /dev/dsk/c8t50060E80004361D1d2s7
# metadb -d /dev/dsk/c8t50060E80004361D1d3s7
# metadb
        flags           first blk       block count
     a        u         16              8192            /dev/dsk/c0t0d0s7
     a        u         16              8192            /dev/dsk/c0t0d1s7
     a        u         16              8192            /dev/dsk/c0t0d2s7
#
```

21. Perform the required settings if you want to mirror or multi-path the post-migration boot disk in the following applications.

    For details on the required settings, see the documentation for the relevant application.

- SDS
- SVM
- VxVM
- Sun StorageTek Traffic Manager Software

3.8 Checking the Path Configuration

HDLM functions, such as load balancing and failovers, are only available for HDLM management-target *devices* that have more than one active path. After you install HDLM or change the hardware configuration, check the structure and statuses of these paths.

To check the path information, use the HDLM command's `view` operation.

The following describes how to check path information by using the `dlnkmgr` command's `view` operation. For details about the view operation, see 6.7.

Specify the `-path` parameter and check the output information:

Execute the following command:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -path > redirect-destination-file-1
```

Open *redirect-destination-file-1* and check the following:

- Make sure that an LU accessed by a path exists.

  A path can be identified with `PathName`. The LU that is accessed by a path can be identified with a combination of `DskName` and `iLU`.

- Make sure that all paths are online.

  Make sure that `PathStatus` is `Online`. If there is a path whose status is not online, `Reduced` will be displayed.

- Make sure that there are enough *host-port-number*.*bus-number* combinations for the number of physical HBA ports.

- Make sure that the combinations of the CHA port (`ChaPort`), through which paths access the same LU, and the HBA port (the host port number and bus number displayed in the `PathName` column) are different.

  The two-digit number from the left of the numbers displayed for `PathName` indicates an host port number. The numbers displayed between the period to the right of the host port number and the next period indicate a bus number.

Specify the `-drv` parameter and check the output information:

Execute the following command:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -drv > redirect-destination-file-2
```

Open *redirect-destination-file-2* and check the `HDevName` and `Device` entries, making sure that an HDLM device corresponding to the sd or ssd device has been created.

3.9 Setting Up HDLM Functions

HDLM has load balancing, automatic failback, error logging, and other functions. You can set up these functions by using the `dlnkmgr` command's `set` operation. The following subsections describe these setup methods.

3.9.1 Checking the current settings

Before setting the HDLM functions by using the `dlnkmgr` command's `set` operation, check the current settings by executing the following command:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -sys -sfunc
HDLM Version              : xx-xx
Service Pack Version        :
Load Balance              : on(rr)
Support Cluster            :
Elog Level               : 3
Elog File Size(KB)          : 9900
Number Of Elog Files        : 2
Trace Level              : 0
Trace File Size(KB)         : 1000
Number Of Trace Files       : 4
Path Health Checking        : on(30)
Auto Failback             : off
Intermittent Error Monitor    : off
KAPL01001-I The HDLM command completed normally. Operation name = view, completion time =
yyyy/mm/dd hh:mm:ss
#
```

3.9.2 Setting up HDLM functions

The recommended and default values for the HDLM functions are shown below in Table 3.17.

Table 3.17 The default and recommended values for each function

Function	Default value	Recommended value
Load-balancing	on Algorithm is round robin	on The recommended algorithm depends on the operating environment.
Path health checking	on 30-minute check interval	on The recommended checking interval depends on the operating environment.
Automatic failback	off	off
Intermittent Error Monitor	off	off
Logging level	3: Collect all information for errors with the level "Information" or higher	3: Collect all information for errors with the level "Information" or higher
Trace level	0: Do not output a trace	0: Do not output a trace
File size (for Error log)	9900 (kilobytes)	9900 (kilobytes)
Number of files (for Error log)	2	2
File size (for Trace)	1000 (kilobytes)	1000 (kilobytes)
Number of files (for Trace)	4	4

3.9.2.1 Setting load balancing

You can select whether to enable the load-balancing function.

The following is an example command to set load-balancing.

```
# /opt/DynamicLinkManager/bin/dlnkmgr set -lb on -lbtype rr
```

Set it to `on` to enable load balancing. Otherwise, set it to `off`. When you set this value to `on`, specify `rr` for round robin or `exrr` for extended round robin after the `-lbtype` option. The type of algorithm specified by the `-lbtype` parameter remains stored in the system, even when you disable the load balancing function by specifying `-lb off`. Therefore, when you re-enable the load balancing function without specifying an algorithm, load balancing will be executed according to the algorithm stored in the system.

3.9.2.2 Setting path health checking

You can select whether to enable the path-health-checking function.s

The following is an example command to set path health checking.

```
# /opt/DynamicLinkManager/bin/dlnkmgr set -pchk on -intvl 10
```

To enable path health checking, set it to `on`. Otherwise, set it to `off`. When you set this value to `on`, you can use the `-intvl` parameter to specify the *checking* interval. The value that was specified previously will be applied if the *checking* interval is not specified. For example, specify the path health check as `off` after specifying the *checking* interval as `15` minutes and executing. Then, when executing after specifying the path health check as `on` without specifying the *checking* interval, the `15` minutes that were specified previously will be applied again.

3.9.2.3 Setting automatic failback

You can select whether to enable the automatic-failback function.

When intermittent error monitoring is enabled and the number of error occurrence is `2` or more, the following condition must be satisfied.

```
error-monitoring-interval >= checking-interval-for-automatic-failback x
number-of-times-error-is-to-occur-during-intermittent-error-monitoring
```

If this condition is not satisfied, an error occurs and the warning KAPL01080-W message is output.

In such a case, change any of the following settings: the checking interval for automatic failback, intermittent error monitoring interval, or the number of times that the error is to occur.

When you set the number of times that the error is to occur to `1`, the condition above does not need to be satisfied.

The following shows an example of executing the command to set automatic failback:

```
# /opt/DynamicLinkManager/bin/dlnkmgr set -afb on -intvl 10
```

To enable automatic failback, set it to `on` . Otherwise, set it to `off`. When `on` is specified, you can use the `-intvl` parameter to specify the *checking* interval. If the *checking* interval is not specified, the value that was previously specified will be applied. For example, specify the auto failback as `off` after specifying the *checking* interval to five minutes and executing. Then, when executing after specifying the auto failback as `on` without specifying the *checking* interval, the five minutes that were specified previously will be applied again.

3.9.2.4 Settings for intermittent error monitoring

You can set the intermittent error monitoring only when automatic failback is enabled. To prevent an intermittent error from reducing I/O performance, we recommend that you monitor intermittent errors when automatic failback is enabled.

When intermittent error monitoring is enabled, you can specify intermittent error conditions (the conditions used by the system to determine whether an intermittent error is occurring). The default value for the intermittent error monitoring interval is `30`. The default value for the number of error occurrences is `3`.

The system assumes that an intermittent error is occurring if the specified number of times that the error is to occur is reached during the monitoring interval (from the time that the monitoring interval starts, until the specified interval ends). A path that is assumed to have an intermittent error is excluded from automatic failback. Intermittent error monitoring starts when the path is recovered from the error by using automatic failback. Monitoring is performed on individual paths.

When a value of `2` or more is specified in the number of error occurrence, make sure that the condition shown in 3.9.2.3 is satisfied. To determine whether an intermittent error is occurring in a path, you can use the execution results of the `dlnkmgr` command's `view` operation.

The following shows an example of executing the command to set intermittent error monitoring.

```
# /opt/DynamicLinkManager/bin/dlnkmgr set -iem on -intvl 20 -iemnum 2
```

To enable intermittent error monitoring, set it to `on`. Otherwise, set it to `off`. When `on` is specified, you can use the `-intvl` and `-iemnum` parameters to specify intermittent error conditions (the conditions used by the system to determine whether an intermittent error is occurring). Specify the monitoring interval for an intermittent error in the `-intvl` parameter, and the number of times that the error is to occur in the `-iemnum` parameter. When these parameters are omitted and three or more errors occur within 30 minutes, the system assumes that an intermittent error is occurring.

3.9.2.5 Setting the error log collection level

The error log (the HDLM manager log (`dlmmgr`*n*`.log` (*n* indicates a file number from `1` to `16`))) collection level can be set.

Table 3.18 lists and describes values for the error log collection level setting.

Table 3.18 Values for the error log collection level setting

Value	Description
0	No error logs are collected.
1	All information for errors with the level "Error" or higher is collected.
2	All information for errors with the level "Warning" or higher is collected.
3	All information for errors with the level "Information" or higher is collected.
4	All information for errors with the level "Information" or higher (including maintenance information) is collected.

If an error occurs, you might have to change the collection level to `1` or higher to collect the log information.

The higher this value is set, the more log information will be output. When the output log is large, it takes less time to overwrite the old error log information with the new information.

The following shows an example of executing the command to set the error log collection level:

```
# /opt/DynamicLinkManager/bin/dlnkmgr set -ellv 2
```

Specify the error log collection level in a number.

3.9.2.6 Setting the trace level

You can set the trace output level.

A trace file in which you can set a trace level is `hdlmtr`*n*`.log` (*n* indicates a file number from `1` to `64`).

Table 3.19 lists and describes values for the trace level setting.

Table 3.19 Values for the trace level setting

Value	Description
0	No trace is output.
1	Only error information is output.
2	Program operation summaries are output.
3	Program operation details are output.
4	All information is output.

If an error occurs, you might have to change the trace level to 1 or higher to collect the log information.

The higher this value is set, the more log information will be output.

When the output log is large, it takes less time to overwrite the old error log information with the new information.

The following is an example command to set the trace level:

```
# /opt/DynamicLinkManager/bin/dlnkmgr set -systflv 1
```

Specify the trace level in a number.

3.9.2.7 Setting the error log file size

The error log file size (the HDLM manager log (dlmmgr*n*.log (*n* indicates a file number from 1 to 16))) can be set.

You can specify a value (in kilobytes) from 100 to 2000000 for the error log file size. The specified value is applied for HDLM manager logs.

If each error log file reaches the specified size, the information in the old error log file will be wrapped around and will be overwritten with new information. By specifying both the log file size and the number of log files, you can collect up to 32000000 kilobytes (approximately 30GB) of error logs in total.

The following shows an example of executing the command to set the error log file size.

```
# /opt/DynamicLinkManager/bin/dlnkmgr set -elfs 1000
```

Specify the size of the error log file in kilobytes.

3.9.2.8 Setting the number of error log files

The number of the error log files (the HDLM manager log (dlmmgr*n*.log (*n* indicates a file number from 1 to 16))) can be set.

You can specify a value from 2 to 16 for the number of error log files (log files for the HDLM manager).

By specifying both the log file size and the number of log files, you can collect up to 32000000 kilobytes (approximately 30GB) of error logs in total.

The following shows an example command that is used to set the number of error log files:

```
# /opt/DynamicLinkManager/bin/dlnkmgr set -elfn 5
```

Specify the number of error log files in numbers.

3.9.2.9 Setting the trace file size

You can set the trace file size.

Trace files for which you can set a trace file size are `hdlmtr`*n*`.log` (*n* indicates a file number from `1` to `64`). The length of a trace file is fixed. Thus, even if the amount of written trace information is less than the set file size, the file size of each output trace file is always fixed.

For the trace file size, you can specify a value (in kilobytes) from 100 to 16000. If you specify a value smaller than the setting value, the message is displayed to confirm execution, and the trace file is temporarily deleted.

Once trace data has been written to all of the trace files, the oldest file is overwritten with new trace data.

By specifying both the trace file size and the number of trace files, you can collect up to 1024000 kilobytes of trace data in total.

The following shows an example command that is used to set the trace file size:

```
# /opt/DynamicLinkManager/bin/dlnkmgr set -systfs 2000
```

Specify the size of the trace file in kilobytes.

3.9.2.10 Setting the number of trace files

You can set the number of the trace files.

Trace files for which you can set the number of files are `hdlmtr`*n*`.log` (*n* indicates a file number from `1` to `64`).

For the number of the trace files, you can specify a value from 2 to 64. If you specify a value smaller than the setting value, the KAPL01097-W message is displayed to confirm execution, and the trace file is temporarily deleted.

By specifying both the trace file size and the number of trace files, you can collect up to 1024000 kilobytes of trace data in total.

The following shows an example command that is used to set the number of trace files:

```
# /opt/DynamicLinkManager/bin/dlnkmgr set -systfn 10
```

Specify the number of trace files in numbers.

3.9.3 Checking the new settings

This section describes how to check the new settings by using the `dlnkmgr` command's `set` operation after the new settings are applied.

When you change these settings, you can display information about all HDLM function settings. The following an example of the command used to display the settings information and its execution results:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -sys -sfunc
HDLM Version                  : xx-xx
Service Pack Version          :
Load Balance                  : on(rr)
Support Cluster               :
Elog Level                    : 2
Elog File Size(KB)            : 1000
Number Of Elog Files          : 5
Trace Level                   : 1
Trace File Size(KB)           : 2000
Number Of Trace Files         : 10
Path Health Checking          : on(10)
Auto Failback                 : on(10)
Intermittent Error Monitor    : on(2/20)
KAPL01001-I The HDLM command completed normally. Operation name = view, completion time =
yyyy/mm/dd hh:mm:ss
#
```

3.10 Setting up Integrated traces

When HDLM is used, `dlnkmgr` commands logs are output to the *integrated trace information files* of Hitachi Network Objectplaza Trace Library (HNTRLib2) (`/var/opt/hitachi/HNTRLib2/spool/hntr2`*n*`.log` (*n* indicates a file number)).

If there is a significant amount of integrated trace information output, information may end up deleted in a short amount of time. Also, if a large amount of integrated trace information is output at once, the integrated trace information that overflowed the buffer might not be saved in integrated trace files. To save as much information as possible, change the settings for Hitachi Network Objectplaza Trace Library to increase the integrated trace file size and buffer size. Note that specifying the values too large places a heavy load on the system. When determining these values, consider these operational tradeoffs.

Table 3.20 shows the default and recommended values for the integrated trace file settings.

Table 3.20 Default and recommended values for the integrated trace file settings

Setting		Default value	Recommended value
Integrated trace file size		256 (kilobytes)	4096 (kilobytes)
Number of integrated trace files		4	8
Buffer size per monitoring interval	Monitoring interval	10 (seconds)	5 (seconds)
	Buffer size	64 (kilobytes)	256 (kilobytes)
Number of messages to be output per monitoring interval	Monitoring interval	0 (seconds)	0 (seconds)
	Number of messages to be output	0 (transactions)	0 (transactions)

If Hitachi Network Objectplaza Trace Library (HNTRLib2) is already installed, the existing settings will be inherited. If you change these settings, keep in mind that they are used by programs other than HDLM.

3.10.1 Notes on Using the Hitachi Network Objectplaza Trace Library

Note the following when using Hitachi Network Objectplaza trace library:

- If HNTRLib has already been installed when you install HDLM, the settings in the trace library will not be inherited by HNTRLib2. HDLM uses the HNTRLib2 default settings.
- If HNTRLib2 has already been installed on a host when you install HDLM, the settings in the trace library will be inherited.
- If a different Hitachi product is using HNTRLib2 when you attempt to uninstall HDLM, HNTRLib2 will not be uninstalled

3.10.2 Displaying the Hitachi Network Objectplaza Trace Library setup menu

To display the Hitachi Network Objectplaza Trace Library setup menu:

1. Log in as a user with root privileges.
2. Execute the following command:

```
# /opt/hitachi/HNTRLib2/bin/hntr2utl2
```

The Hitachi Network Objectplaza Trace Library setup menu appears.

```
Hitachi Network Objectplaza Trace Library 2 - Configuration Utility  Rel 2.0

  Select the item you want to change.  (Type 1-7 or e)

        [Log Files]
     1: Size of a log file.      256 KB
     2: Number of log files.       4
     3: Name of log files.      /var/opt/hitachi/HNTRLib2/spool/hntr2*.log

        [Monitor]
     4: Size of buffer.           64 KB
     5: Interval timer.           10 Sec

        [Logging Restriction]
     6: Lookout span.              0 Sec
     7: Max messages per span.     0

     e: Exit

    Enter the number>
```

If you do not want to change the settings, type `e` and then press the **Enter** key to quit the menu.

The following explains how to modify each setting.

3.10.3 Changing the size of integrated trace files

To change the size of integrated trace files:

1. In the Hitachi Network Objectplaza Trace Library setup menu, type `1` and then press the **Enter** key.

   A screen to set the size of the integrated trace file appears. The current value is displayed in **Current Size(KB)**.

```
Hitachi Network Objectplaza Trace Library 2 - Configuration Utility  Rel 2.0

  Type new file size [8-8192]       (Type '!' to return)


    Current Size(KB): 256
        New Size(KB):
```

2. Enter the desired size in **New Size(KB)**.

   You can specify a value from 8 to 8192 (KB). The default value is 256. Set this to a value larger than the buffer value set in step 2 of 3.10.5. The recommend value is 4096 for collecting an integrated trace.

   If you do not want to change the integrated trace file size, leave **New Size(KB)** blank, type `!`, and then press the **Enter** key to return to the Hitachi Network Objectplaza Trace Library setup menu.

3. Press the **Enter** key.

   The new setting is applied and the Hitachi Network Objectplaza Trace Library setup menu appears again.

3.10.4 Changing the number of integrated trace files

To change the number of integrated trace files:

1. In the Hitachi Network Objectplaza Trace Library setup menu, type `2` and then press the **Enter** key.

   A screen to set the number of integrated trace files appears. The current value is displayed in **Current Number(KB)**.

```
Hitachi Network Objectplaza Trace Library 2 - Configuration Utility  Rel 2.0

   Type the number of files [1-16]       (Type '!' to return)


     Current Number(KB):  4
         New Number(KB):
```

2. Enter the desired number in **New Number(KB)**.

   The default value is 4. You can specify a value from 1 to 16. The The value set here becomes the maximum of *n* in `/var/opt/hitachi/HNTRLib2/spool/hntr2`*n*`.log`. The recommended value for integrated trace collection is 8.

   If you do not want to change the number of integrated trace files, leave **New Number(KB)** blank, type `!`, and then press the **Enter** key to return to the Hitachi Network Objectplaza Trace Library setup menu.

3. Press the **Enter** key.

   The new setting is applied and the Hitachi Network Objectplaza Trace Library setup menu appears again.

3.10.5 Changing the buffer size per monitoring interval duration

To change the buffer size per monitoring interval:

1. In the Hitachi Network Objectplaza Trace Library setup menu, type `4`, and then press the **Enter** key.

   A screen to set the buffer size appears. The current value is displayed in **Current Size(KB)**.

```
Hitachi Network Objectplaza Trace Library 2 - Configuration Utility  Rel 2.0

   Type new buffer size [8-2048]     (Type '!' to return)


     Current Size(KB):  64
         New Size(KB):
```

2. Enter a larger size in **New Size(KB)**.

Set a new buffer size to fit the monitoring interval set in **5: Interval Timer**. The default value is 64. You can specify a value from 8 to 2048 (KB). Set this to a value smaller than the integrated trace file size set in step 2 of 3.10.3. The recommended value is 256 for collecting an integrated trace.

If you do not want to change the buffer size, leave **New Size(KB)** blank, type `!`, and then press the **Enter** key to return to the Hitachi Network Objectplaza Trace Library setup menu.

3. Press the **Enter** key.

   The new setting is applied and the Hitachi Network Objectplaza Trace Library setup menu appears again.

4. In the Hitachi Network Objectplaza Trace Library setup menu, type `5` and then press the **Enter** key.

   A screen to set the monitoring interval appears. The current value is displayed in **Current Span(sec)**.

```
Hitachi Network Objectplaza Trace Library 2 - Configuration Utility  Rel 2.0

  Type the value of interval timer for the monitor [1-300](Type '!' to return)


    Current Span(sec): 10
        New Span(sec):
```

5. Enter a shorter interval in **New Span(sec)**.

   The default value is 10. You can specify a value from 1 to 300 (seconds). The recommended value is 5 for collecting an integrated trace.

   If you do not want to change the monitoring interval, leave **New Span(sec)** blank, enter `!` , and then press the **Enter** key to return You will be returned to the Hitachi Network Objectplaza Trace Library setup menu.

6. Press the **Enter** key.

   The new setting is applied and the Hitachi Network Objectplaza Trace Library setup menu appears again.

3.10.6 Adjusting the number of messages to be output per monitoring interval

To adjust the number of messages to fit a particular monitoring interval:

1. In the Hitachi Network Objectplaza Trace Library setup menu, type `6` and then press the **Enter** key.

   A screen to set the monitoring interval for the number of messages output to the integrated trace file will appears. The current value is displayed in **Current Span (sec)**.

```
Hitachi Network Objectplaza Trace Library 2 - Configuration Utility  Rel 2.0

  Type the number of lookout span [1-3600 or 0]     (Type '!' to return)


    Current Span(sec):   0
        New Span(sec):
```

2. Enter a desired interval in **New Span(sec)**.

   You can specify a value from 0 to 3600 (seconds). The default and recommended value is 0.

   If you do not want to change the monitoring interval, leave **New Span(sec)** blank, type `!`, and then press the **Enter** key to return to the Hitachi Network Objectplaza Trace Library setup menu.

   Note that when you specify a monitoring interval of 0, even if you specify the maximum number of messages in **7: Max messages per span**, the amount of integrated trace information to be output will not be adjusted.

3. Press the **Enter** key.

   The new setting is applied and the Hitachi Network Objectplaza Trace Library setup menu appears again.

4. In the Hitachi Network Objectplaza Trace Library setup menu, type `7`, and then press the **Enter** key.

   A screen to set the maximum number of messages output to the integrated trace file based on the monitoring interval specified in **6: Lookout span** appears. The current value is displayed in **Current Max(sec)**.

```
Hitachi Network Objectplaza Trace Library 2 - Configuration Utility  Rel 2.0

  Type the number of max messages [0-500]     (Type '!' to return)


    Current Max(sec):   0
        New Max(sec):
```

5. Adjust the maximum number of messages output to the integrated trace files in **New Max(sec)**.

   The default is 0. You can specify a value from 0 to 500 (messages). If you want to increase the number of messages that are output to the integrated trace file as much as possible, we recommend that you set this value to 0.

   When you specify a monitoring interval of 0 in **6: Lookout span**, the value set in **New Max(sec)** will be disregarded.

   Also, when you specify a value of 0 for **New Max(sec)**, even if you specify the monitoring interval in **6: Lookout span**, the maximum number of messages output will not be adjusted.

   If you do not want to change the maximum number of messages output, leave **New Max(sec)** blank, enter `!`, and then press the **Enter** key to return to the Hitachi Network Objectplaza Trace Library setup menu.

6. Press the **Enter** key.

   The new setting is applied and the Hitachi Network Objectplaza Trace Library setup menu appears again.

3.10.7 Finishing the Hitachi Network Objectplaza Trace Library settings

This section explains how to close the Hitachi Network Objectplaza Trace Library setup menu when you are finished:

1. In the Hitachi Network Objectplaza Trace Library setup menu, type `e`, and then press the **Enter** key.

   You will be asked to if you wish to save the new settings.

```
Save or not? (Yes/No)>
```

2. To save the new settings, type `Yes`, otherwise, type `No`.

3.10.8 Applying the Hitachi Network Objectplaza Trace Library settings

To apply the settings, after you change the amount of integrated trace information by using Hitachi Network Objectplaza Trace Library:

1. Log in as a user with root privileges.
2. Execute the following command to check the programs using HNTRLib2.

   In the following example, only HDLM uses HNTRLib2.

```
# /opt/hitachi/HNTRLib2/etc/hntr2dgetname
Hitachi Dynamic Link Manager
#
```

3. Stop the programs that are using HNTRLib2.

   If programs other than HDLM are displayed in step 2, stop the programs, and then go to step 4. You do not need to stop the HDLM manager.

   If you do not know how to stop and start programs other than HDLM, do not perform the following steps, and restart the host.

4. Execute the following command to stop the integrated trace collection process:

```
# /opt/hitachi/HNTRLib2/bin/hntr2kill
```

5. Execute the following command to delete the memory mapped file:

```
# rm /opt/hitachi/HNTRLib2/mmap/hntr2mmap.mm
```

6. Execute the following command to start the integrated trace collection process:

```
# /opt/hitachi/HNTRLib2/bin/hntr2mon -d &
```

```
```

7. Start the programs stopped in step 3.

   If you stopped programs other than HDLM in step 3, start them.

3.11 Creating File Systems for HDLM (when volume management software is not used)

When you are not using any volume management software, you need to either mount HDLM-managed devices as file systems or use them as character devices. To use an HDLM-managed device as a character device, use the logical device file `/dev/rdsk/c`*U*`t`*X*`d`*Y*`s`*Z*.

This section explains how to specify the logical device file name of an HDLM device and mount the HDLM-managed device as a file system.

To mount the HDLM-managed device as a file system:

1. Create a file system by using an HDLM device.

   If the file system already exists, skip to step 2.

   The following shows an example of the command that is used to create a file system:

```
# newfs /dev/rdsk/c2t7d16s0
```

2. Create a directory on which to mount the file system.

   Execute a command such as the following:

```
# mkdir mount-point
```

3. Mount the file system.

   Execute a command such as the following:

```
# mount -F ufs /dev/dsk/c2t7d16s0 mount-point
```

4. Enable automatic mounting.

   To automatically mount the HDLM-managed device when the host starts, edit the `/etc/vfstab` file.

   Execute the following command to display the current setting:

```
# mount -p
```

   Figure 3.11 shows an output example of the command.

```
# mount -p
/dev/dsk/c0t8d0s0 - / ufs - no rw,intr,largefiles,onerror=panic,suid,dev=800038
/proc - /proc proc - no dev=4080000
fd - /dev/fd fd - no rw,suid,dev=4140000
mnttab - /etc/mnttab mntfs - no dev=4240000
swap - /var/run tmpfs - no dev=1
swap - /tmp tmpfs - no dev=2
/dev/dsk/c0t0d0s7 - /Sol7_export ufs - no rw,intr,largefiles,onerror=panic,suid,dev=800007
/dev/dsk/c0t0d0s0 - /Sol7_sys ufs - no rw,intr,largefiles,onerror=panic,suid,dev=800000
/dev/dsk/c0t8d0s7 - /export/home ufs no rw,intr,largefiles,onerror=panic,suid,dev=80003f
/dev/dsk/c2t7d16s0 - /mntpt  ufs - no rw,intr,largefiles,onerror=panic,suid,dev=37c0000
#
```

Figure 3.11 An output example for the mount -p command

Edit the `/etc/vfstab` file as appropriate based on the settings output. The shaded portion indicates the part of the `/etc/vfstab` file to set.

An example of how to edit the `/etc/vfstab` file is shown below in Figure 3.12.

```
#device          device          mount      FS   fsck mount   mount
#to mount        to fsck         point      type pass at boot options
#
#/dev/dsk/c1d0s2  /dev/rdsk/c1d0s2  /usr       ufs  1   yes     -
fd               -               /dev/fd    fd   -   no      -
/proc            -               /proc      proc -   no      -
/dev/dsk/c0t8d0s1 -              -          swap -   no      -
/dev/dsk/c0t8d0s0 /dev/rdsk/c0t8d0s0/          ufs  1   no      -
/dev/dsk/c0t8d0s7 /dev/rdsk/c0t8d0s7/export/home ufs  2   yes     -
swap             -               /tmp       tmpfs -   yes     -
/dev/dsk/c0t0d0s0 -              /Sol7_sys   ufs  -   yes     -
/dev/dsk/c0t0d0s7 -              /Sol7_export ufs  -   yes     -
/dev/dsk/c2t7d16s0 /dev/rdsk/c2t7d16s0 /mntpt   ufs  2   yes     -
```

Figure 3.12 An example of how to edit the /etc/vfstab file

Add the shaded portion to the file.

3.12 Setting Up VxVM

When you are using VxVM, VxVM accesses an LU by using an HDLM device.

VxVM will automatically recognize HDLM devices. Perform setup as you would for any other device, by following the VxVM documentation.

3.12.1 Creating a disk group

This subsection mainly describes the preparations for registering the HDLM device in the VxVM disk group.

To register the HDLM device in the VxVM disk group, you must remove non-HDLM devices from VxVM. If you have not yet removed non-HDLM devices, first remove them as described in 3.12.3.2 or 3.12.3.3 in 3.12.3, and then perform the following procedure:

To register an HDLM device:

1. Make sure that VxVM recognizes HDLM devices.

   Execute the following commands sequentially to check whether VxVM recognizes HDLM devices.

   - The `dlnkmgr` command's `view -path` operation

     Check the HDLM devices that must be recognized by VxVM. The following shows an example of executing the command:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -path
Paths:000006 OnlinePaths:000006
PathStatus   IO-Count    IO-Errors
Online       0           0

PathID PathName                          DskName
iLU            ChaPort Status     Type  IO-Count    IO-Errors  DNum  HDevName
000000 0000.0000.0000000000000001.0000 HITACHI .OPEN-3      -SUN.15001
0672           2C      Online     Own           0            0     -  c6t1d0
000001 0000.0000.0000000000000001.0001 HITACHI .OPEN-3      -SUN.15001
0673           2C      Online     Own           0            0     -  c6t1d1
000002 0000.0000.0000000000000001.0002 HITACHI .OPEN-3      -SUN.15001
0679           2C      Online     Own           0            0     -  c6t1d2
000003 0001.0000.0000000000000004.0000 HITACHI .OPEN-3      -SUN.15001
0672           1C      Online     Own           0            0     -  c6t1d0
000004 0001.0000.0000000000000004.0001 HITACHI .OPEN-3      -SUN.15001
0673           1C      Online     Own           0            0     -  c6t1d1
000005 0001.0000.0000000000000004.0002 HITACHI .OPEN-3      -SUN.15001
0679           1C      Online     Own           0            0     -  c6t1d2
KAPL01001-I The HDLM command completed normally. Operation name = view,
completion time = yyyy/mm/dd hh:mm:ss
#
```

     The following describes the procedure for checking whether the shaded HDLM devices are recognized.

   - `vxdisk` command

     Make sure that the devices are recognized by VxVM. The following shows an example of executing the command:

```
# vxdisk list
DEVICE       TYPE   DISK       GROUP    STATUS
HDS9980V0_0  sliced  -          -        online
HDS9980V0_1  sliced  -          -        online
HDS9980V0_2  sliced  -          -        online
c1t1d0s2     sliced  -          -        error
c1t2d0s2     sliced  rootdisk   rootdg   online
#
```

     If an enclosure name (`HDS9980V0_1`) is displayed in the `DEVICE` column for the device, execute the following command to check the device name, in the format `c`*U*`t`*X*`d`*Y*`s`*Z*, corresponding to the enclosure name (HDS9980V0_1).

```
# vxdisk list HDS9980V0_1
Device:     HDS9980V0_1
devicetag:  HDS9980V0_1
type:       sliced
...
Multipathing information:
numpaths:   1
c6t1d1s2  state=enabled   type=secondary
#
```

Note

The enclosure name corresponding to the LU might vary before and after installing HDLM. Make sure that you perform step 1 to check the correspondence between the enclosure name and the LU.

If the `vxdisk` command displays the value of `HDevName` (c*U*t*X*d*Y*) output by the `dlnkmgr` command's `view -path` operation, the HDLM device is recognized by VxVM. Go to step 2. If the `vxdisk` command does not display the value of `HDevName` (c*U*t*X*d*Y*) output by the `dlnkmgr` command's `view -path` operation, the HDLM device is not recognized by VxVM. In this case, perform step 2, then re-execute step 1 to make sure that the HDLM device is recognized.

2. Execute the `vxdmpadm` command for the devices displayed by the `vxdisk` command to make sure that only the HDLM devices are `ENABLED`.

   To check whether the HDLM devices are `ENABLED`, see the value of `HDevName (c`*U*t*X*d*Y*`)` output by the `dlnkmgr` command's `view -path` operation. Also, make sure that all non-HDLM devices are removed from VxVM management and not displayed.

   The following shows an example of executing the command:

```
# vxdmpadm getsubpaths dmpnodename=HDS9980V0_1
NAME          STATE          PATH-TYPE  CTLR-NAME  ENCLR-TYPE    ENCLR-NAME
=============================================================================
c6t1d1s2      ENABLED        -          c6         HDS9980V      HDS9980V0
```

Note

The load balancing and failover functions will not work normally if the `vxdmpadm` command outputs non- HDLM devices (devices that are not displayed for HDevName of the `dlnkmgr` command's `view -path` operation), and if `STATE` is `ENABLED`. Therefore, non-HDLM devices must be removed from VxVM management. To remove non- HDLM devices from VxVM management, see 3.12.3.2 or 3.12.3.3 of 3.12.3. To apply the settings performed in procedure (2) or (3) in 3.12.3, restart the host.

3. Create a disk group by referring to the VxVM documentation.

   An HDLM device can be added to VxVM as a disk of the `sliced` type. When using VxVM 4.0 or later, you can also add the HDLM device to VxVM as a disk of the `cdsdisk` type.

   Note

   When using VxVM 3.1.1, VxVM 3.2, or VxVM 3.5, do not add the HDLM device to `rootdg`.

When using VxVM 4.0 or later, do not add the HDLM device to the disk group registered as `bootdg`. Also, do not register the disk group containing the HDLM device in `bootdg`.

4. If you are using Sun Cluster, register the created disk group into a Sun Cluster disk device group.

   For details, see 3.12.5.

3.12.2 Creating VxVM volumes

A VxVM volume can be created by using the `vxassist` (or `vxmake`) command or VERITAS Enterprise Administrator. For details, refer to the VxVM documentation.

Note

If you have created a volume for a disk group registered in Sun Cluster, register the change made to the VxVM disk group configuration into Sun Cluster. For details, see 3.12.5.

When creating a VxVM volume, a block device file and character device file for accessing volumes will be created:

- Block device file

  `/dev/vx/dsk/`*disk-group-name*`/`*volume-name*

- Character device file

  `/dev/vx/rdsk/`*disk-group-name*`/`*volume-name*

When creating a file system, mounting a file system, or specifying the layout of a database using a raw partition, use the device file appropriate to each situation.

In the following example, the commands create and mount a file system.

```
# newfs /dev/vx/rdsk/dlmdg/vol01
# mount /dev/vx/dsk/dlmdg/vol01 /mnt
```

3.12.3 Removing devices from VxVM

VxVM uses HDLM devices to access LUs managed by HDLM. Therefore, you must remove non-HDLM devices (such as sd or ssd devices), which correspond to the LU, from VxVM management.

In this section, excluding devices from VxVM management is called *removing devices*.

The following describes devices that must be removed from VxVM, and the procedure for removing such devices.

3.12.3.1 Devices to be removed from VxVM

Figure 3.13 shows devices that are recognized by VxVM before they are removed from VxVM.

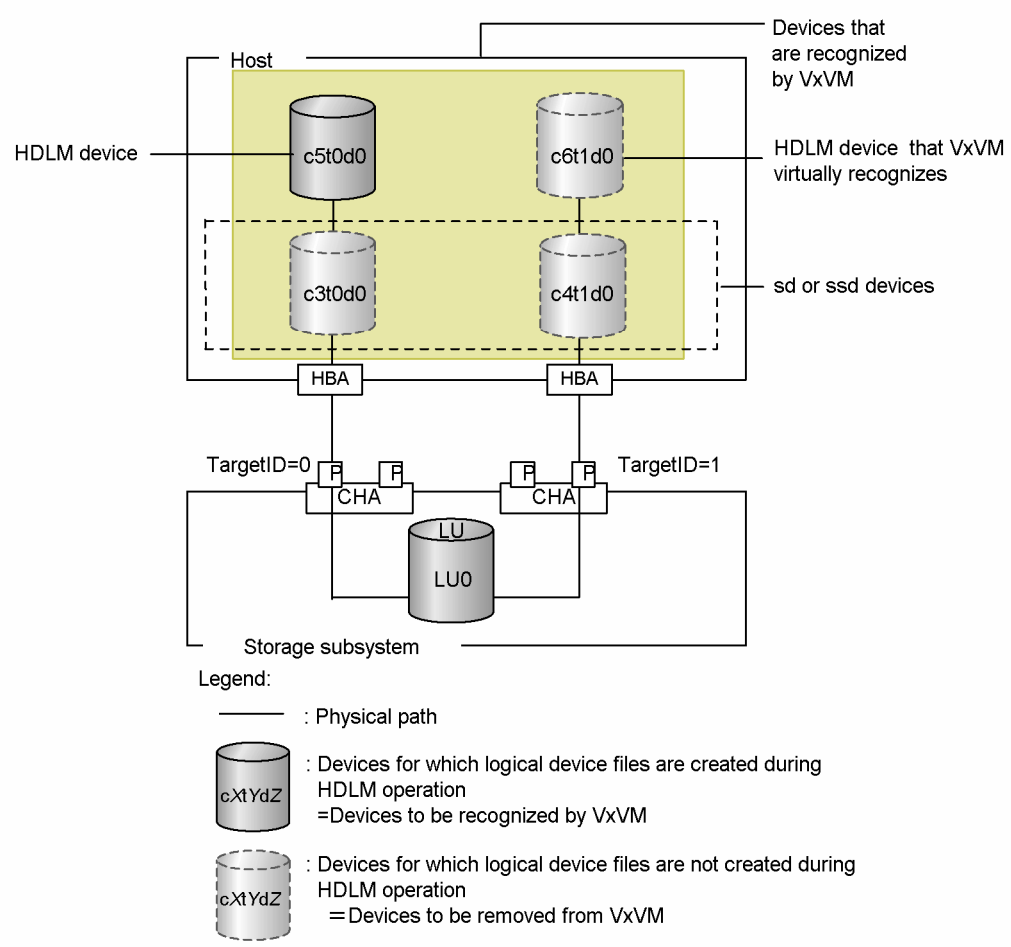

Figure 3.13 Devices That Are Recognized by VxVM before They Are Removed from VxVM

You must remove the following non-HDLM devices that are recognized by VxVM. If the following devices are being managed by VxVM, HDLM functions, such as load balancing and failover, will work normally.

- sd or ssd devices (c3t0d0, c4t1d0)

  When an HDLM device is created, the logical device file of the corresponding sd or ssd device will be deleted. However, VxVM recognizes the sd or ssd device as still existing even if the logical device file was deleted.

- HDLM devices assumed by VxVM (c6t1d0) (when VxVM 3.1.1, or VxVM 3.5 or later is being used)

  Devices that VxVM virtually recognizes as HDLM devices. These devices are not actually created.

Only one HDLM device is created for one LU#. However, for one LU, VxVM recognizes the same number of HDLM devices as the number of physical paths.

#

For logical device files of HDLM devices, the number of such files to be created is the same as the number of slices for one LU.

Figure 3.14 shows the devices that are recognized by VxVM after removal of the non-HDLM devices from VxVM.

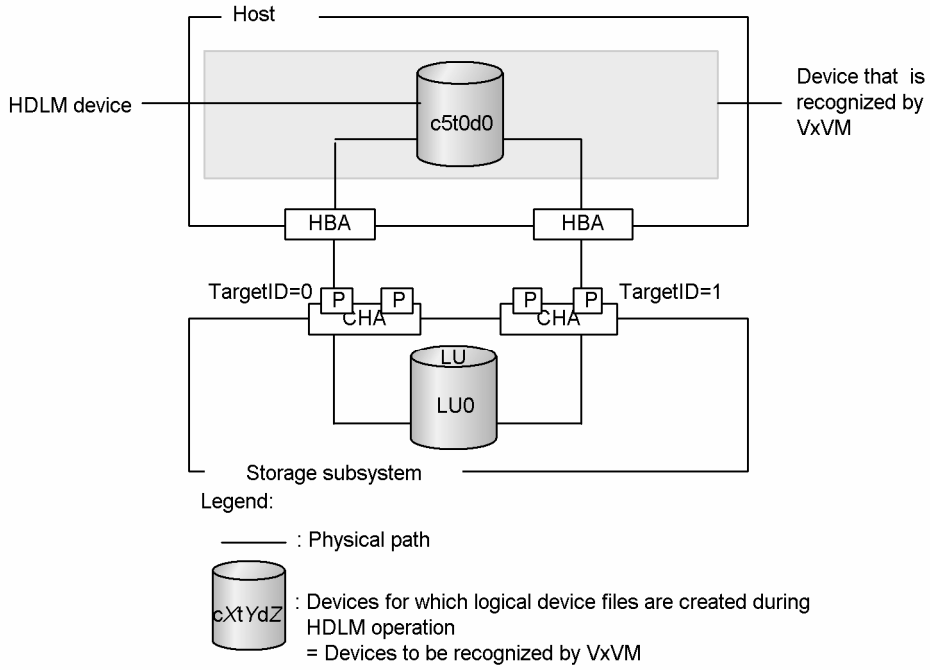

Figure 3.14 Devices Recognized by VxVM after Removal of Devices from VxVM

Only HDLM devices (c5t0d0) that are actually created are recognized. In this case, HDLM functions, such as load balancing and failover, work normally.

Note

After VxVM is installed, if the `vxdiskadm` utility has not been used to suppress a device or DMP from VxVM, the settings files such as `/etc/vx/vxvm.exclude` may be absent. If these settings files do not exist, you may be unable to suppress devices correctly.

The settings files necessary for device suppression are created by selecting menu task `17` (`Prevent multipathing/Suppress devices from VxVM's view`) or `18` (`Allow multipathing/Unsuppress devices from VxVM's view`) from the `vxdiskadm` main menu, and then by selecting menu task `8` (`List currently suppressed/non-multipathed devices`) from the `VolumeManager/Disk/ExcludeDevices` or `VolumeManager/Disk/IncludeDevices` menu.

3.12.3.2 Removing devices from VxVM on a controller basis

If physical paths connected to one LU go through different HBAs, devices are removed from VxVM on a controller or path basis. The procedure for removing devices from VxVM on a controller basis is described below by using the configuration example shown in Figure 3.15. This figure shows a configuration that is recognized by VxVM before devices are removed.

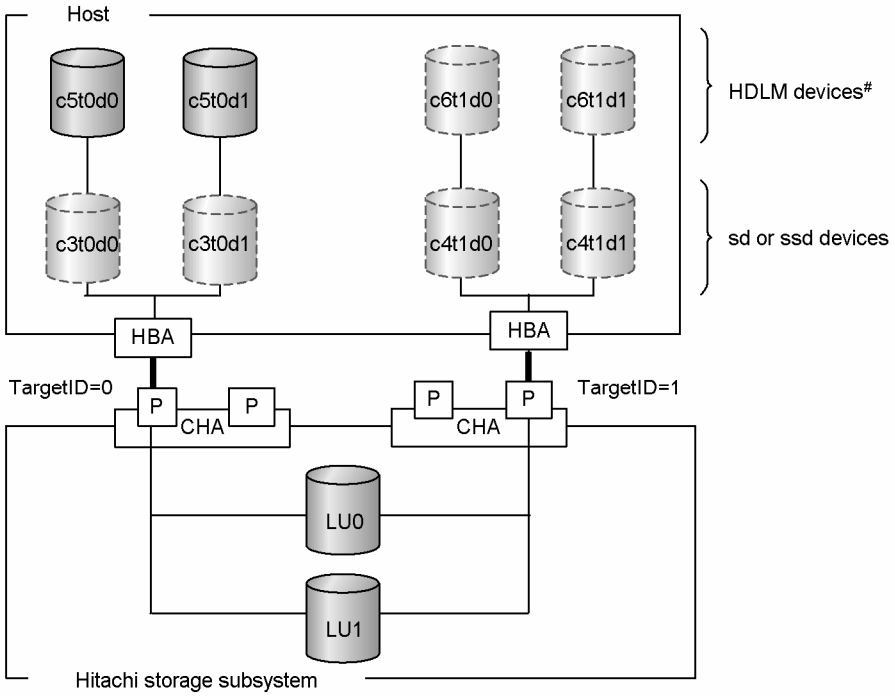

Legend:

: Fiber cable or SCSI cable

: Physical path

cXtYdZ : Devices for which logical device files are created during HDLM operation
=Devices to be recognized by VxVM

cXtYdZ : Devices for which logical device files are not created during HDLM operation
=Devices to be removed from VxVM

\# c6t1d0 and c6t1d1 will be recognized only when VxVM 3.1.1, VxVM 3.5, or VxVM 4.0 is used.
These are devices that VxVM virtually recognizes and are not actually created.

Figure 3.15 When HBA ports on the alternate path differ from those on the primary path connected to the Hitachi storage subsystem

To remove devices from VxVM, in the above configuration:

1. Use the `dlmvxexclude` utility to acquire the mapping between controller numbers and devices.

The devices to be removed from VxVM are output to the paths section in the execution result of the dlmvxexclude utility.

For details on the dlmvxexclude utility, see 7.6.

The following shows an example of executing the dlmvxexclude utility:

```
# /opt/DynamicLinkManager/bin/dlmvxexclude
exclude_all 0
paths
c3t0d0s2 /pci@1f,4000/pci@4/scsi@4/fp@0,0/sd@0,0
c3t0d1s2 /pci@1f,4000/pci@4/scsi@4/fp@0,0/sd@0,1
c4t1d0s2 /pci@1f,4000/pci@4/scsi@5/fp@0,0/sd@1,0
c4t1d1s2 /pci@1f,4000/pci@4/scsi@5/fp@0,0/sd@1,1
c6t1d0s2 /pseudo/dlmndrv@2/dlmfdrv@1,0
c6t1d1s2 /pseudo/dlmndrv@2/dlmfdrv@1,1
#
controllers
#
product
#
pathgroups
#
KAPL10704-I dlmvxexclude completed successfully.
```

The shaded portion, in the next step, is contained in the /etc/vx/vxvm.exclude file.

2. Make an entry for the controller of the device you wish to remove from VxVM in the controllers section of the /etc/vx/vxvm.exclude file.

   In the /etc/vx/vxvm.exclude file, write the shaded portion of the execution result of the dlmvxexclude utility shown in step 1. The settings will take effect when the host is restarted.

   Figure 3.16 shows an example of such a description in the /etc/vx/vxvm.exclude file.

```
exclude_all 0
paths
#
controllers
c3 /pci@1f,4000/pci@4/scsi@4/fp@0,0
c4 /pci@1f,4000/pci@4/scsi@5/fp@0,0
c6 /pseudo/dlmndrv@2
#
product
#
pathgroups
#
```

Figure 3.16 Example of a description in the /etc/vx/vxvm.exclude (controllers section)

The shaded portion indicates the specified entry.

3. Execute the following command to restart the host and apply the settings.

```
# reboot
```

4. Execute the `dlnkmgr` command's `view -path` operation to identify the HDLM devices that must be recognized by VxVM.

   The following shows an example of executing the command:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -path
Paths:000004 OnlinePaths:000004
PathStatus   IO-Count    IO-Errors
Online       72          0

PathID PathName                          DskName
iLU    ChaPort Status        Type  IO-Count   IO-Errors  DNum HDevName
000000 0000.0000.0000000000000000.0000  HITACHI  .OPEN-3
-SUN.15001
0654   1C      Online        Own         25           0     - c5t0d0
000001 0000.0000.0000000000000000.0001  HITACHI  .OPEN-3
-SUN.15001
0655   1C      Online        Own         25           0     - c5t0d1
000002 0001.0000.0000000000000001.0000  HITACHI  .OPEN-3
-SUN.15001
0654   4j      Online        Own         11           0     - c5t0d0
000003 0001.0000.0000000000000001.0001  HITACHI  .OPEN-3
-SUN.15001
0655   4j      Online        Own         11           0     - c5t0d1
KAPL01001-I The HDLM command completed normally. Operation name = view,
completion time = yyyy/mm/dd hh:mm:ss
#
```

   The devices displayed for `HDevName` (the shaded portion) are the HDLM devices that must be recognized by VxVM. In the above example, two devices, c5t0d0 and c5t0d1, must be recognized.

5. Execute the `vxdisk` command to list devices that are recognized by VxVM.

   The following shows an example of executing the command:

```
# vxdisk list
DEVICE       TYPE      DISK         GROUP        STATUS
c1t1d0s2     sliced     -            -           error
c1t2d0s2     sliced     rootdisk     rootdg      online
c5t0d0s2     sliced     -            -           online
c5t0d1s2     sliced     -            -           online
```

   If the HDLM devices output by the `dlnkmgr` command's `view -path` operation (the shaded portion) are displayed, the HDLM devices have been recognized by VxVM.

   If the HDLM devices are not displayed, perform the procedure described in 3.12.3.4 of 3.12.3.

6. Execute the `vxdmpadm` command on the HDLM devices listed by the `vxdisk` command. A list of devices that access the same LU is displayed. Execute the `vxdmpadm` command on each of the devices listed by the `vxdisk` command.

The following shows an example of executing the command:

```
# vxdmpadm getsubpaths dmpnodename=c5t0d0s2
NAME          STATE           PATH-TYPE  CTLR-NAME  ENCLR-TYPE   ENCLR-NAME
=============================================================================
c5t0d0s2      ENABLED         -          c5         HDS9980V     HDS9980V0
```

If any non-HDLM devices are `ENABLED`, perform the procedure described in 3.12.3.4 of 3.12.3 to set up `vxvm.exclude` again.

When the execution result of the `vxdmpadm` command is the status shown in Table 3.21, you do not need to set up `vxvm.exclude`.

Table 3.21 When setup of vxvm.exclude is unnecessary

Display of the vxdmpadm command	Why suppression is unnecessary
HDLM devices are not displayed.	They are non-HDLM-managed devices such as a built-in disk.
Only HDLM devices are displayed.	Among the devices corresponding to HDLM-managed devices, the non-HDLM devices have been successfully removed from VxVM.
Devices other than HDLM devices are `DISABLED`.	Among the devices corresponding to HDLM-managed devices, the non-HDLM devices have been successfully removed from VxVM (`DISABLED` is displayed because the host has not been restarted after setup).

7. If you have set up `vxvm.exclude` again, execute the following command to restart the host and apply the setting.

```
# reboot
```

If you did not set `vxvm.exclude`, go to step 8.

8. Use the `vxdiskadm` utility of VxVM to check the suppression settings.

Start the utility and select menu task `19` (`List currently suppressed/non-multipathed devices`) from the `vxdiskadm` main menu. A list of suppressed devices is displayed.

3.12.3.3 Removing devices from VxVM on a path basis

If all the physical paths to one LU go through the same HBA, devices are removed on a path basis. The procedure for removing devices on a path basis is described below by using the configuration example shown in Figure 3.17. This figure shows a configuration that is recognized by VxVM before devices are removed.

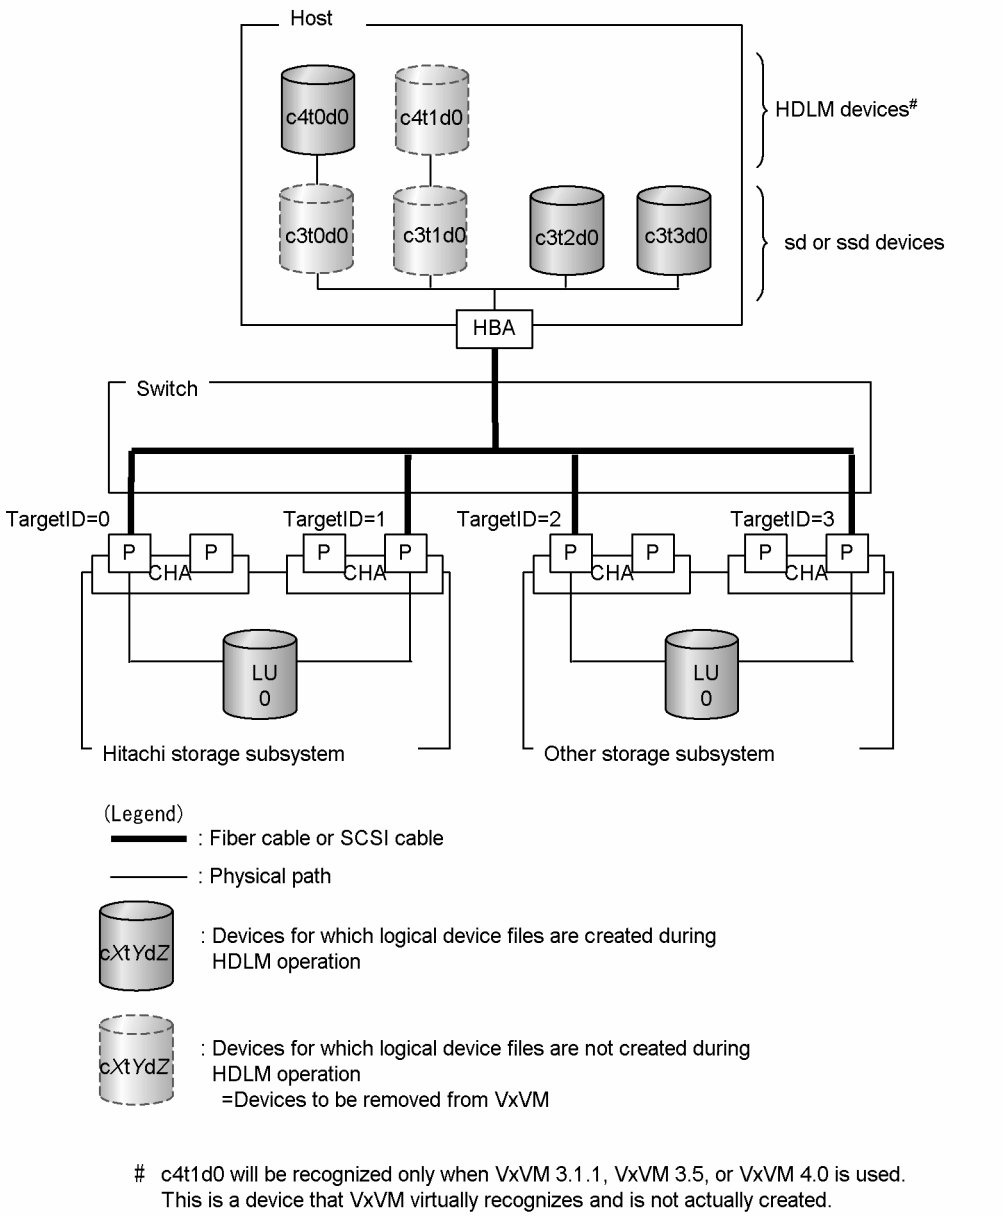

Figure 3.17 When the HBA Port of the Physical Path Connected to a Hitachi Storage Subsystem Is the Same As the One Connected to a Third-party Storage Subsystem

To remove devices from VxVM, in the above configuration:

1. Use the `dlmvxexclude` utility to acquire the mapping between controller numbers and devices.

   The devices to be removed from VxVM are output to the `paths` section in the execution result of the `dlmvxexclude` utility. For details about the `dlmvxexclude` utility, see 7.6.

   The following shows an example of executing the `dlmvxexclude` command:

```
# /opt/DynamicLinkManager/bin/dlmvxexclude
exclude_all 0
paths
c3t0d0s2 /pci@1f,4000/fibre-channel@2/sd@0,0
c3t1d0s2 /pci@1f,4000/fibre-channel@2/sd@1,0
c4t1d0s2 /pseudo/dlmndrv@1/dlmfdrv@1,0
#
controllers
#
product
#
pathgroups
#
KAPL10704-I dlmvxexclude completed successfully.
```

The shaded portion indicates the devices that need to be removed from VxVM.

2. Make an entry for the device you wish to removed from VxVM in the `paths` section of the `/etc/vx/vxvm.exclude` file.

   In the `/etc/vx/vxvm.exclude` file, write the shaded portion of the execution result of the `dlmvxexclude` utility shown in step 1. The settings will take effect when the host is restarted.

   Figure 3.18 shows an example of such a description in the `/etc/vx/vxvm.exclude` file.

```
exclude_all 0
paths
c3t0d0s2 /pci@1f,4000/fibre-channel@2/sd@0,0
c3t1d0s2 /pci@1f,4000/fibre-channel@2/sd@1,0
c4t1d0s2 /pseudo/dlmndrv@1/dlmfdrv@1,0
#
controllers
#
product
#
pathgroups
#
```

Figure 3.18 Example of a description in the /etc/vx/vxvm.exclude (paths section)

The shaded portion indicates the specified entry.

3. Execute the following command to restart the host and apply the settings.

```
# reboot
```

4. Execute the `dlnkmgr` command's `view -path` operation to identify the HDLM devices that must be recognized by VxVM.

   The following shows an example of executing the command:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -path
Paths:000008 OnlinePaths:000008
PathStatus   IO-Count    IO-Errors
Online       116         0

PathID PathName                          DskName
iLU    ChaPort Status     Type IO-Count    IO-Errors  DNum  HDevName
000000 0000.0000.0000000000000000.0000 HITACHI .OPEN-3     -SUN.15001
0654   1C      Online     Own        25            0     -  c4t0d0
000001 0001.0000.0000000000000001.0000 HITACHI .OPEN-3     -SUN.15001
0654   4j      Online     Own        11            0     -  c4t0d0
...
KAPL01001-I The HDLM command completed normally. Operation name =
view, completion time = yyyy/mm/dd hh:mm:ss
#
```

The devices displayed for `HDevName` (the shaded portion) are the HDLM devices that must be recognized by VxVM. In the above example, it is c4t0d0.

5. Execute the `vxdisk` command to list devices that are recognized by VxVM.

   The following shows an example of executing the command:

```
# vxdisk list
DEVICE       TYPE       DISK         GROUP        STATUS
c1t1d0s2     sliced     -             -           error
c1t2d0s2     sliced     rootdisk      rootdg      online
c4t0d0s2     sliced     -             -           online
...
#
```

   If the HDLM devices output by the `dlnkmgr` command's `view -path` operation (the shaded portion) are displayed, the HDLM devices have been recognized by VxVM.

   If the HDLM devices are not displayed, perform the procedure described in 3.12.3.4 of 3.12.3.

6. Execute the `vxdmpadm` command on the HDLM devices listed by the `vxdisk` command. A list of devices that access the same LU is displayed. Execute the `vxdmpadm` command on each of the devices listed by the `vxdisk` command.

   The following shows an example of executing the command:

```
# vxdmpadm getsubpaths dmpnodename=c4t0d0s2
NAME         STATE          PATH-TYPE  CTLR-NAME  ENCLR-TYPE   ENCLR-NAME
=============================================================================
c4t0d0s2     ENABLED          -         c4         HDS9980V     HDS9980V0
```

   If any non-HDLM devices are `ENABLED`, perform the procedure described in 3.12.3.4 of 3.12.3 to set up `vxvm.exclude` again.

   When the execution result of the `vxdmpadm` command is the status shown in Table 3.22, you do not need to set up `vxvm.exclude`.

Table 3.22 When setup of vxvm.exclude is unnecessary

Display of the vxdmpadm command	Why suppression is unnecessary
HDLM devices are not displayed.	They are non-HDLM-managed devices such as a built-in disk.
Only HDLM devices are displayed.	Among the devices corresponding to HDLM-managed devices, the non-HDLM devices have been successfully removed from VxVM.
Devices other than HDLM devices are `DISABLED`.	Among the devices corresponding to HDLM-managed devices, the non-HDLM devices have been successfully removed from VxVM (`DISABLED` is displayed because the host has not been restarted after setup).

7. If you have set up `vxvm.exclude` again, execute the following command to restart the host and apply the setting.

```
# reboot
```

   If you did not set `vxvm.exclude`, go to step 8.

8. Execute the `vxdiskadm` utility of the VxVM to check the suppression settings.

   Execute the utility and select menu task `19` (`List currently suppressed/non-multipathed devices`) from the `vxdiskadm` main menu. A list of suppressed devices is displayed.

3.12.3.4 Actions to be taken if an sd or ssd device has not been suppressed from VxVM

When you create a `vxvm.exclude` file using the output results of the `dlmvxexclude` utility with the following conditions, you might not be able to acquire the correct controller number of an sd or ssd device even after restarting the host. This can result in the sd or ssd device not being removed from VxVM normally.

- An sd or ssd device corresponding to an HDLM-managed device is connected to a non-Sun HBA.
- An HDLM device is running.

The following shows examples of when an sd or ssd device has not been correctly removed. The description below assumes that the shaded HDLM device in the execution results of the `view -path` operation has not been recognized.

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -path
Paths:000004 OnlinePaths:000004
PathStatus   IO-Count    IO-Errors
Online       72          0

PathID PathName                          DskName
iLU    ChaPort Status        Type  IO-Count   IO-Errors  DNum HDevName
000000 0000.0000.0000000000000000.0000  HITACHI  .OPEN-3
-SUN.15001
0654   1C      Online        Own         25           0     - c5t0d0
000001 0000.0000.0000000000000000.0001  HITACHI  .OPEN-3
-SUN.15001
0655   1C      Online        Own         25           0     - c5t0d1
000002 0001.0000.0000000000000001.0000  HITACHI  .OPEN-3
-SUN.15001
0654   4j      Online        Own         11           0     - c5t0d0
000003 0001.0000.0000000000000001.0001  HITACHI  .OPEN-3
-SUN.15001
0655   4j      Online        Own         11           0     - c5t0d1
KAPL01001-I The HDLM command completed normally. Operation name = view,
completion time = yyyy/mm/dd hh:mm:ss
#
```

Example 1:

The `vxdisk` command does not display HDLM devices but displays the sd or ssd devices corresponding to those HDLM devices.

```
# vxdisk list
DEVICE       TYPE      DISK         GROUP        STATUS
c1t1d0s2     sliced     -            -            error
c1t2d0s2     sliced     rootdisk     rootdg       online
c3t0d0s2     sliced     -            -            online
c3t0d1s2     sliced     -            -            online
...
#
```

The `vxdisk` command displays the SCSI device c3t0d0, which corresponds to an HDLM device, instead of the HDLM device c5t0d0.

Example 2:

The `vxdmpadm` command outputs a non-HDLM device, and the `STATE` column shows `ENABLED`.

```
# vxdisk list
DEVICE       TYPE      DISK         GROUP        STATUS
c1t1d0s2     sliced     -            -           error
c1t2d0s2     sliced     rootdisk     rootdg      online
c5t0d0s2     sliced     -            -           online
c3t0d1s2     sliced     -            -           online
# vxdmpadm getsubpaths dmpnodename=c5t0d0s2
NAME         STATE           PATH-TYPE  CTLR-NAME  ENCLR-TYPE   ENCLR-NAME
=============================================================================
c3t0d0s2     ENABLED          -          c3         HDS9980V      HDS9980V0
c4t1d0s2     ENABLED          -          c4         HDS9980V      HDS9980V0
c5t0d0s2     ENABLED          -          c5         HDS9980V      HDS9980V0
```

The `vxdmpadm` command displays c3t0d0 and c4t1d0 in addition to the HDLM device c5t0d0, and the `STATE` column shows `ENABLED`.

If sd or ssd devices have not been removed normally from VxVM, as shown in Examples 1 and 2, take the following action:

To suppress sd or ssd devices:

1. Execute the `vxdmpadm` command.

   In the following example, HDLM device c5t0d0 has been created and a non-HDLM device is `ENABLED`.

   As shown in Example 1, if the HDLM device (c5t0d0s2) is not displayed, in the `vxdmpadm` command specify the SCSI device (c3t0d0s2) that corresponds to the HDLM device.

```
# vxdmpadm getsubpaths dmpnodename=c5t0d0s2
NAME      STATE       PATH-TYPE  CTLR-NAME  ENCLR-TYPE  ENCLR-NAME
============================================================================
c3t0d0s2  ENABLED      -          c3         HDS9980V    HDS9980V0
c4t1d0s2  ENABLED      -          c4         HDS9980V    HDS9980V0
c5t0d0s2  ENABLED      -          c5         HDS9980V    HDS9980V0
```

2. Open the `/etc/vx/vxvm.exclude` file with a text editor such as vi.
3. Compare the logical device file names (`c`*U*`t`*X*`d`*Y*`s2` format) of the devices shown below to find the devices for which the `t`*X*`d`*Y* portions are the same.
   - Non-HDLM devices in the `ENABLED` state that are displayed by the `vxdmpadm` command
   - Devices in `/etc/vx/vxvm.exclude` that do not include `dlmndrv@`
4. For devices whose `t`*X*`d`*Y* portions are the same, change the controller number described in `/etc/vx/vxvm.exclude` to one of the devices displayed by the `vxdmpadm` command.

   Figure 3.19 shows the `/etc/vx/vxvm.exclude` file before editing. Figure 3.20 shows the `/etc/vx/vxvm.exclude` file after editing.

```
exclude_all 0
paths
c3t1d0s2 /pci@1f,0/pci@1/fibre-channel@3/sd@1,0
c3t1d1s2 /pci@1f,0/pci@1/fibre-channel@3/sd@1,1
c4t0d0s2 /pci@1f,0/pci@1/fibre-channel@4/sd@0,0
c4t0d1s2 /pci@1f,0/pci@1/fibre-channel@4/sd@0,1
c6t1d0s2 /pseudo/dlmndrv@2/dlmfdrv@1,0
c6t1d1s2 /pseudo/dlmndrv@2/dlmfdrv@1,1
#
controllers
#
product
#
pathgroups
#
```

Figure 3.19 The /etc/vx/vxvm.exclude file (before editing)

```
exclude_all 0
paths
c4t1d0s2 /pci@1f,0/pci@1/fibre-channel@3/sd@1,0
c3t1d1s2 /pci@1f,0/pci@1/fibre-channel@3/sd@1,1
c3t0d0s2 /pci@1f,0/pci@1/fibre-channel@4/sd@0,0
c4t0d1s2 /pci@1f,0/pci@1/fibre-channel@4/sd@0,1
c6t1d0s2 /pseudo/dlmndrv@2/dlmfdrv@1,0
c6t1d1s2 /pseudo/dlmndrv@2/dlmfdrv@1,1
#
controllers
#
product
#
pathgroups
#
```

Figure 3.20 The /etc/vx/vxvm.exclude file (after editing)

Among the non-HDLM devices in the ENABLED state, there might be two or more devices whose t*X*d*Y* portions in the logical device file name (c*U*t*X*d*Y*s2 format) are the same. In this case, change the currently assigned controller number in /etc/vx/vxvm.exclude.

5. After modifying all the sd or ssd device entries that have not been suppressed in /etc/vx/vxvm.exclude, execute the vxdctl command to restart the VxVM daemon.

   The following shows an example of executing the command:

```
# vxdctl enable
```

6. Execute the vxdmpadm command, specifying a device output by the vxdisk command, to check whether the device has been removed from VxVM.

   If the STATE column shows DISABLED, the device has been removed from VxVM. The following shows an example of executing the command:

```
# vxdmpadm getsubpaths dmpnodename=c3t0d0s2
NAME          STATE          PATH-TYPE  CTLR-NAME  ENCLR-TYPE   ENCLR-NAME
=============================================================================
c3t0d0s2      DISABLED       -          c3         HDS9980V     HDS9980V0
c4t1d0s2      DISABLED       -          c4         HDS9980V     HDS9980V0
c5t0d0s2      ENABLED        -          c5         HDS9980V     HDS9980V0
```

If there is more than one candidate for the controller number to be assigned in step 4, you might not be able to remove devices from VxVM in one operation. In this case, change the controller number you changed in step 4 to another controller number that was displayed using the `vxdmpadm` command, and then perform steps 5 and 6 again.

3.12.4 Introducing VxVM while using HDLM

To introduce VxVM in an environment in which HDLM is being used:

1. When using VxVM 3.1.1, VxVM 3.2, or VxVM 3.5, create `rootdg` on a disk not managed by HDLM.

   For details, refer to the VxVM documentation.

2. Create a disk group in accordance with the procedure described in 3.12.1.

   When using VxVM 4.0 or later, do not register the disk group containing the HDLM device as `bootdg`.

   For information on creating a disk group or a volume for HDLM devices, refer to the VxVM documentation.

3. To use Sun Cluster, perform step 5 in 3.12.5.

3.12.5 Linking VxVM and Sun Cluster

To create a file system on an HDLM device managed by VxVM and use it with Sun Cluster, perform the following procedure to create the file system:

1. On each node, execute the `vxdisk list` command to verify that VxVM recognizes the HDLM device.

2. On one of the nodes, execute the following command to check the device ID of the HDLM device:

   ```
   # /usr/cluster/bin/scdidadm -L
   ```

3. If a VxVM disk group has not been created, execute the `vxdiskadm` utility of VxVM on one of the nodes to create a disk group for the HDLM device.

   For details, see 3.12.1.

   Note

   Use the `scswitch` and `shutdown` commands to restart the node.

   The following shows an example of executing the command:

   ```
   # /usr/cluster/bin/scswitch -S -h phys-schost-1
   ```

```
# shutdown -g0 -y -i6
```

4. Verify that the VxVM disk group has been imported into only one of the nodes that constitute a cluster.

   If the VxVM disk group has been deported from all nodes, import the VxVM disk group to be registered in Sun Cluster into only one of the nodes.

5. If the VxVM disk group is not registered in Sun Cluster, execute the `scsetup` utility from the node into which you imported the VxVM disk group to be registered in Sun Cluster, and then register the VxVM disk group into a Sun Cluster disk device group.

   For details, refer to the explanation of system administration in the Sun Cluster documentation.

   The following shows the procedure:

   - Start the `scsetup` utility.

     ```
     # /usr/cluster/bin/scsetup
     ```

     The `Main Menu` appears.

   - Enter 4 (Device group and volumes) to perform operations on the VxVM disk device group.

     The `Device Group Menu` appears.

   - Enter 1 (Register a VxVM disk group as a device group) to register the VxVM disk device group.

     Following the prompts, specify the VxVM disk group name to be registered as a Sun Cluster disk device group, the priority level of the master node, and whether or not to perform switchback.

   Note

   After you enter this information, the following message may be output to syslog:

   ```
   KAPL05023-E Could not process the
   IOCTL(00004D03).(00000000:00000000)
   ```

   Even if this message is output to syslog, no problems exist operationally, provided that the Status of the device group newly registered in the following step is `Online`.

   - Confirm that the disk device group is registered, and that its Status is `Online`.

     If the disk device group was registered appropriately, use of the following command will display information about the new disk device group:

     ```
     # /usr/cluster/bin/scstat -D
     ```

   Note

   After registering the VxVM disk group into a Sun Cluster disk device group, do not use the VxVM command to import or deport that VxVM disk group.

6. If the VxVM volume has not been created, create a VxVM volume in the disk group from the master node (the node that currently acts as master for the disk device group). The following shows an example of executing the command:

```
# vxassist -g dlmdg make vol-dg 10g
```

7. To create a volume or otherwise modify the VxVM disk group after it has been registered into a Sun Cluster disk device group, execute the `scsetup` utility to register the change in the VxVM disk group configuration into Sun Cluster.

   This ensures that the Sun Cluster global device namespace status is correct.

   For details, refer to the explanation of system administration in the Sun Cluster documentation.

   The following shows the procedure:

   - Start the `scsetup` utility.

     ```
     # /usr/cluster/bin/scsetup
     ```

     The `Main Menu` appears.
   - Enter 4 (Device group and volumes) to perform operations on the VxVM disk device group.
   - Enter 2 (Synchronize volume information for a VxVM device group) to register the change in the configuration.

     Following the prompt, enter the name of the VxVM device group whose configuration was changed.

8. If the file system has not been created, execute the `newfs` command to create the file system.

   The following shows an example of executing the command:

```
# newfs /dev/vx/rdsk/dlmdg/vol-dg
```

3.13 Setting Up SDS

If SDS is installed on an HDLM host, SDS uses the logical device file name of the HDLM device to access the HDLM-managed device. Therefore, after you install HDLM, you need to register the logical device file name of the HDLM device. When uninstalling HDLM, unregister the logical device file name of the HDLM device and set the logical device file name of the sd or ssd device.

When using SDS, keep the following in mind.

3.13.1 Notes

- Even if you use a local diskset or shared diskset, do not specify the `-r` option when changing the path configuration and executing the `dlmsetconf` utility for creating the HDLM driver configuration definition file. When you specify the `-r` option, the minor number of the HDLM device used on the SDS will be changed, and you will no longer be able to use the metadevice.

- Even if you use a local diskset or shared diskset, you cannot use the DiskSuite tool (*metatool*). Use the command line utility instead. Also, when you set up a metadevice, do not use the `-a` option of the `metainit` command. Register the database on the command line, and then register the metadevice.
- Specify the same name for the logical device file name (`c`*U*`t`*X*`d`*Y*) of the HDLM device on all the hosts that belong to a single shared diskset.
- When a shared diskset is used, specify the same major number and minor number for the physical device files of the HDLM devices on all hosts that belong to that shared diskset.

  Executing the following command enables you to check the major number and minor number for the physical device file of the HDLM device.

```
# ls -lR /devices/pseudo | grep dlmfdrv
```

  To specify the same major number and minor number for the physical device files of the HDLM device on each host using the shared diskset:

1. View the `/etc/name_to_major` file on each host, and then select a major number that is not used on any host.

   Edit the `/etc/name_to_major` file to set the major number of the HDLM filter driver (`dlmfdrv`) to the number you chose above.

   Restart all the hosts, specifying either of the following commands:

   `# reboot -- -r`

   or

   `boot -r` (at the `ok` prompt)

For more information on SDS settings, refer to the SDS documentation.

3.13.2 Registering HDLM devices

If both HDLM and SDS are used on the same host, perform the following settings after installing HDLM.

3.13.2.1 To use a local metadevice

To use a local metadevice:

2. Create a replica of the state database.

   Create three replicas of the local state database on the disk (slice). Create these replicas on a disk that is not managed by HDLM, not shared, and has not been connected to a Fibre Channel.

   The following shows an example of executing the command:

```
host1# metadb -a -f c1t1d0s1 c1t1d0s3 c1t1d0s7
```

3. Create a local metadevice.

Use the HDLM device (slice) to create a local metadevice (logical disk).

The types of metadevices include coupled, stripe, mirrored, RAID5, and transmeta devices.

The following shows an example of executing the command to create local metadevice (d0) of type stripe (1) comprised of two HDLM devices (slices).

```
host1# metainit d0 1 2 c5t1d0s7 c5t2d1s7
```

4. Create a file system on the local metadevice.

   The following shows an example of executing the command:

```
host1# newfs /dev/md/rdsk/d0
```

5. Mount the created file system to the mount point.

   The following shows an example of executing the command:

```
host1# mkdir -p /mnt/filesystem1
host1# mount /dev/md/dsk/d0 /mnt/filesystem1
```

3.13.2.2 To use a shared diskset

For disks registered with the SDS shared diskset, SDS automatically allocates the minimal number of state databases that store the configuration and status of shared metadevices created on a shared diskset. When using Sun Cluster, you use the device ID in the device name of the disk.

Note

When you register a disk into a shared diskset, the existing partitions may be reconfigured. For details, refer to the SDS documentation.

To use a shared diskset:

1. Create three replicas of the local state database on the disk (slice).

   Create these replicas on all hosts that use the shared diskset. Make sure that all the following conditions are satisfied for the creation-target disk:

   - The disk is not managed by HDLM.
   - The disk is not shared.
   - The disk is not connected via Fibre Channel.

   The following shows an example of executing the command:

```
# metadb -a -f c1t1d0s1 c1t1d0s3 c1t1d0s7
```

2. Create a shared diskset.

   The following shows examples of executing the commands.

```
host1# metaset -s diskset1 -a -h host1 host2
```

3. Register a disk into the shared diskset.

The following shows examples of executing the commands:

When not using Sun Cluster:

```
host1# metaset -s diskset1 -a c5t1d2 c5t1d3
```

When using Sun Cluster:

```
host1# metaset -s diskset1 -a /dev/did/dsk/d10 /dev/did/dsk/d11
```

4. Create a shared metadevice in the shared diskset.

   The following shows examples of executing the commands:

   When not using Sun Cluster (stripe type):

   ```
   host1# metainit -s diskset1 d1 1 2 c5t1d2s0 c5t1d3s0
   ```

   When using Sun Cluster (stripe type):

   ```
   host1# metainit -s diskset1 d1 1 2 /dev/did/dsk/d10s0 /dev/did/dsk/d11s0
   ```

5. Create a file system on the shared metadevice of the shared diskset.

   The following shows an example of executing the commands:

   ```
   host1# newfs /dev/md/diskset1/rdsk/d1
   ```

6. Mount the created file system to the mount point.

   The following shows examples of executing the commands. One is for a normal mount, when Sun Cluster is not being used, and the other is for a global mount, when Sun Cluster is being used.

   When not using Sun Cluster:

   ```
   host1# mkdir /diskset1
   host1# mount /dev/md/diskset1/dsk/d1 /diskset1
   ```

   When using Sun Cluster:

   ```
   host1# mkdir -p /global/diskset1/d1
   host2# mkdir -p /global/diskset1/d1
   host1# mount -g /dev/md/diskset1/dsk/d1 /global/diskset1/d1
   ```

3.14 Setting Up SVM

If SVM is installed on an HDLM host, SVM uses the logical device file name of the HDLM device to access the HDLM-managed device. Therefore, after installing HDLM, you need to register the logical device file name of the HDLM device in SVM. When uninstalling HDLM, unregister the logical device file name of the HDLM device and set the logical device file name of the sd or ssd device.

3.14.1 Notes

When using SVM, keep the following in mind.

- Even if you use a local diskset or shared diskset, do not specify the `-r` option when changing the path configuration and executing the `dlmsetconf` utility for creating the HDLM driver configuration definition file. When you specify the `-r` option, the minor number of the HDLM device used on the SVM will be changed, and you will no longer be able to use the volume.
- Even if you use a local diskset or shared diskset, you cannot use the Solaris management console. Use the command line utility instead. Also, when you set up a volume, do not use the `-a` option of the `metainit` command. Register the database on the command line, and then register the volume.
- Specify the same name for the logical device file name (`c`*U*`t`*X*`d`*Y*) of the HDLM device on all the hosts that belong to a single shared diskset.
- When a shared diskset is used, specify the same major number and minor number for the physical device files of the HDLM devices on all the hosts that belong to that shared diskset.

  Executing the following command enables you to check the major number and minor number for the physical device file of the HDLM device.

```
# ls -lR /devices/pseudo | grep dlmfdrv
```

  To specify the same major number and minor number for the physical device files of the HDLM device on each host using the shared diskset:

1. View the `/etc/name_to_major` file on each host, and then select a major number that is not used on any host.

   Edit the `/etc/name_to_major` file to set the major number of the HDLM filter driver (`dlmfdrv`) to the number you chose above.

   Restart all the hosts, specifying either of the following commands:

   ```
   # reboot -- -r
   ```

   or

   `boot -r` (at the `ok` prompt)

For more information on SVM settings, refer to the SVM documentation.

3.14.2 Registering HDLM devices

If SVM is used, perform the following settings after installing HDLM, and register the SVM volume.

3.14.2.1 To use a local volume

To use a local volume:

1. Create a replica of the state database.

Create three replicas of the local state database on the disk (slice). Create these replicas on a disk that is not managed by HDLM, not shared, and has not been connected to a Fibre Channel.

The following shows an example of executing the command:

```
host1# metadb -a -f c1t1d0s1 c1t1d0s3 c1t1d0s7
```

2. Create a local volume.

   Use the HDLM device (slice) to create a local volume (logical disk).

   The types of volumes include coupled, stripe, mirrored, and RAID5.

   The following shows an example of executing the command to create a local volume (`d0`) of type stripe (`1`) comprised of two HDLM devices (slices).

```
host1# metainit d0 1 2 c5t1d0s7 c5t2d1s7
```

3. Create a file system on the local volume.

   The following shows an example of executing the command:

```
host1# newfs /dev/md/rdsk/d0
```

4. Mount the created file system to the mount point.

   The following shows an example of executing the commands:

```
host1# mkdir -p /mnt/filesystem1
host1# mount /dev/md/dsk/d0 /mnt/filesystem1
```

3.14.2.2 To use a shared diskset

For disks registered with the shared diskset, SVM automatically allocates the minimal number of state databases that store the configuration and status of shared volumes created on a shared diskset. When using Sun Cluster, you use the device ID in the name of the device.

Note

When you register a disk into a shared diskset, existing partitions may be reconfigured. For detail, refer to the SVM documentation.

To use a shared diskset:

a) Create three replicas of the local state database on the disk (slice).

Create these replicas on all hosts that use the shared diskset. Make sure that all the following conditions are satisfied for the creation-target disk:

- The disk is not managed by HDLM.
- The disk is not shared.
- The disk is not connected via Fibre Channel.

The following shows an example of executing the command:

```
# metadb -a -f c1t1d0s1 c1t1d0s3 c1t1d0s7
```

b) Create a shared diskset.

The following shows an example of executing the commands:

```
host1# metaset -s diskset1 -a -h host1 host2
```

c) Register a disk into the shared diskset.

The following shows examples of executing the commands.

When not using Sun Cluster:

```
host1# metaset -s diskset1 -a c5t1d2 c5t1d3
```

When using Sun Cluster:

```
host1# metaset -s diskset1 -a /dev/did/dsk/d10 /dev/did/dsk/d11
```

d) Create a shared volume in the shared diskset.

The following shows examples of executing the commands:

When not using Sun Cluster (stripe type):

```
host1# metainit -s diskset1 d1 1 2 c5t1d2s0 c5t1d3s0
```

When using Sun Cluster (stripe type):

```
host1# metainit -s diskset1 d1 1 2 /dev/did/dsk/d10s0 /dev/did/dsk/d11s0
```

e) Create a file system on the shared volume of the shared diskset.

The following shows an example of executing the commands:

```
host1# newfs /dev/md/diskset1/rdsk/d1
```

f) Mount the created file system to the mount point.

The following shows examples of executing the commands. One is for a normal mount, when Sun Cluster is not being used, and the other is for a global mount, when Sun Cluster is being used.

When not using Sun Cluster:

```
host1# mkdir /diskset1
host1# mount /dev/md/diskset1/dsk/d1 /diskset1
```

When using Sun Cluster:

```
host1# mkdir -p /global/diskset1/d1
host2# mkdir -p /global/diskset1/d1
host1# mount -g /dev/md/diskset1/dsk/d1 /global/diskset1/d1
```

3.15 Setting Up VCS

When VCS is installed on an HDLM host, you can register HDLM devices, or VxVM disk groups containing HDLM devices, as resource groups.

For details on registering resource groups, refer to the VCS documentation.

Note

The I/O fencing function of VCS version 4.0 or later is not supported.

3.16 Uninstalling HDLM

This section explains how to return the HDLM environment to the way it was before HDLM was installed, and describes each step of the process.

3.16.1 Overview of HDLM uninstallation

Return the HDLM environment to the way it was before HDLM was installed by following the flow shown below in Figure 3.21.

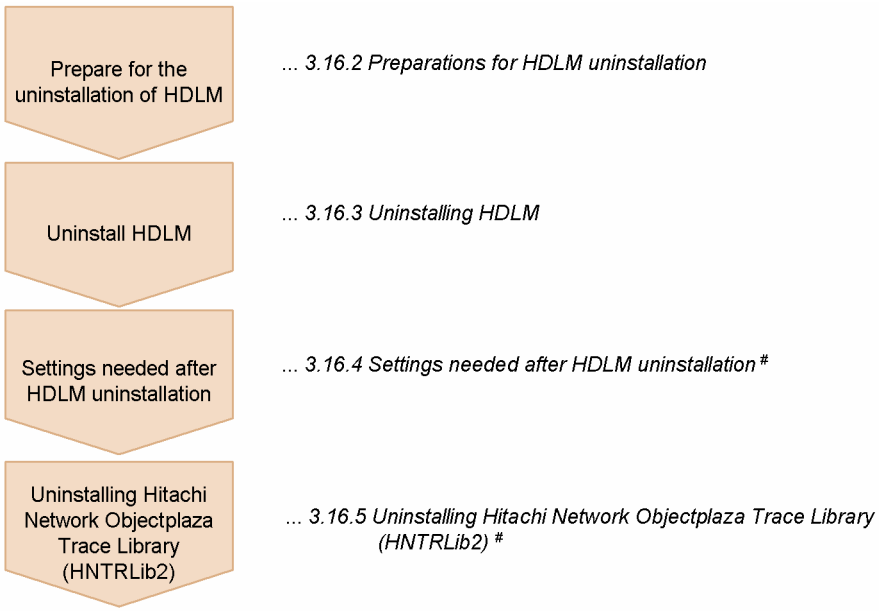

Figure 3.21 Overview of HDLM uninstallation

#

Perform as required.

3.16.2 Preparations for HDLM uninstallation

3.16.2.1 Performing operations on HDLM-managed devices

To perform operations on HDLM-managed disks:

1. If necessary, back up the contents of the HDLM-managed devices to tape or some other storage medium.

2. For the HDLM devices, collect all relevant information on the correspondence between the logical device file names and the internal LUNs of the LUs in the storage subsystem. This information will enable you to keep the correspondence with the logical device file names after HDLM has been uninstalled.

   You can collect this correspondence information by outputting the results of the `dlnkmgr` command's `view -path` operation to a file. The following shows an example of executing the command:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -path > /var/tmp/old_view_path_list.txt
```

   Note

   Do not specify any of the directories shown below as the output destination of the results of the `dlnkmgr` command's `view -path` operation. If you specify one of these directories, the file will be removed when you uninstall HDLM.

   - `/opt/DynamicLinkManager`
   - `/opt/hitachi`
   - `/var/opt/DynamicLinkManager`
   - `/var/opt/hitachi`

3. Back up the logical device files.

   When you uninstall HDLM, logical device files used by HDLM will be deleted from the `/dev/dsk` and `/dev/rdsk` directories.

   Back up all information pertaining to the logical device files within the `/dev/dsk` and `/dev/rdsk` directories to another directory.

   The following shows an example of executing the commands:

```
# mkdir /var/tmp/bkfile2
# ls -l /dev/dsk > /var/tmp/bkfile2/devices1
# ls -l /dev/rdsk > /var/tmp/bkfile2/devices2
```

4. Unregister the HDLM-managed devices.

   If HDLM-managed devices are registered with applications, including cluster software programs (other than Volume Manager), unregister those devices. Since HDLM-managed devices are accessed by using the logical device file names of the HDLM devices, these applications will no longer be able to access those HDLM-managed devices by using the names currently used.

5. Unmount the disks.

   If any HDLM-managed devices are mounted, unmount them. Also, if the disks are set to be mounted automatically when the host starts, delete this setting in the `/etc/vfstab` file.

6. Unregister HDLM-managed devices from the volume management software.

   If any HDLM-managed devices are registered in the volume management software, unregister them.

3.16.2.2 Remove Sun Cluster settings

If you are using Sun Cluster, before you uninstall HDLM, perform the following to remove the disk device group settings and the Quorum device settings.

If you are using SDS or SVM, delete the disk device group to which the disksets that use HDLM belong to unregister them. If you are using SDS, see 3.16.2.5 of 3.16.2. If you are using SVM, see 3.16.2.6 of 3.16.2. Unregistering a diskset deletes data in the diskset. Make sure that you first back up the data. If you are using VxVM, delete the disk device group, delete the disk device group that contains the disk group that uses HDLM to unregister the disk group.

The following procedure shows how to unregister a disk group.

To unregister a disk group:

1. Place the disk device group offline.

```
# /usr/cluster/bin/scswitch -F -D disk-device-group-name
```

2. Start the `scsetup` utility.

```
# /usr/cluster/bin/scsetup
```

3. Enter `4` (*Device group and volumes*) to work with VxVM disk device groups.

   `Device Group Menu` is displayed.

4. To unregister a VxVM disk device group, enter `3` (*Unregister a VxVM device group*).

   As the utility instructs, enter the name of the VxVM disk device group you want to unregister.

If the HDLM device has been configured as a Quorum device, perform the following operations to remove the Quorum device settings.

1. Check the device ID of the Quorum device.

   The following shows an example of executing the command:

```
# /usr/cluster/bin/scconf -pv | grep Quorum
# /usr/cluster/bin/scdidadm -L
```

2. Since the device ID may have changed after HDLM uninstallation, on each node check the correspondence of the device with the internal LUNs of the LUs in the storage subsystem.

   The following shows an example of executing the command:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -path
```

3. When two or more Quorum devices exist, execute the following command until the number of Quorum devices becomes one. When the number of Quorum devices finally becomes one, go to step 4.

```
# /usr/cluster/bin/scconf -r -q globaldev=Quorum-device's-device-ID
```

4. Switch to become the superuser in the node that is to be changed to the maintenance mode in order to release the settings of the last Quorum device.

5. Back up all resource groups and HDLM devices from the node that are to be changed to the maintenance mode.

   The following shows an example of executing the command:

```
host1# /usr/cluster/bin/scswitch -S -h host1
```

6. Shut down the node that is to be changed to the maintenance mode so that the `ok` prompt appears.

   The following shows an example of executing the command:

```
host1# shutdown -g0 -y -i0
```

7. At another node in the same cluster, switch to become the superuser, and then place the node that was stopped in step 6 into the maintenance mode.

   The following shows an example of executing the command:

```
host2# /usr/cluster/bin/scconf -c -q node=host1,maintstate
```

8. Confirm that the cluster node enters the maintenance mode.

   Make sure that the node that is in the maintenance mode has the following status:

   - The node is in the Offline status.
   - Both the current `Quorum votes present` and `Quorum votes possible` are 0.

   The following shows an example of executing the command:

```
host2# /usr/cluster/bin/scstat -q
```

9. If the cluster is a two -node configuration, enter the install mode.

   If the cluster consists of three or more nodes, this step is not necessary so go to step 10.
   The following shows an example of executing the command:

```
host2# /usr/cluster/bin/scconf -c -q installmode
```

10. Release the settings of the last Quorum device.

    The following shows an example of executing the command:

```
host2# /usr/cluster/bin/scconf -r -q globaldev=Quorum-device's-device-ID
```

11. At a node that is not in the maintenance mode, switch to become the superuser.

    If the cluster consists of two nodes, go to step 13. If the cluster consists of three or more nodes, go to step 12.

12. If the cluster consists of three or more nodes, reset the cluster Quorum votes.

    The following shows an example of executing the command:

```
host2# /usr/cluster/bin/scconf -c -q node=host1,reset
```

13. Change the mode of the node, which you changed into the maintenance mode, back to the previous mode.

    At the `ok` prompt, execute `boot`.

14. If the cluster consists of two nodes, make sure that the node (which was changed into the maintenance mode) that you want to change back to the previous mode is active, and then reset the cluster Quorum votes. If the cluster consists of three or more nodes, this step is not necessary.

    The following shows an example of executing the command:

```
host2# /usr/cluster/bin/scconf -c -q node=host1,reset
```

3.16.2.3 Remove VCS settings

If you are using VCS, before uninstalling HDLM, unregister the HDLM devices and VxVM disk groups containing HDLM devices, from the resource groups. For details on how to remove the VCS settings, refer to the VCS documentation.

3.16.2.4 Remove VxVM settings

If you are using both HDLM and VxVM on the same host, before you uninstall HDLM, perform the following to delete the HDLM devices from VxVM.

Note

If you are using Sun Cluster, before deleting HDLM devices from VxVM, remove the settings for the disk device group (see 3.16.2.2 of 3.16.2).

- For HDLM version 04-01 or earlier:

  a) Execute the following command to unmount the file systems on which the VxVM volumes are mounted:

  ```
  # umount mount-point
  ```

  Execute the following command to delete all VxVM volumes:

  ```
  # vxedit -r -f rm volume-name
  ```

  Execute the following command to delete all disk groups:

  ```
  # vxdg destroy disk-group-name
  ```

  Execute the following command to delete all HDLM devices:

  ```
  # vxdisk rm chWtXdYsZ
  ```

- For HDLM version 05-00 or later

  b) Execute the following command to deactivate all volumes used in VxVM:

  ```
  # vxvol -g disk-group-name stopall
  ```

  The deactivated volumes become unavailable.

When there is already a created disk group, execute the following command to deport the disk group:

```
# vxdg deport disk-group-name
```

To enable VxVM management, delete the entries for sd or ssd devices and the entries for HDLM devices from the `/etc/vx/vxvm.exclude` file.

After uninstalling HDLM, the VxVM volumes that were used with HDLM devices can be used with sd or ssd devices. For details, see 3.16.4.2 of 3.16.4.

3.16.2.5 Remove SDS settings

If you are using SDS, before you uninstall HDLM, perform the following to delete the registration of HDLM devices from SDS.

Note

Deleting a metadevice deletes data in the metadevice. Make sure that you first back up the data.

- When using a local metadevice

  c) Execute the following command to unmount the file system

  The following shows an example of executing the command:

  ```
  host1# umount /mnt/filesystem1
  ```

  Delete the local metadevice that is using the HDLM device.

  The following shows an example of executing the command to delete local metadevice `d0`.

  ```
  host1# metaclear d0
  ```

- When using a shared diskset

  d) Delete the shared metadevice.

  The following shows an example of executing the command.

  When not using Sun Cluster:

  ```
  host1# umount /diskset1
  host1# metaclear -s diskset1 d1
  ```

  When using Sun Cluster:

  ```
  host1# umount /global/diskset1/d1
  ```

```
host1# metaclear -s diskset1 d1
```

Delete the disk from the shared diskset.

The following shows an example of executing the command.

When not using Sun Cluster:

```
host1# metaset -s diskset1 -d -f c5t1d2 c5t1d3
```

When using Sun Cluster:

```
host1# metaset -s diskset1 -d -f /dev/did/dsk/d10 /dev/did/dsk/d11
```

Delete the shared diskset.

The following shows an example of executing the command.

```
host1# metaset -s diskset1 -d -f -h host1 host2
```

3.16.2.6 Remove SVM settings

If you are using SVM, before you uninstall HDLM, perform the following to delete the registration of HDLM devices from SVM.

Note

Deleting a volume deletes data in the volume. Make sure that you backup the data.

- When using a local volume

e) Execute the following command to unmount the file system:

The following shows an example of executing the command:

```
host1# umount /mnt/filesystem1
```

Delete the SVM local volume that is using the HDLM device.

The following shows an example of executing the command to delete SVM local volume `d0`.

```
host1# metaclear d0
```

- When using a shared diskset

f) Delete the shared volume.

The following shows an example of executing the command.

When not using Sun Cluster:

```
host1# umount /diskset1
```

```
host1# metaclear -s diskset1 d1
```

When using Sun Cluster:

```
host1# umount /global/diskset1/d1
host1# metaclear -s diskset1 d1
```

Delete the disk from the shared diskset.

The following shows an example of executing the command.

When not using Sun Cluster:

```
host1# metaset -s diskset1 -d -f c5t1d2 c5t1d3
```

When using Sun Cluster:

```
host1# metaset -s diskset1 -d -f /dev/did/dsk/d10 /dev/did/dsk/d11
```

Delete the shared diskset.

The following shows an example of executing the command.

```
host1# metaset -s diskset1 -d -f -h host1 host2
```

3.16.3 Uninstalling HDLM

3.16.3.1 Uninstalling HDLM from the local boot disk environment

When you uninstall HDLM, if the KAPL09019-E or KAPL09020-E message is output, follow the procedure described directions in 3.16.5 to uninstall HNTRLib2. However, if the KAPL09026-I message is output, since a program other than HDLM is using Hitachi Network Objectplaza Trace Library (HNTRLib2), only HDLM will be uninstalled.

Note

When uninstalling HDLM on a host where a Device Manager Agent 5.0 or later is installed, do not execute any of the following commands of Device Manager Agent during uninstallation. Also, do not uninstall HDLM while executing any of the following Device Manager Agent commands:

```
hbsasrv, HiScan, hdvmagt_account, hdvmagt_schedule, hldutil, TIC
```

To uninstall HDLM:

1. Log in to Solaris as a user with root privileges.
2. Stop all services, DBMSs, and processes that are using HDLM-managed paths.
3. Execute the following command to unmount all HDLM devices:

```
# umount mount-point
```

4. If automatic mounting is enabled, disable it by removing the definitions of logical device files for the HDLM devices (c*U*t*X*d*Y*s*Z*) from the /etc/vfstab file.

5. Execute the following command to uninstall HDLM:

```
# pkgrm DLManager
```

6. Type y in response to the following message:

```
The following package is currently installed:
   DLManager  Dynamic Link Manager
              (sparc) xx-xx

Do you want to remove this package? [y,n,?,q]
```

7. Type y in response to the following message:

```
This package contains scripts which will be executed with super-user
permission during the process of removing this package.

Do you want to continue with the removal of this package [y,n,?,q]
```

   Uninstalling will start.

8. When you are uninstalling HDLM version 04-00 or earlier, manually uninstall Hitachi Network Objectplaza Trace Library (HNTRLib).

   When HDLM is the only program that uses HNTRLib, uninstall HNTRLib as directed in 3.16.6.

9. Restart the host, with the reconfiguration option specified.

   When Sun Cluster is not being used:

   Execute one of the following commands:

   - # reboot -- -r
   - boot -r (at the ok prompt)

   When Sun Cluster is being used:

   Execute the following command on one of the nodes:

```
# /usr/cluster/bin/scshutdown -g0 -y
```

   Then, execute the following command at the ok prompt on all nodes:

```
boot -r
```

   Note

   While a node is starting, the message below may be output to the console or syslog:

```
Could not stat: ../../devices/pseudo/dlmndrv@1/dlmfdrv@0,0:c,raw
path not loaded.
```

This message is output because HDLM is uninstalled and the HDLM devices are deleted. If you do not want this message to be output next time the node starts, update the device ID information by referring to steps 1 to 4 in 3.16.4.1 of 3.16.4. For details, refer to the Sun Cluster documentation.

3.16.3.2 Uninstalling HDLM from the boot disk environment

1. Migrate to the local boot disk environment, following the procedure described in *Migrating from a Boot Disk Environment to the Local Boot Disk Environment*.
2. Uninstall HDLM, following the procedure described in *Uninstalling HDLM from the local boot disk environment*.

3.16.4 Settings needed after HDLM uninstallation

3.16.4.1 Sun Cluster settings

After restarting the host, perform steps 1 to 4 on all nodes.

When setting Quorum devices, perform step 5.

When setting a disk device group, perform step 6.

1. Remove the device ID information of devices that do not exist. The following shows an example of executing the command:

```
# /usr/cluster/bin/scdidadm -C
```

2. Assign device IDs. The following shows an example of executing the command:

```
# /usr/cluster/bin/scdidadm -r
```

3. Notify the server of the new device ID information. The following shows an example of executing the command:

```
# /usr/cluster/bin/scdidadm -ui
```

4. Update the Sun Cluster global device namespace. The following shows an example of executing the command:

```
# /usr/cluster/bin/scgdevs
```

5. Set up a Quorum device at any of the nodes that constitute the cluster.

   Check the device ID of the sd or ssd device, which is used for the Quorum device. The following shows an example of executing the command:

```
# /usr/cluster/bin/scdidadm -L
```

Set up a Quorum device. The following shows an example of executing the command:

```
# /usr/cluster/bin/scconf -a -q globaldev=d4
```

In a two-node configuration, execute the following command to release the cluster installation mode:

```
# /usr/cluster/bin/scconf -c -q reset
```

6. When you are using volume management software, specify the settings for the disk device group.

   If you are using SDS or SVM, you must re-create the diskset. For details on how to re-create a diskset, refer to the Sun Cluster documentation.

   If you are using VxVM, you can use a VxVM disk group that was used with HDLM. To use such a VxVM disk group, re-register it in Sun Cluster's disk device group. To re-register it, import the VxVM disk group to any node in the cluster, and then use the `scsetup` utility to register the Sun Cluster disk device group (see 3.16.4.2 of 3.16.4). For details, refer to the Sun Cluster documentation.

3.16.4.2 VxVM settings

In HDLM 05-00 or later, VxVM disk groups that were used with HDLM devices can still be used after uninstallation. To use such disk groups, perform the following procedure after uninstalling HDLM.

To use VxVM's disk groups that were used with HDLM devices in HDLM 05-00 or later:

1. Execute the following command to import the disk group:

   ```
   # vxdg import disk-group-name
   ```

2. To use a volume that belongs to a disk group, execute the following command to activate the volume:

   ```
   # vxvol -g disk-group-name startall
   ```

3.16.4.3 SDS settings

Create meta devices. For details, refer to the SDS documentation.

3.16.4.4 SVM settings

Create volumes. For details, refer to the SVM documentation.

3.16.5 Uninstalling Hitachi Network Objectplaza Trace Library (HNTRLib2)

When you uninstall HDLM, if the KAPL09019-E or KAPL09020-E message is output, follow the procedure below to uninstall HNTRLib2.

To uninstall HNTRLib2:

1. Log in to Solaris as a user with root privileges.
2. Execute the following command to unregister the name of the bundled program products:

```
# /opt/hitachi/HNTRLib2/etc/hntr2cancel "Hitachi Dynamic Link Manager"
```

3. Execute the following command:

```
# /opt/hitachi/HNTRLib2/etc/hntr2setup
```

   The HNTRLib2 setup menu will be displayed.

4. From the setup menu, select `9`.

   HNTRLib2 will be uninstalled.

   If HNTRLib2 is not being used by any other programs:

   HNTRLib2 will be uninstalled normally, and the following message be displayed.

   ```
   Unsetup is complete.
   ```

   If HNTRLib2 is being used by another program:

   HNTRLib2 will not be uninstalled, and the following message will appear.

   ```
   Because a bundle PP name is registered,
   I did not do the Uninstall.
   ```

   If HNTRLib2 was not uninstalled, execute the following command to check what programs are using it:

   ```
   # /opt/hitachi/HNTRLib2/etc/hntr2getname
   ```

   If you are unable to complete uninstallation even though no programs other than HDLM are using HNTRLib2, contact your HDLM vendor or maintenance company if there is a maintenance contract for HDLM.

Note

If the log output directory set in HNTRLib2 was not the default directory, the log files will not be deleted during uninstallation. In this case delete these files after uninstallation.

3.16.6 Uninstalling Hitachi Network Objectplaza Trace Library (HNTRLib)

After you uninstall HDLM version 04-00 or earlier, check the manuals and *Release Notes* that came with the software installed on the host. If no applications other than HDLM are using HNTRLib, uninstall it by following the procedure below:

To uninstall HNTRLib:

1. Make sure that no applications other than HDLM are using HNTRLib.

   See the manuals and documentation for each program to check whether the program is using Hitachi Network Objectplaza Trace Library (HNTRLib).

2. Log in to Solaris as a user with root privileges.

3. Execute the following command:

```
# /opt/hitachi/HNTRLib/etc/hntrsetup
```

   The HNTRLib setup menu will be displayed.

4. From the setup menu, select `9`.

   HNTRLib will be uninstalled.

5. Delete the HNTRLib common library files and the directory in which they are stored.

   Execute the commands shown below:

   - To delete the shared library files (symbolic links):

```
# rm /opt/hitachi/common/lib/libhntrs*
# rm /opt/hitachi/common/lib/libhntrt*
```

   - To delete the directory that contains the shared library files:

```
# rm -r /opt/hitachi/common/lib/D001
```

   If the `/opt/hitachi` directory contains only the files and directories shown above, delete the `/opt/hitachi` directory.

Note

- If the log output directory set in HNTRLib was not the default directory, the log files will not be deleted during uninstallation. In this case delete these files after uninstallation.
- Even if you attempt to uninstall HNTRLib2 when HNTRLib is installed, HNTRLib will not be uninstalled. If no other programs are using HNTRLib, delete it manually.

Chapter 4 HDLM Operation

This chapter describes operating procedures for HDLM, including how to operate HDLM and the HDLM manager, and how to change the configuration of the operating environment.

- 4.1 Notes on using HDLM
- 4.2 HDLM operations using commands
- 4.3 Starting and stopping the HDLM manager
- 4.4 Changing the configuration of the HDLM operating environment

4.1 Notes on using HDLM

This section explains the notes on using HDLM and using the environment in which HDLM is installed. Make sure that you read this section.

4.1.1 Displaying path information

- The path IDs displayed by using the `dlnkmgr` command's `view` operation differ depending on the order in which paths were detected when the host started. For this reason, you should always use the path name to specify a path.
- When displaying the path status for a new LU that was dynamically added by executing the `dlmcfgmgr` utility for managing the HDLM configuration, a disconnected path might be displayed as `Online(E)` even though another online path is connected to the LU.

  If this situation occurs, perform the online operation for the path displayed as `Online(E)`, and then update the path status.

4.1.2 When a path error is detected

When a path error is detected by HDLM, you must immediately resolve the error and restore the path.

A check for path errors is performed whenever an I/O is issued. If there are any paths through which I/O is not normally issued, such as a non-owner path, you should enable path health checking in order to detect errors even when there is no I/O. For details on path health checking, see 2.10.

4.1.3 iostat command

The `iostat` command lists devices that are not represented by sd or ssd device logical device file names, but by sd or ssd device types and instance numbers.

By executing the `dlnkmgr` command's `view -drv` operation, you can obtain the mapping between the logical device file names of HDLM devices and the combinations of sd or ssd device types and instance numbers listed by the `iostat` command.

Shown below are examples of output from the `iostat` command and the `dlnkmgr` command's `view -drv` operation. It is assumed that one LU and two paths exist. In the lists, `c0t6d0` and `c1t12d0` represent a CD-ROM drive and a built-in disk, respectively.

```
# iostat -xtcn  interval 2
    tty         cpu
 tin tout  us  sy wt  id
  0   1    0   0  1   98
                    extended device statistics
  r/s   w/s   kr/s   kw/s  wait  actv  wsvc_t  asvc_t  %w   %b   device
  0.0   0.0   0.0    0.0   0.0   0.0   0.0     0.0     0    0   c0t6d0
  3.3   4.9   28.3   68.1  0.0   0.0   0.0     4.3     0    2   c1t12d0
  0.0   0.0   0.0    0.0   0.0   0.0   0.0     0.0     0    0   ssd2
  0.0   0.0   0.0    0.0   0.0   0.0   0.0     0.0     0    0   ssd3
  0.0   0.0   0.0    0.0   0.0   0.0   0.0     0.0     0    0   hostname:vold(pid466)

# /opt/DynamicLinkManager/bin/dlnkmgr view -drv
PathID HDevName                 Device  LDEV
000000 c5t50060E80033A9912d95 ssd2    9970/9980.15001.065F
000001 c5t50060E80033A9912d95 ssd3    9970/9980.15001.065F
KAPL01001-I The HDLM command completed normally. Operation name = view, completion
time = yyyy/mm/dd hh:mm:ss
#
```

The shaded portion indicates mapping from combinations of sd or ssd device types and instance numbers to HDLM device logical device file names.

4.1.4 Storage subsystem

- Start up the storage subsystem before you start up the hosts, so that Solaris can detect the storage subsystem.
- Do not change the vendor ID or product ID of a storage subsystem. If you change either of them, HDLM will not be able to recognize the storage subsystem.

4.1.5 Command execution

- When a path error occurs while a file system is being created or formatted or the `fsck` command is being executed, the operation might not have been completed. If so, recover from the path error, and then perform the operation again.

4.1.6 Using Sun HBA

If all of the following conditions are satisfied, the KAPL08019-E and KAPL08022-E messages indicating a path error may be output to `/var/adm/messages` during boot processing (even when the path is normal):

- Solaris 8 or Solaris 9 is used.
- A Sun HBA that the ssd driver recognizes is used.
- The host and storage subsystem are connected via Fabric.

If this situation occurs, execute the `dlnkmgr view -path` command to check the path status after the completion of booting. There are no problems if the path is in `Online` status.

4.1.7 Starting the host, system, and HDLM in a Solaris 10 environment

- When starting the host

  If the host is started in single-user mode, the displayed path status indicates an error even when the actual path status is normal. Therefore, before you use HDLM devices in single-user mode, execute the procedure described in *(1) Initializing HDLM when the host is started in single-user mode*in 4.1.9 *Operation in single-user mode*.

4.1.8 Upgrading Solaris

In a local boot disk environment:

When upgrading Solaris, first uninstall HDLM, and then upgrade Solaris. After the Solaris upgrade has finished, re-install a version of HDLM appropriate to the new version of Solaris.

In a boot disk environment

You cannot upgrade Solaris in a boot disk environment. If you need to upgrade Solaris, follow these steps:

Migrate to the local boot disk environment, following the procedure described in 3.16.3.2, and then uninstall HDLM.

Upgrade Solaris.

Reconfigure the boot disk environment, following the procedure described in 3.6.

4.1.9 Operation in single-user mode

The Solaris-related, HDLM-related, and application-related tasks listed in 4.1.9.2 can still be performed when the host is started in single-user mode, and even if the host is subsequently changed to run level 1. Before you perform any of these tasks, execute the procedure described in 4.1.9.1.

4.1.9.1 Initializing HDLM when the host is started in single-user mode

If the host is started in single-user mode, initialize HDLM by following the procedure below.

To initialize HDLM:

1. Mount the file system for the tasks to be performed.

   To perform HDLM-related tasks, mount the `/var`, `/opt`, and `/tmp` directories.

2. Execute the following HDLM command to place the paths online:

```
# /opt/DynamicLinkManager/bin/dlnkmgr online -s
```

Even if the paths are physically connected, this command may fail if executed immediately after the boot processing completes. In this case, wait at least 10 seconds and then re-execute the command. If still unsuccessful, take the action indicated in the displayed message.

3. Execute the following command to set the HDLM logical device file name in the HDLM filter driver:

```
# /opt/DynamicLinkManager/bin/dlmsetdevname
```

4. Execute the following command to start the HDLM manager:

```
# /etc/init.d/DLMManager start
```

4.1.9.2 Tasks that can be performed in single-user mode

By executing the procedure described in 4.1.9.1, you can perform the following tasks in single-user mode while applications are inactive.

For Solaris

You can perform the following tasks on the server on which HDLM is installed, in the same manner as on a server without HDLM installed:

- Apply patches
- Add or delete packages
- Change basic settings in Solaris, etc.

For HDLM

- Installing HDLM
- Configuring a boot disk environment
- Setting up HDLM functions
- Setting up an integrated trace#
- Creating a file system
- Executing an HDLM command
- Using an HDLM device
- Changing the configuration of the HDLM operating environment
- Applying an HDLM service pack

#

In single-user mode, HDLM command activity is not logged to integrated trace files.

For applications

- Install an application
- Upgrade an application
- Set or change the configuration

4.1.10 Maintenance tasks on devices connected by paths in the boot disk environment

Note the following points when performing maintenance on devices connected to the boot disk in a boot disk environment running Solaris:

- Ensure that at least one path connected to the boot disk is online when you perform a maintenance task.
- If the maintenance task entails disconnecting a path, change the status of that path to `Offline(C)` by executing the `offline` operation of the `dlnkmgr` command before you begin. If the path is not in `Offline(C)` status, an error will be detected when the path is disconnected during maintenance, and access to the boot disk will be temporarily disrupted, with a consequent deterioration in system performance. When you finish the maintenance work, immediately place the path online again by executing the `online` operation of the `dlnkmgr` command.

4.2 HDLM operations using commands

This section explains how to use the HDLM command operations to utilize HDLM functions. For details on the command operations, see Chapter 6.

4.2.1 Notes on using commands

- Execute the command as a user with root privileges.
- To specify a value containing a space in a parameter, enclose the entire value in double quotation marks (").

4.2.2 Viewing path information

This subsection explains how to display path information by using the `dlnkmgr` command.

To display path information, execute the `dlnkmgr` command's `view` operation with the `-path` parameter specified. The following example shows how to execute the command:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -path
```

To display information only for the paths accessing the specified host device, execute the `dlnkmgr` command's `view` operation with the `-path` and `-hdev` parameters specified. The following example shows how to execute the command:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -path  -hdev c6t5d0
Paths:000002 OnlinePaths:000002
PathStatus   IO-Count    IO-Errors
Online       15          0

PathID PathName                          DskName
iLU              ChaPort Status     Type IO-Count IO-Errors DNum
HDevName
000000 0000.0000.0000000000000005.0000 HITACHI .OPEN-3          .15001
0670             1H      Online     Own        12         0    -
c6t5d0
000001 0001.0000.0000000000000006.0000 HITACHI .OPEN-3          .15001
0670             2H      Online     Own         3         0    -
c6t5d0
KAPL01001-I The HDLM command completed normally. Operation name = view,
completion time = yyyy/mm/dd hh:mm:ss
#
```

For details on the displayed items and their descriptions, see 6.7.

4.2.3 Changing the status of paths

This subsection explains how to change the status of paths.

4.2.3.1 When changing the status of a path to online

To change the status of a path to online:

1. Check the current status of the path.

   To change the status of the path for each HBA port, channel adapter port, or path to online, check the path name or AutoPATH_ID.

   The following shows an example of executing the command:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -path
```

2. To change the status of a path to online, execute the `dlnkmgr` command's `online` operation.

   You can specify the paths you want to place online by specifying an HBA port or channel adapter port to which the target paths are connected, or by specifying a single path. For details on how to specify a path, see 6.5.

   For example, if you want to place all paths that pass through a specific HBA port online, execute the `dlnkmgr` command's `online` operation with the `-hba` parameter specified. The following shows an example of executing the command:

```
# /opt/DynamicLinkManager/bin/dlnkmgr online -hba 0001.0000
KAPL01057-I All the paths which pass the specified HBA will be changed to the Online status.
Is this OK? [y/n]:y
KAPL01061-I 3 path(s) were successfully placed online. 0 path(s) could not be placed online.
Operation name = online
#
```

3. Check the changed status of the path.

The following shows an example of executing the command:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -path
```

4.2.3.2 When changing the status of a path to Offline(C)

To change the status of a path to Offline(C):

1. Check the current status of the path.

   To change the status of the path for each HBA port, channel adapter port, or path to Offline(C), check the path name or AutoPATH_ID.

   The following shows an example of executing the command:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -path
```

2. To change the status of the path to Offline(C), execute the `dlnkmgr` command's `offline` operation.

   You can specify the paths you want to place online by specifying an HBA port or channel adapter port to which the target paths are connected, or by specifying a single path. For details on how to specify a path, see 6.4.

   For example, if you want to place all paths that pass through a specific HBA port offline, execute the `dlnkmgr` command's `offline` operation with the `-hba` parameter specified. The following shows an example of executing the command:

```
# /opt/DynamicLinkManager/bin/dlnkmgr offline -hba 0001.0000
KAPL01055-I All the paths which pass the specified HBA port will be changed to the Offline(C)
status. Is this OK? [y/n]:y
KAPL01056-I If you are sure that there would be no problem when all the paths which pass
the specified HBA are placed in the Offline(C) status, enter y.  Otherwise, enter n. [y/n]:y
KAPL01061-I 3 path(s) were successfully placed offline(C). 0 path(s) could not be placed
offline(C). Operation name = offline
#
```

3. Check the changed status of the path.

   The following shows an example of executing the command:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -path
```

4.2.4 Viewing LU information

This subsection explains how to display LU information by using the `dlnkmgr` command.

To display LU information, execute the `dlnkmgr` command's `view` operation with the `-lu` parameter specified. The following shows an example of executing the command:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -lu
Product      : 9500V
SerialNumber  : 0051
LUs          : 3
```

```
iLU  HDevName Device PathID Status
0001 c6t2d0   sd61   000003 Online
              sd64   000009 Online
0002 c6t2d1   sd62   000004 Online
              sd65   000010 Online
0003 c6t2d2   sd63   000005 Online
              sd66   000011 Online

Product       : 9970/9980
SerialNumber  : 15001
LUs           : 3

iLU  HDevName Device PathID Status
0670 c6t5d0   sd74   000000 Online
              sd80   000006 Offline(E)
0671 c6t5d1   sd75   000001 Online
              sd81   000007 Offline(E)
0672 c6t5d2   sd76   000002 Online
              sd82   000008 Offline(C)
KAPL01001-I The HDLM command completed normally. Operation name = view, completion time =
yyyy/mm/dd hh:mm:ss
#
```

For details on the displayed items and their descriptions, see 6.7.

4.2.5 Displaying corresponding information about an HDLM device, sd or ssd device, and LDEV

This subsection explains how to display corresponding information about an HDLM device, sd or ssd device, and LDEV by using the `dlnkmgr` command.

You can display corresponding information about an HDLM device, sd or ssd device, and LDEV by executing the `dlnkmgr` command's `view` operation with the `-drv` parameter specified. For details on the `view` operation, see 6.7 The execution result of this operation is displayed on a single line for each path.

The following shows an example of executing the `view` operation:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -drv
PathID HDevName Device LDEV
000000 c6t5d0   sd74   9970/9980.15001.0670
000001 c6t5d1   sd75   9970/9980.15001.0671
000002 c6t5d2   sd76   9970/9980.15001.0672
000003 c6t5d3   sd77   9970/9980.15001.0673
000004 c6t5d4   sd78   9970/9980.15001.0674
000005 c6t5d5   sd79   9970/9980.15001.0675
000006 c6t5d0   sd80   9970/9980.15001.0670
000007 c6t5d1   sd81   9970/9980.15001.0671
000008 c6t5d2   sd82   9970/9980.15001.0672
000009 c6t5d3   sd83   9970/9980.15001.0673
000010 c6t5d4   sd84   9970/9980.15001.0674
000011 c6t5d5   sd85   9970/9980.15001.0675
KAPL01001-I The HDLM command completed normally. Operation name = view, completion time =
yyyy/mm/dd hh:mm:ss
#
```

For details on the displayed items and their descriptions, see 6.7.

4.2.6 Initializing statistical information for paths

This subsection explains how to initialize statistical information (I/O counts and I/O errors) for all paths managed by HDLM.

This procedure is useful when you wish to check the number of I/O hits and I/O errors that have occurred since the last time the I/O counts and I/O errors were initialized to the value `0`.

To initialize statistical information for paths:

1. Check the current status of the path.

   The following shows an example of executing the command:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -path
```

2. To initialize statistical information (I/O counts and I/O errors) for all paths managed by HDLM, execute the `dlnkmgr` command's `clear` operation with the `-pdst` parameter specified.

   The following shows an example of executing the command:

```
# /opt/DynamicLinkManager/bin/dlnkmgr clear -pdst
KAPL01049-I Would you like to execute the operation? Operation name = clear [y/n]:y
KAPL01001-I The HDLM command completed normally. Operation name = clear, completion time
= yyyy/mm/dd hh:mm:ss
#
```

3. Check whether the statistical information for the paths has been initialized.

   The following shows an example of executing the command:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -path
```

4.2.7 Viewing and setting up the operating environment

This subsection explains how to display and set up the HDLM operating environment.

4.2.7.1 Viewing the operating environment

To display the operating environment, execute the `dlnkmgr` command's `view` operation with the `-sys` and `-sfunc` parameters specified.

The following shows an example of executing the command:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -sys -sfunc
HDLM Version                : xx-xx
Service Pack Version        :
Load Balance                : on(rr)
Support Cluster             :
Elog Level                  : 3
Elog File Size(KB)          : 9900
Number Of Elog Files        : 2
```

```
Trace Level                    : 0
Trace File Size(KB)            : 1000
Number Of Trace Files          : 4
Path Health Checking           : on(30)
Auto Failback                  : off
Intermittent Error Monitor     : off
KAPL01001-I The HDLM command completed normally. Operation name = view, completion time =
yyyy/mm/dd hh:mm:ss
#
```

For details on the displayed items and their descriptions, see 6.7.

4.2.7.2 Setting up the operating environment

To set up the HDLM operating environment, execute the `dlnkmgr` command's `set` operation. This operation allows you to set the following functions:

- Load balancing
- Path health checking
- Automatic failback
- Intermittent error monitoring
- Error log collection level
- Trace level
- Error log file size
- The number of error log files
- Trace file size
- The number of trace files

For details on how to set each function, see 6.6.

For example, to set up the error log collection level, execute the `dlnkmgr` command's `set` operation with the `-ellv` parameter specified. When the confirmation message is displayed, enter `y` to execute, or `n` to cancel the command.

The following shows an example of executing the command:

```
# /opt/DynamicLinkManager/bin/dlnkmgr set -ellv 1
KAPL01049-I Would you like to execute the operation? Operation name = set [y/n]: y
KAPL01001-I The HDLM command completed normally. Operation name = set, completion time =
yyyy/mm/dd hh:mm:ss
#
```

To check whether the settings have been applied, perform the procedure described in see 4.2.7.1 from 4.2.7.

4.2.8 Viewing license information

This subsection explains how to display license information.

To display license information, execute the `dlnkmgr` command's `view` operation with the `-sys` and `-lic` parameters specified.

The following shows an example of executing the command:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -sys -lic
License Type Expiration
Permanent    -
KAPL01001-I The HDLM command completed normally. Operation name = view, completion time =
yyyy/mm/dd hh:mm:ss
#
```

For details on the displayed items and their descriptions, see 6.7.

4.2.9 Updating the license

This subsection explains how to update the license.

To update the license, execute the `dlnkmgr` command's `set` operation with the `-lic` parameter specified. When the confirmation message is displayed, enter `y` to execute, or `n` to cancel the command. If the license key file does not exist, a message asking you to enter the license key appears, so enter the license key.

Note

When you use the `dlnkmgr` command's `set` operation with the `-lic` parameter to install a license, you can only execute one command at a time. If you attempt to execute more than the `dlnkmgr` command containing the `set` operation with the `-lic` parameter more than once, a core file is created and the following message may appear:

`KAPL01075-E A fatal error occurred in HDLM. The system environment is invalid.`

If this message appears, execute the `dlnkmgr` command's `view` operation with the `-sys -lic` parameter to make sure that the license is installed correctly.

The following shows an example of executing the command:

```
# /opt/DynamicLinkManager/bin/dlnkmgr set -lic
KAPL01049-I Would you like to execute the operation? Operation name = set [y/n]: y
KAPL01071-I A permanent license was installed.
KAPL01001-I The HDLM command completed normally. Operation name = set, completion time =
yyyy/mm/dd hh:mm:ss
#
```

4.2.10 Viewing HDLM version information

This subsection explains how to display HDLM version information.

To display HDLM version information, execute the `dlnkmgr` command's `view` operation with the `-sys` parameter specified. The following shows an example of executing the command:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -sys
HDLM Version                : xx-xx
Service Pack Version        :
Load Balance                : on(rr)
Support Cluster             :
Elog Level                  : 3
Elog File Size (KB)         : 9900
Number Of Elog Files        : 2
Trace Level                 : 0
Trace File Size(KB)         : 1000
Number Of Trace Files       : 4
Path Health Checking        : on(30)
Auto Failback               : off
Intermittent Error Monitor  : off
HDLM Manager Ver      WakeupTime
Alive       xx-xx    yyyy/mm/dd hh:mm:ss
HDLM Alert Driver Ver      WakeupTime          ElogMem Size
Alive            xx-xx    yyyy/mm/dd hh:mm:ss
HDLM Driver Ver      WakeupTime
Alive       xx-xx    yyyy/mm/dd hh:mm:ss
License Type Expiration
Permanent    -
KAPL01001-I The HDLM command completed normally. Operation name = view, completion time =
yyyy/mm/dd hh:mm:ss
#
```

The value displayed in `HDLM version` indicates the HDLM version.

4.2.11 Viewing HDLM component information

This subsection explains how to display HDLM component information.

To display HDLM component information, execute the `dlnkmgr` command's `view` operation with the `-sys` parameter specified. The following shows an example of executing the command is executed:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -sys
HDLM Version                : xx-xx
Service Pack Version        :
Load Balance                : on(rr)
Support Cluster             :
Elog Level                  : 3
Elog File Size (KB)         : 9900
Number Of Elog Files        : 2
Trace Level                 : 0
Trace File Size(KB)         : 1000
Number Of Trace Files       : 4
Path Health Checking        : on(30)
Auto Failback               : off
Intermittent Error Monitor  : off
HDLM Manager Ver      WakeupTime
Alive       xx-xx    yyyy/mm/dd hh:mm:ss
HDLM Alert Driver Ver      WakeupTime          ElogMem Size
Alive            xx-xx    yyyy/mm/dd hh:mm:ss
HDLM Driver Ver      WakeupTime
Alive       xx-xx    yyyy/mm/dd hh:mm:ss
License Type Expiration
Permanent    -
```

```
KAPL01001-I The HDLM command completed normally. Operation name = view, completion time =
yyyy/mm/dd hh:mm:ss
#
```

Among the displayed items, HDLM Manager, HDLM Alert Driver, and HDLM Driver indicate the HDLM component information.

Also, you can view information for each HDLM component. Execute the dlnkmgr command's view operation with the -sys and subsequent parameter specified. The following shows an example of executing the command:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -sys -msrv
# /opt/DynamicLinkManager/bin/dlnkmgr view -sys -adrv
# /opt/DynamicLinkManager/bin/dlnkmgr view -sys -pdrv
```

4.3 Starting and stopping the HDLM manager

If an error occurs in the system, such as in an HDLM program, you may need to manually stop or start the HDLM manager to recover from the error.

4.3.1 Starting the HDLM manager

To start the HDLM manager, log in to Solaris as a user with root privileges and then execute the following command.

```
# /etc/init.d/DLMManager start
```

Lower-case characters can be used for the command name (dlmmanager).

```
# /etc/init.d/dlmmanager start
```

The startup script that was set up during HDLM installation runs, and starts the HDLM manager.

Use the following HDLM command's view operation to confirm that the HDLM manager is running:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -sys -msrv
HDLM Manager Ver      WakeupTime
Alive        xx-xx    yyyy/mm/dd hh:mm:ss
KAPL01001-I The HDLM command completed normally. Operation name = view, completion time =
yyyy/mm/dd hh:mm:ss
#
```

When the HDLM Manager column shows Alive, the HDLM manager is active.

4.3.2 Stopping the HDLM manager

To stop the HDLM manager, log in to Solaris as a user with root privileges, and then execute the following command:

```
# /etc/init.d/DLMManager stop
```

Lower-case characters can be used for the command name (dlmmanager).

```
# /etc/init.d/dlmmanager stop
```

The stop script that was set up during HDLM installation runs, and stops the HDLM manager.

Execute the following dlnkmgr command's view operation to confirm that the HDLM manager has stopped.

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -sys -msrv
HDLM Manager Ver      WakeupTime
Dead
KAPL01001-I The HDLM command completed normally. Operation name = view, completion time =
yyyy/mm/dd hh:mm:ss
#
```

When the HDLM Manager column shows Dead, the HDLM manager is inactive.

4.4 Changing the configuration of the HDLM operating environment

This section describes how to reconfigure the environment for running HDLM. Whenever you add a new logical unit (LU), or specify whether or not an LU is managed by HDLM, you must modify the configuration file of HDLM.

4.4.1 Precautions regarding changes to the configuration of an HDLM operating environment

Keep the following in mind when changing the configuration of an HDLM operating environment.

4.4.1.1 Changing the configuration of a system that uses HDLM

After changing the configuration of a system that uses HDLM, perform these steps:

1. Delete all unnecessary logical device files in the /dev/dsk and /dev/rdsk directories.
2. Back up the HDLM driver configuration definition file (/kernel/drv/dlmfdrv.conf).
3. Execute the dlmsetconf utility for creating the HDLM driver configuration definition file.

   Edit the HDLM driver unconfiguration definition file (/opt/DynamicLinkManager/config/dlmfdrv.unconf) as required, and then execute the dlmsetconf utility.

   For details on how to edit the HDLM driver unconfiguration definition file (/opt/DynamicLinkManager/config/dlmfdrv.unconf), see 3.5.4.
4. Depending on the changes you have made to the HDLM system configuration, restart the host with the reconfiguration option specified, or execute the dlmcfgmgr utility for managing the HDLM configuration.

4.4.1.2 Switching the kernel mode

If you switch the kernel mode from 32 bits to 64 bits when only the 32-bit version of HDLM is installed, you must re-install HDLM after the switch is completed.

For details on how to switch the kernel mode currently being used, see 3.5.1.5.

4.4.1.3 When the path configuration is changed

When the path configuration of an LU is changed and any of the following operations occur, logical device files of HDLM devices might be re-created with a different name:

- The `-r` parameter was specified when executing the `dlmsetconf` utility for creating the HDLM driver configuration definition file.
- One of the following changes was made to the construction of the path used by HDLM to determine the logical device file name of the HDLM device:
  - A path was deleted.
  - The target ID was changed. Alternatively, when Sun HBA had been used with fabric connection (using Sun StorageTek SAN Foundation Software), the connection destination was changed to a different storage port.
  - A LUN was changed.
  - The HBA model was changed.
  - The HBA mounting slot was changed.
  - The contents of the `/etc/path_to_inst` file were altered to change the instance number of the HBA driver.
  - The FCode was installed on an HBA where the FCode had not been previously installed. Alternatively, the FCode was uninstalled from an HBA.

If one of these operations occurs, you must change the logical device file setting for the HDLM devices that are set in the application program. If the logical device file of an HDLM device is re-created with a different name, the logical device file before re-creation might remain in the `/dev/dsk` and `/dev/rdsk` directories. In this case, you must delete the logical device file before re-creation. Delete the logical device file before its re-creation, and then follow the instructions in 4.4.1.1 in 4.4.1

The logical device file names of HDLM devices on each LU are determined by using the path (physical device file name corresponding to the sd or ssd device) that can be obtained by using the following procedure:

To obtain the path to an LU:

1. Obtain the instance number of the HDLM driver corresponding to the LU.

   Execute the following command to display the minor number of the HDLM device of the desired LU (HDLM device on slice `0`). The value obtained by dividing the displayed minor number by `8` is the instance number of HDLM.

```
# /usr/bin/ls -lL HDLM-logical-device-file-name-on-slice-0
```

```
```

The following shows an example of executing the command.

This example obtains the instance number of the HDLM driver of the LU whose logical device file name is `/dev/dsk/c4t21d0s0`. In this case, since the minor number is `8`, the instance number is `1`.

```
# /usr/bin/ls -lL /dev/dsk/c4t21d0s0
brw-r-----   1 root     sys      248,  8 Jan 24 21:03 /dev/dsk/c4t21d0s0
#
```

2. Obtain path information from the instance number of the HDLM driver.

   Execute the following command to display the driver configuration information, and then refer to the instance field of the HDLM driver obtained in step 1 (the driver name of the HDLM driver is `dlmfdrv`). Identify the path used to determine the HDLM logical device file name of the LU from the `FilterPath` property.

```
# /usr/sbin/prtconf -vP
```

   The following shows an example of executing the command.

   This example shows the information on instance 1 of the HDLM driver (the driver name of the HDLM driver is `dlmfdrv`), which is contained in the output results of the `prtconf` command.

   The shaded portion is the physical device file name of the sd or ssd driver of the path used to determine the HDLM logical device file name.

```
# /usr/sbin/prtconf -vP
        ...
    dlmfdrv, instance #1
        System properties:
            name <inq-83h> length <511>
                ...
            name <FilterPath> length <38>
                value '/pci@1f,4000/pci@2/QLGC,qla@4/sd@15,0'
                ...
#
```

4.4.1.4 When the path configuration is change in a boot disk environment

If you change the path used by HDLM for determining the name of a logical device file of an HDLM device, perform the steps below. If you do not do so, the disk may fail to boot because the physical device file associated with the HDLM logical device file has changed. You can find out the path used by HDLM to determine the logical device file name of an HDLM device by executing the procedure described in 4.4.1.3.

1. Migrate to the local boot disk environment, following the procedure described in 3.7.

2. Change the path configuration.
3. Return to the HDLM boot disk environment, following the procedure described in 3.6.

4.4.2 Overview of reconfiguring the HDLM device

Whether or not the host must be restarted after the reconfiguration depends on the contents of the reconfiguration. This section describes the reconfigurations that does and does not require the restarting of the host. It also gives notes for reconfiguration.

4.4.2.1 Reconfiguring the HDLM device

You can change the HDLM device configuration in one of the following two ways:

- Restart the host, with the reconfiguration option specified.

  Execute the `dlmsetconf` utility for creating the HDLM driver configuration definition file to update the HDLM driver configuration definition file (`/kernel/drv/dlmfdrv.conf`), and then restart the host, with the reconfiguration option specified.

- Use the `dlmcfgmgr` utility for managing the HDLM configuration, without restarting the host.

  Execute the `dlmsetconf` utility to update the HDLM driver configuration definition file (`/kernel/drv/dlmfdrv.conf`), and then execute the `dlmcfgmgr` utility.

Modifying the configuration of the HDLM device without restarting the host while HDLM is running is called a *dynamic reconfiguration of the HDLM device*.

Dynamic reconfiguration of the HDLM device is applicable to the following operations:

- Adding a new logical unit (LU)

  A *new LU* is one that has been added to the host, and recognized by the HDLM driver for the first time. When a new LU is added, the path connected to that LU is also added.

- Reconfiguring an LU that is not managed by HDLM, as an HDLM-managed device

  Reconfigure an LU that has already been added to the host, but is not managed by HDLM, as an HDLM-managed device.

The following operations are not supported in the dynamic reconfiguration of the HDLM device. To enable the following operations, restart the host with the reconfiguration option specified.

- Re-adding and deleting an existing LU
- Changing the setting of an existing LU
- Adding or deleting a path to an existing LU
- Reconfiguring an existing LU (an HDLM-managed device), as a non-HDLM-managed device
- Adding and replacing hardware (including host bus adapter and switches)
- Placing the following under HDLM management: LUs and paths recognized by a host through an HBA port not used by HDLM.

For details on the dlmsetconf utility, see 7.5 in Chapter 7.

For details on the dlmcfgmgr utility, see 7.3 in Chapter 7.

4.4.2.2 Notes on reconfiguring the HDLM device

The following describes the notes on reconfiguring the HDLM device:

- Dynamic reconfiguration of HDLM devices can only be performed for Solaris 8 version 1/01 or later, or Solaris 9.
- When adding a new LU, since the HBA driver and either the sd or ssd driver recognize the reconfiguration, you may need to restart the host, depending on the following conditions:
  - What kind of HBA and HBA driver are used
  - The configuration for a newly-added LU

  For details, refer to the documentation for the HBA driver being used.
- When the HBA used for accessing an HDLM-managed LU is removed or replaced, processing to unconfigure or reconfigure the HBA driver and either the sd or ssd driver might not complete until the host is restarted.
- When you modify disk device configurations (such as deleting or removing LUs), or server-side hardware configurations, execute the dlmsetconf utility, and then restart the host, with the reconfiguration option specified. Also, when the hardware supports dynamic reconfiguration for the disk and you want to dynamically add new LUs, execute the dlmsetconf utility, and then execute the dlmcfgmgr utility for managing the HDLM configuration.
- Without restarting the system, when you add a new path to the LU that HDLM manages, the logical device file of the sd or ssd device corresponding to the path will not be created. When you want to use the newly added path in HDLM, execute the dlmsetconf utility and then restart the host, with the reconfiguration option specified.

4.4.3 Adding a new logical unit

This section describes the operations for configuring a new LU recognized by the host, as an HDLM-managed device.

In a Solaris 8 version1/01 or later or Solaris 9 environment that supports dynamic reconfiguration of HDLM devices, you can modify settings so that new LUs are recognized automatically by the HDLM driver, without restarting the host. In this case, you need to use the dlmcfgmgr utility for managing the HDLM configuration in order to have the HDLM driver recognize the contents of the HDLM driver configuration definition file (/kernel/drv/dlmfdrv.conf) without having to restart the host. For details on the dlmcfgmgr utility, see 7.3.

4.4.3.1 Notes

The following describes the notes on adding new LUs:

- If you do not want to configure an LU that was newly added to the host, as an HDLM managed device, first set the LU as a non-HDLM-managed device, and then execute the `dlmcfgmgr` utility. If the newly added LU is set as a non-HDLM-managed device, reconfiguration of the LU for the HDLM driver will not be performed.
- You cannot perform dynamic reconfiguration when you add a new LU to HDLM and an existing instance number is allocated to the path of the LU. In such a case, restart the host.
- If the major and minor numbers of the sd or ssd device of an LU to be newly added to HDLM matches the major or minor number of an sd or ssd device previously registered in HDLM, dynamic reconfiguration cannot be performed. You must restart the host.
- Dynamic reconfiguration cannot be performed when using VxVM, because the host must be restarted in order to use (in VxVM) an LU that has been newly added to HDLM.
- After you execute the `dlmcfgmgr` utility to add a new LU, when you execute the path status display command (the `view -path` operation), the HDevName of the added LU may be displayed as a hyphen (-).

4.4.3.2 Adding a new LU (when not using Sun Cluster)

To add a new LU:

1. Set up the storage subsystem and switch so that Solaris recognizes the LU.

   Specify the following settings to have Solaris recognize the newly added LU:

   - Storage subsystem settings (to map the internal LU to the storage port LUN)
   - Connection between the storage port and the host HBA port
   - Fibre Channel switch settings
   - Settings for the host HBA, and for the sd or ssd driver

   For details on how to perform these settings, refer to the documentation for the HBA being used. You may need to restart the host, depending on the HBA being used, its drivers, and the path configuration of the newly added LU.

2. Check the newly recognized LU.

   You can check the LU by using the `format` command. Figure 4.1 shows an example of executing the `format` command. The shaded portion represents the newly recognized LU.

```
# format
Searching for disks...done

AVAILABLE DISK SELECTIONS:
       0. c1t9d0 <FUJITSU-MAP3735NP-0107 cyl 47994 alt 2 hd 4 sec 747>
          /pci@8,700000/scsi@6,1/sd@9,0
       1. c2t1d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424>
          /pci@8,600000/SUNW,qlc@4/fp@0,0/ssd@w21000004cfaf4b6e,0
       2. c2t2d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424>
          /pci@8,600000/SUNW,qlc@4/fp@0,0/ssd@w21000004cfb3a627,0
       3. c3t50060E80033A9902d2 <HITACHI-OPEN-3-SUN-2105 cyl 3336 alt 2 hd 15 sec 96>
          /pci@8,700000/SUNW,qlc@2/fp@0,0/ssd@w50060e80033a9902,2
       4. c3t50060E80033A9938d2 <HITACHI-OPEN-3-SUN-2105 cyl 3336 alt 2 hd 15 sec 96>
          /pci@8,700000/SUNW,qlc@2/fp@0,0/ssd@w50060e80033a9938,2
       5. c4t50060E80033A9902d2 <HITACHI-OPEN-3-SUN-2105 cyl 3336 alt 2 hd 15 sec 96>
          /pci@8,700000/SUNW,qlc@3/fp@0,0/ssd@w50060e80033a9902,2
       6. c4t50060E80033A9938d2 <HITACHI-OPEN-3-SUN-2105 cyl 3336 alt 2 hd 15 sec 96>
          /pci@8,700000/SUNW,qlc@3/fp@0,0/ssd@w50060e80033a9938,2
       7. c5t50060E80033A9902d0 <HITACHI-OPEN-3-SUN-2105 cyl 3336 alt 2 hd 15 sec 96>
          /pseudo/dlmndrv@1/dlmfdrv@w50060e80033a9902,0
       8. c5t50060E80033A9902d1 <HITACHI-OPEN-3-SUN-2105 cyl 3336 alt 2 hd 15 sec 96>
          /pseudo/dlmndrv@1/dlmfdrv@w50060e80033a9902,1
Specify disk (enter its number):
```

Figure 4.1 Example of executing the format command

3. Execute the `dlmsetconf` utility for creating the HDLM driver configuration definition file to update the HDLM driver configuration definition file (`/kernel/drv/dlmfdrv.conf`).

```
# /opt/DynamicLinkManager/bin/dlmsetconf
```

4. Set up LUs not to be managed by HDLM.

   If you do not want HDLM to manage one of the newly recognized LUs, set up the LU so that it is excluded from HDLM management. See 3.5.4.

   If you do not want to specify settings to exclude such LUs from management, skip to step 5.

5. If you are using VxVM, you need to disable, from VxVM, any devices (such as sd or ssd devices) other than the HDLM devices corresponding to the added LU.

   Specify settings as shown in 3.12.3.2 or 3.12.3.3.

6. Force the HDLM driver to recognize the added LU.

   In environments supporting dynamic reconfiguration of HDLM devices, execute the `dlmcfgmgr` utility for managing the HDLM configuration. HDLM device will be reconfigured without restarting the host.

```
# /opt/DynamicLinkManager/bin/dlmcfgmgr -a
```

   In environments not supporting dynamic reconfiguration of HDLM devices, restart the host, with the reconfiguration option specified.

```
# reboot -- -r
```

7. Check the status of the path for accessing the added LU.

   The following shows an example of executing the command. The shaded portion represents the path that accesses the added LU.

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -path
Paths:000012 OnlinePaths:000012
PathStatus   IO-Count    IO-Errors
Online       24           0

PathID PathName                          DskName
  iLU             ChaPort Status   Type IO-Count  IO-Errors  DNum HDevName
000000 0000.0000.00000000000111A6.0000 HITACHI .OPEN-3      -SUN.15001
  0659            1C      Online   Own         6          0     - c5t50060E80033A9902d0
000001 0000.0000.00000000000011080.0000 HITACHI .OPEN-3      -SUN.15001
  0659            4j      Online   Own         2          0     - c5t50060E80033A9902d0
000002 0000.0000.00000000000111A6.0001 HITACHI .OPEN-3      -SUN.15001
  065A            1C      Online   Own         6          0     - c5t50060E80033A9902d1
000003 0000.0000.00000000000011080.0001 HITACHI .OPEN-3      -SUN.15001
  065A            4j      Online   Own         2          0     - c5t50060E80033A9902d1
000004 0001.0000.00000000000111A6.0000 HITACHI .OPEN-3      -SUN.15001
  0659            1C      Online   Own         2          0     - c5t50060E80033A9902d0
000005 0001.0000.00000000000011080.0000 HITACHI .OPEN-3      -SUN.15001
  0659            4j      Online   Own         2          0     - c5t50060E80033A9902d0
000006 0001.0000.00000000000111A6.0001 HITACHI .OPEN-3      -SUN.15001
  065A            1C      Online   Own         2          0     - c5t50060E80033A9902d1
000007 0001.0000.00000000000011080.0001 HITACHI .OPEN-3      -SUN.15001
  065A            4j      Online   Own         2          0     - c5t50060E80033A9902d1
000008 0000.0000.00000000000111A6.0002 HITACHI .OPEN-3      -SUN.15001
  065B            1C      Online   Own         0          0     - c5t50060E80033A9902d2
000009 0000.0000.00000000000011080.0002 HITACHI .OPEN-3      -SUN.15001
  065B            4j      Online   Own         0          0     - c5t50060E80033A9902d2
000010 0001.0000.00000000000111A6.0002 HITACHI .OPEN-3      -SUN.15001
  065B            1C      Online   Own         0          0     - c5t50060E80033A9902d2
000011 0001.0000.00000000000011080.0002 HITACHI .OPEN-3      -SUN.15001
  065B            4j      Online   Own         0          0     - c5t50060E80033A9902d2
KAPL01001-I The HDLM command completed normally. Operation name = view, completion time =
yyyy/mm/dd hh:mm:ss
#
```

8. Register the added LU to the application.

4.4.3.3 Adding a new LU by restarting the nodes (when using Sun Cluster)

Note

Have the HDLM drivers on all nodes in the cluster recognize the LU newly recognized by the node, and make the configuration the same for each node.

To add a new LU:

1. Set up the storage subsystem and switch so that Solaris recognizes the LU.

2. Check the LU newly recognized by the node.

   You can check the LU by using the `format` command. Figure 4.2 shows an example of executing the `format` command. The shaded portion represents the LU newly recognized by the node.

```
# format
Searching for disks...done

AVAILABLE DISK SELECTIONS:
       0. c1t9d0 <FUJITSU-MAP3735NP-0107 cyl 47994 alt 2 hd 4 sec 747>
          /pci@8,700000/scsi@6,1/sd@9,0
       1. c2t1d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424>
          /pci@8,600000/SUNW,qlc@4/fp@0,0/ssd@w21000004cfaf4b6e,0
       2. c2t2d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424>
          /pci@8,600000/SUNW,qlc@4/fp@0,0/ssd@w21000004cfb3a627,0
       3. c3t50060E80033A9902d2 <HITACHI-OPEN-3-SUN-2105 cyl 3336 alt 2 hd 15 sec 96>
          /pci@8,700000/SUNW,qlc@2/fp@0,0/ssd@w50060e80033a9902,2
       4. c3t50060E80033A9938d2 <HITACHI-OPEN-3-SUN-2105 cyl 3336 alt 2 hd 15 sec 96>
          /pci@8,700000/SUNW,qlc@2/fp@0,0/ssd@w50060e80033a9938,2
       5. c4t50060E80033A9902d2 <HITACHI-OPEN-3-SUN-2105 cyl 3336 alt 2 hd 15 sec 96>
          /pci@8,700000/SUNW,qlc@3/fp@0,0/ssd@w50060e80033a9902,2
       6. c4t50060E80033A9938d2 <HITACHI-OPEN-3-SUN-2105 cyl 3336 alt 2 hd 15 sec 96>
          /pci@8,700000/SUNW,qlc@3/fp@0,0/ssd@w50060e80033a9938,2
       7. c5t50060E80033A9902d0 <HITACHI-OPEN-3-SUN-2105 cyl 3336 alt 2 hd 15 sec 96>
          /pseudo/dlmndrv@1/dlmfdrv@w50060e80033a9902,0
       8. c5t50060E80033A9902d1 <HITACHI-OPEN-3-SUN-2105 cyl 3336 alt 2 hd 15 sec 96>
          /pseudo/dlmndrv@1/dlmfdrv@w50060e80033a9902,1
Specify disk (enter its number):
```

Figure 4.2 Example of executing the format command

3. If setup is performed for Quorum devices and disk device groups, on an LU newly recognized by the node, release these settings.

   When you unregister the SDS or SVM diskset, the diskset data will be deleted so first back up any necessary data. For details on how to remove settings for Quorum devices and VxVM disk groups, see 3.16.2.2 of 3.16.2. For details on how to remove SDS or SVM disksets, see 3.16.2.5 of 3.16.2 or 3.16.2.6 of 3.16.2.

4. Make sure that the key used for disk reservation has been registered for the LU newly recognized by the node.

   Execute the following command on any node:

```
# /usr/cluster/lib/sc/scsi -c inkeys -d
the-name-of-the-logical-device-file-on-slice-2-of-the-LU-newly-recognized-by-the-node
```

   The following shows an example of executing the command:

```
# /usr/cluster/lib/sc/scsi -c inkeys -d /dev/rdsk/c3t50060E80033A9902d2s2
```

If the key is displayed, perform the steps beginning with step 5. Even when you release the setup for Quorum devices and disk device groups, since the key used for disk reservation will remain registered, you will need to restart the node in non-cluster mode and unregister the key.

If the key is not displayed, perform the steps beginning with step 9.

5. Simultaneously stop all nodes in the cluster.

   Execute the following command on any node:

```
# /usr/cluster/bin/scshutdown -g0 -y
```

6. Start all nodes in the cluster in non-cluster mode.

   At the `ok` prompt, execute the `boot -x` command.

7. Unregister the key for the LU newly recognized by the node.

   Execute the following command on any node:

```
# /usr/cluster/lib/sc/scsi -c scrub -d
the-name-of-the-logical-device-file-on-slice-2-of-the-LU-newly-recognized-by-the-node
```

   The following shows an example of executing the command:

```
# /usr/cluster/lib/sc/scsi -c scrub -d /dev/rdsk/c3t50060E80033A9902d2s2
```

8. Make sure that the key for the LU newly recognized by the node has been unregistered.

   Perform step 4 again on any node, and check that the key is not displayed.

9. Make sure that the logical device files for the sd or ssd device are backed up.

   When you restart the node, with the reconfiguration option specified, after running the `dlmsetconf` utility in the next step (step 10), the logical device files of the sd or ssd device to be managed by HDLM will be deleted. Make sure that these files are backed up before you execute the `dlmsetconf` utility to restart the node. For details on how to back up files, see step 4 in 3.5.1.1 of 3.5.1.

10. Use the following `dlmsetconf` utility to create the HDLM driver configuration definition file (`/kernel/drv/dlmfdrv.conf`) and the HDLM driver unconfiguration definition file (`/opt/DynamicLinkManager/config/dlmfdrv.unconf`).

```
# /opt/DynamicLinkManager/bin/dlmsetconf
```

    The device mentioned in 2.1 will be set as an HDLM-managed device.

    When the `dlmsetconf` utility is executed, the KAPL10242-I message will be displayed prompting you to confirm the deletion of the logical device files for the sd or ssd device.

    If you respond by entering `y`, and then execute the `dlmsetconf` utility and restart the node, with the reconfiguration option specified, the logical device files of the sd or ssd device to be managed by HDLM will be deleted. If the logical device files are already backed up, enter `y` and execute the `dlmsetconf` utility. If the files are not backed up, respond by entering `n`, and then cancel processing for the `dlmsetconf` utility. In this case, back up the logical device files and execute the `dlmsetconf` utility, by referring to step 4 in 3.5.1.1 of 3.5.1.

When the dlmsetconf utility is executed and then the host is restarted with the reconfiguration option specified, all sd or ssd devices connected to the host will be set to be managed by HDLM.

Note

Restart the nodes simultaneously as shown in steps 12 and 13, after the setup for LUs newly recognized on all nodes has completed.

11. If you are using VxVM, you need to disable, from VxVM, any devices (such as sd or ssd devices) other than the HDLM devices corresponding to the added LU.

    Specify settings according to 3.12.3.2 of 3.12.3 or 3.12.3.3 of 3.12.3.

12. Stop all nodes in the cluster.

    If the key is displayed in step 4, and nodes are running in non-cluster mode in step 6, execute the following commands on all nodes:

```
# shutdown -g0 -y -i0
```

    When running nodes in cluster mode, execute the following command on any node:

```
# /usr/cluster/bin/scshutdown -g0 -y
```

13. Start all nodes in the cluster simultaneously, in cluster mode.

    At the ok prompt, execute the boot -r command.

    Note

    The following message might be output to the console window or syslog while the node is running:

```
Could not read symbolic link for:
/dev/rdsk/c3t50060E80033A9902d2s2 path not loaded
```

    This message is output when HDLM deletes a logical device file for an sd or ssd device managed by HDLM. Perform the procedure for updating the device ID information (steps 17 to 20) so that this message is not output next time the node starts. For details, refer to the Sun Cluster documentation.

14. Execute the dlnkmgr command's view operation on all nodes in the cluster, and display the status of each program.

    The following shows an example of executing the command:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -sys
HDLM Version               : xx-xx
Service Pack Version       :
Load Balance               : on(rr)
Support Cluster            :
Elog Level                 : 3
Elog File Size (KB)        : 9900
Number Of Elog Files       : 2
Trace Level                : 0
Trace File Size(KB)        : 1000
Number Of Trace Files      : 4
```

```
Path Health Checking         : on(30)
Auto Failback                : off
Intermittent Error Monitor   : off
HDLM Manager Ver      WakeupTime
Alive        xx-xx    yyyy/mm/dd hh:mm:ss
HDLM Alert Driver Ver      WakeupTime          ElogMem Size
Alive             xx-xx    yyyy/mm/dd hh:mm:ss
HDLM Driver  Ver      WakeupTime
Alive        xx-xx    yyyy/mm/dd hh:mm:ss
License Type Expiration
Permanent    -
KAPL01001-I The HDLM command completed normally. operation name = view, completion time =
yyyy/mm/dd hh:mm:ss
#
```

- Even when cluster software is used, it is not displayed in `Support Cluster`. Nevertheless, no problem exists with the functionality for cluster support.

15. Use the execution results from the `view` operation to check that the programs are operating normally.

    If `HDLM Manager`, `HDLM Alert Driver`, and `HDLM Driver` are all `Alive`, the programs are operating normally.

16. Check the status of the path accessing the LU recognized for the first time by the HDLM driver.

    The following shows an example of executing the command. The shaded portion indicates the path accessing the LU recognized for the first time by the HDLM driver.

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -path
Paths:000012 OnlinePaths:000012
PathStatus   IO-Count    IO-Errors
Online       24           0

PathID PathName                        DskName
  iLU              ChaPort  Status     Type IO-Count   IO-Errors  DNum HDevName
000000 0000.0000.00000000000111A6.0000 HITACHI .OPEN-3      -SUN.15001
  0659             1C       Online     Own         6           0     - c5t50060E80033A9902d0
000001 0000.0000.00000000000011080.0000 HITACHI .OPEN-3      -SUN.15001
  0659             4j       Online     Own         2           0     - c5t50060E80033A9902d0
000002 0000.0000.00000000000111A6.0001 HITACHI .OPEN-3      -SUN.15001
  065A             1C       Online     Own         6           0     - c5t50060E80033A9902d1
000003 0000.0000.00000000000011080.0001 HITACHI .OPEN-3      -SUN.15001
  065A             4j       Online     Own         2           0     - c5t50060E80033A9902d1
000004 0000.0000.00000000000111A6.0002 HITACHI .OPEN-3      -SUN.15001
  065B             1C       Online     Own         0           0     - c5t50060E80033A9902d2
000005 0000.0000.00000000000011080.0002 HITACHI .OPEN-3      -SUN.15001
  065B             4j       Online     Own         0           0     - c5t50060E80033A9902d2
000006 0001.0000.00000000000111A6.0000 HITACHI .OPEN-3      -SUN.15001
  0659             1C       Online     Own         2           0     - c5t50060E80033A9902d0
000007 0001.0000.00000000000011080.0000 HITACHI .OPEN-3      -SUN.15001
  0659             4j       Online     Own         2           0     - c5t50060E80033A9902d0
000008 0001.0000.00000000000111A6.0001 HITACHI .OPEN-3      -SUN.15001
  065A             1C       Online     Own         2           0     - c5t50060E80033A9902d1
000009 0001.0000.00000000000011080.0001 HITACHI .OPEN-3      -SUN.15001
  065A             4j       Online     Own         2           0     - c5t50060E80033A9902d1
000010 0001.0000.00000000000111A6.0002 HITACHI .OPEN-3      -SUN.15001
  065B             1C       Online     Own         0           0     - c5t50060E80033A9902d2
000011 0001.0000.00000000000011080.0002 HITACHI .OPEN-3      -SUN.15001
  065B             4j       Online     Own         0           0     - c5t50060E80033A9902d2
KAPL01001-I The HDLM command completed normally. Operation name = view, completion time =
yyyy/mm/dd hh:mm:ss
#
```

17. Delete the device ID information for devices that do not exist.

    Execute the following command:

```
# /usr/cluster/bin/scdidadm -C
```

18. Allocate the device ID.

    Execute the following command:

```
# /usr/cluster/bin/scdidadm -r
```

19. Notify the node of the most recent device ID information.

    Execute the following command:

```
# /usr/cluster/bin/scdidadm -ui
```

20. Update the Sun Cluster global device namespace.

    Execute the following command:

```
```

```
# /usr/cluster/bin/scgdevs
```

21. When using, as a Quorum device, an HDLM device of an LU newly recognized for the first time by HDLM driver, set up the Quorum device on any node in the cluster.

    The following shows an example of executing the command:

    Check the device ID of the HDLM device.

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -path
# /usr/cluster/bin/scdidadm -L
```

    Set up the Quorum device:

```
# /usr/cluster/bin/scconf -a -q globaldev=d4
```

    For cluster configurations with two nodes, execute the following command to clear the install mode for the cluster:

```
# /usr/cluster/bin/scconf -c -q reset
```

22. Execute the following command to view the correspondence between HDevName and iLU:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -lu
```

    Check the correspondence between `HDevName` (the name of the logical device file of the HDLM device) and `iLU` (the internal LU number) from the execution results of the command, compare the correspondence with the correspondence that took place before the HDLM installation, and then re-register HDLM devices with the programs that will use them.

    If you want to use volume management software to create a file system on an HDLM-managed device, skip to step 24.

23. Perform registration for the disk device group.

    If you are using VxVM, see 3.12.5.

    If you are using SDS, see 3.13.2.2 of 3.13.2.

    If you are using SVM, see 3.14.2.2 of 3.14.2.

24. Create the file system and set up the global mount as necessary.

    For details, refer to the Sun Cluster documentation.

4.4.3.4 Adding a new LU via dynamic reconfiguration(When using Sun Cluster)

Note

- A new LU cannot be added via dynamic reconfiguration when you are using VxVM, because the host must be restarted in order to use (in VxVM) an LU that has been newly added to HDLM. Perform the procedure described in 4.4.3.3.
- For details on the precautions for the `dlmcfgmgr` utility, see 7.3.

To add a new LU:

1. Set up the storage subsystem and switch so that Solaris recognizes the LU.
2. Check the LU newly recognized by the node.

   Use the `format` command to check the LU. Figure 4.3 shows an example of executing the `format` command. The shaded portion represents the newly recognized LU.

```
# format
Searching for disks...done

AVAILABLE DISK SELECTIONS:
       0. c1t9d0 <FUJITSU-MAP3735NP-0107 cyl 47994 alt 2 hd 4 sec 747>
          /pci@8,700000/scsi@6,1/sd@9,0
       1. c2t1d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424>
          /pci@8,600000/SUNW,qlc@4/fp@0,0/ssd@w21000004cfaf4b6e,0
       2. c2t2d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424>
          /pci@8,600000/SUNW,qlc@4/fp@0,0/ssd@w21000004cfb3a627,0
       3. c3t50060E80033A9902d2 <HITACHI-OPEN-3-SUN-2105 cyl 3336 alt 2 hd 15 sec 96>
          /pci@8,700000/SUNW,qlc@2/fp@0,0/ssd@w50060e80033a9902,2
       4. c3t50060E80033A9938d2 <HITACHI-OPEN-3-SUN-2105 cyl 3336 alt 2 hd 15 sec 96>
          /pci@8,700000/SUNW,qlc@2/fp@0,0/ssd@w50060e80033a9938,2
       5. c4t50060E80033A9902d2 <HITACHI-OPEN-3-SUN-2105 cyl 3336 alt 2 hd 15 sec 96>
          /pci@8,700000/SUNW,qlc@3/fp@0,0/ssd@w50060e80033a9902,2
       6. c4t50060E80033A9938d2 <HITACHI-OPEN-3-SUN-2105 cyl 3336 alt 2 hd 15 sec 96>
          /pci@8,700000/SUNW,qlc@3/fp@0,0/ssd@w50060e80033a9938,2
       7. c5t50060E80033A9902d0 <HITACHI-OPEN-3-SUN-2105 cyl 3336 alt 2 hd 15 sec 96>
          /pseudo/dlmndrv@1/dlmfdrv@w50060e80033a9902,0
       8. c5t50060E80033A9902d1 <HITACHI-OPEN-3-SUN-2105 cyl 3336 alt 2 hd 15 sec 96>
          /pseudo/dlmndrv@1/dlmfdrv@w50060e80033a9902,1
Specify disk (enter its number):
```

Figure 4.3 Example of executing the format command

3. If setup is performed for Quorum devices and disk device groups, on an LU newly recognized by the node, release these settings.

   When you unregister the SDS or SVM diskset, the diskset data will be deleted so first back up any necessary data. For details on how to remove settings for Quorum devices and VxVM disk groups, see 3.16.2.2 of 3.16.2. For details on how to remove SDS or SVM disksets, see 3.16.2.5 of 3.16.2 or 3.16.2.6 of 3.16.2.

4. Make sure that the key used for disk reservation has been registered for the LU newly recognized by the node.

   Execute the following command on any node:

```
# /usr/cluster/lib/sc/scsi -c inkeys -d
the-name-of-the-logical-device-file-on-slice-2-of-the-LU-newly-recognized-by-the-node
```

   The following shows an example of executing the command:

```
# /usr/cluster/lib/sc/scsi -c inkeys -d /dev/rdsk/c3t50060E80033A9902d2s2
```

Add the new LU to HDLM by following the steps in 4.4.3.3 beginning with step 5. Even when you release the setup for Quorum devices and disk device groups, since the key used for disk reservation will remain registered, you will need to restart the node in non-cluster mode and unregister the key.

If the key is not displayed, perform the steps beginning with step 5.

5. Make sure that the logical device files for the sd or ssd devices are backed up.

   When you execute the `dlmcfgmgr` utility in step 8 after executing the `dlmsetconf` utility in step 6, since the logical device files of the sd or ssd devices to be managed by HDLM will be deleted, make sure that you back up these files beforehand. For details on how to back up files, see step 4 in 3.5.1.1 of 3.5.1.

6. Use the following `dlmsetconf` utility to create the HDLM driver configuration definition file (`/kernel/drv/dlmfdrv.conf`) and the HDLM driver unconfiguration definition file (`/opt/DynamicLinkManager/config/dlmfdrv.unconf`):

   ```
   # /opt/DynamicLinkManager/bin/dlmsetconf
   ```

   The device mentioned in 2.1 will be set to be managed by HDLM.

   When the `dlmsetconf` utility is executed, the KAPL10242-I message will be output.

   If you respond by entering `y`, when you run the `dlmcfgmgr` utility in step 8 after running the `dlmsetconf` utility, the logical device files of the sd or ssd devices to be managed by HDLM will be deleted. If the logical device files are already backed up, respond by entering `y` and running the `dlmsetconf` utility. If the files are not backed up, respond by entering `n`, and then cancel processing for the `dlmsetconf` utility. Then, backup the logical device files by referring to step 4 in 3.5.1.1 of 3.5.1.

   All sd or ssd devices connected to the host will be set as HDLM-managed devices.

7. If you do not want HDLM to manage one of the LUs newly recognized by the node, set up the LU so that it is excluded from HDLM management.

   For details, see 3.5.4. If you do not want to specify settings to exclude the LUs from HDLM management, skip to step 8.

8. Execute the `dlmcfgmgr` utility shown below so that the HDLM driver recognizes the LUs newly recognized by the node:

   ```
   # /opt/DynamicLinkManager/bin/dlmcfgmgr -a
   ```

9. Make sure that the HDLM device of the LU recognized for the first time by the HDLM driver is displayed by the `format` command.

   Also check that the sd or ssd devices managed by HDLM are not displayed by the `format` command. Figure 4.4 shows the example of executing the `format` command. The shaded portion indicates the HDLM device of the LU recognized for the first time by the HDLM driver.

```
# format
Searching for disks...done

AVAILABLE DISK SELECTIONS:
       0. c1t9d0 <FUJITSU-MAP3735NP-0107 cyl 47994 alt 2 hd 4 sec 747>
          /pci@8,700000/scsi@6,1/sd@9,0
       1. c2t1d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424>
          /pci@8,600000/SUNW,qlc@4/fp@0,0/ssd@w21000004cfaf4b6e,0
       2. c2t2d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424>
          /pci@8,600000/SUNW,qlc@4/fp@0,0/ssd@w21000004cfb3a627,0
       3. c5t50060E80033A9902d0 <HITACHI-OPEN-3-SUN-2105 cyl 3336 alt 2 hd 15 sec 96>
          /pseudo/dlmndrv@1/dlmfdrv@w50060e80033a9902,0
       4. c5t50060E80033A9902d1 <HITACHI-OPEN-3-SUN-2105 cyl 3336 alt 2 hd 15 sec 96>
          /pseudo/dlmndrv@1/dlmfdrv@w50060e80033a9902,1
       5. c5t50060E80033A9902d2 <HITACHI-OPEN-3-SUN-2105 cyl 3336 alt 2 hd 15 sec 96>
          /pseudo/dlmndrv@1/dlmfdrv@w50060e80033a9902,2
Specify disk (enter its number):
```

Figure 4.4 Example of executing the format command

10. Check the status of the path accessing the LU recognized for the first time by the HDLM driver.

    The following shows an example of executing the command. The shaded portion indicates the path accessing the LU recognized for the first time by the HDLM driver.

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -path
Paths:000012 OnlinePaths:000012
PathStatus   IO-Count    IO-Errors
Online       24           0

PathID PathName                         DskName
  iLU            ChaPort Status   Type IO-Count   IO-Errors  DNum HDevName
000000 0000.0000.00000000000111A6.0000 HITACHI .OPEN-3      -SUN.15001
  0659           1C      Online   Own        6            0     - c5t50060E80033A9902d0
000001 0000.0000.00000000000011080.0000 HITACHI .OPEN-3     -SUN.15001
  0659           4j      Online   Own        2            0     - c5t50060E80033A9902d0
000002 0000.0000.00000000000111A6.0001 HITACHI .OPEN-3      -SUN.15001
  065A           1C      Online   Own        6            0     - c5t50060E80033A9902d1
000003 0000.0000.00000000000011080.0001 HITACHI .OPEN-3     -SUN.15001
  065A           4j      Online   Own        2            0     - c5t50060E80033A9902d1
000004 0001.0000.00000000000111A6.0000 HITACHI .OPEN-3      -SUN.15001
  0659           1C      Online   Own        2            0     - c5t50060E80033A9902d0
000005 0001.0000.00000000000011080.0000 HITACHI .OPEN-3     -SUN.15001
  0659           4j      Online   Own        2            0     - c5t50060E80033A9902d0
000006 0001.0000.00000000000111A6.0001 HITACHI .OPEN-3      -SUN.15001
  065A           1C      Online   Own        2            0     - c5t50060E80033A9902d1
000007 0001.0000.00000000000011080.0001 HITACHI .OPEN-3     -SUN.15001
  065A           4j      Online   Own        2            0     - c5t50060E80033A9902d1
000008 0000.0000.00000000000111A6.0002 HITACHI .OPEN-3      -SUN.15001
  065B           1C      Online   Own        0            0     - c5t50060E80033A9902d2
000009 0000.0000.00000000000011080.0002 HITACHI .OPEN-3     -SUN.15001
  065B           4j      Online   Own        0            0     - c5t50060E80033A9902d2
000010 0001.0000.00000000000111A6.0002 HITACHI .OPEN-3      -SUN.15001
  065B           1C      Online   Own        0            0     - c5t50060E80033A9902d2
000011 0001.0000.00000000000011080.0002 HITACHI .OPEN-3     -SUN.15001
  065B           4j      Online   Own        0            0     - c5t50060E80033A9902d2
KAPL01001-I The HDLM command completed normally. Operation name = view, completion time =
yyyy/mm/dd hh:mm:ss
#
```

11. Delete the device ID information for devices that do not exist.

    Execute the following command:

    ```
    # /usr/cluster/bin/scdidadm -C
    ```

12. Allocate the device ID.

    Execute the following command:

    ```
    # /usr/cluster/bin/scdidadm -r
    ```

13. Notify the node of the most recent device ID information.

    Execute the following command:

    ```
    # /usr/cluster/bin/scdidadm -ui
    ```

14. Update the Sun Cluster global device namespace.

    Execute the following command:

```
# /usr/cluster/bin/scgdevs
```

15. When using, as a Quorum device, an HDLM device of an LU newly recognized for the first time by HDLM driver, set up the Quorum device on any node in the cluster.

    The following shows an example of executing the command:

    Check the device ID of the HDLM device.

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -path
# /usr/cluster/bin/scdidadm -L
```

    Set up the Quorum device:

```
# /usr/cluster/bin/scconf -a -q globaldev=d4
```

    For cluster configurations with two nodes, execute the following command to clear the install mode for the cluster:

```
# /usr/cluster/bin/scconf -c -q reset
```

16. Execute the following command to view the correspondence between HDevName and iLU:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -lu
```

    Check the correspondence between `HDevName` (the name of the logical device file of the HDLM device) and `iLU` (the internal LU number) from the execution results of the command, compare the correspondence with the correspondence that took place before the HDLM installation, and then re-register HDLM devices with the programs that will use them.

    If you want to use volume management software to create a file system on an HDLM-managed device, skip to step 18.

17. Perform registration for the disk device group.

    If you are using VxVM, see 3.12.5.

    If you are using SDS, see 3.13.2.2 of 3.13.2.

    If you are using SVM, see 3.14.2.2 of 3.14.2.

18. Create the file system and set up the global mount as necessary.

    For details, refer to the Sun Cluster documentation.

4.4.4 Configuration changes such as deleting a logical unit

This section describes the procedure for changing the configuration in any of the following ways:

- Deleting an LU
- Changing an LU[#]
- Changing the host LUN

- Changing, adding, or deleting a channel adapter port
- Changing a target ID
- Adding or replacing a Fibre Channel switch
- Replacing, deleting, or adding a host bus adapter

#

To change an LU, we recommend that you add a new LU and then delete the existing one, rather than changing a LU in the storage subsystem management program. After adding the new LU, move data as required from the LU you are deleting to the new LU. For details on adding a new LU, see 4.4.3. If you use the storage subsystem management program to change the LDEV of the LU, an instance of the HDLM driver will not be created. As a result, it may not be possible to verify the changed configuration until you create an HDLM driver configuration definition file (`/kernel/drv/dlmfdrv.conf`).

1. Back up the data in the file systems, volumes, and metadevices that use the HDLM device you are reconfiguring.
2. Edit the `/etc/vfstab` file.

   If automatic mounting is enabled, disable it by removing the definitions of logical device files for the HDLM devices (`c`*U*`t`*X*`d`*Y*`s`*Z*) you are reconfiguring from the `/etc/vfstab` file.
3. If you are using Sun Cluster, clear the settings for the disk device group that contains the HDLM devices you are reconfiguring. If a Quorum device is among the HDLM devices you are reconfiguring, clear the Quorum device settings.

   For details, see 3.5.1.10.
4. If you are using VCS, unregister the HDLM devices you are reconfiguring and VxVM disk groups containing those HDLM devices, from the resource groups. For details on how to remove the VCS settings, refer to the VCS documentation.
5. If you are using VxVM, unregister the HDLM devices associated with the LUs you are deleting.

   If VxVM uses the LU you are deleting, delete the volumes and VxVM disk groups. To have another LU take over from the LU you are deleting, replace the LU.
6. If you are using SDS or SVM, perform the following as required:
   - If you are using a shared diskset, unregister the HDLM devices you are reconfiguring from the shared diskset. If the path used to reserve an LU is changed, it may not be possible to clear the reservation.
   - To delete an LU when using a local diskset, unregister the HDLM device associated with that LU from SDS or SVM. If SDS or SVM is using the LU, delete the metadevices and volumes. To have another LU take over from the LU you are deleting, replace the LU.
7. Shut down the host.
8. Change the configuration.
9. Boot the host.
10. Make sure that Solaris recognizes the reconfigured operating environment.

    If you added a path:

There is nothing to check.

If you deleted a path:

Execute the `dlnkmgr` command's `view` operation specifying the `-lu` parameter, and check that the deleted path is `Offline(E)`. An execution example is shown below:

```
# dlnkmgr view -lu
...
iLU  HDevName              Device PathID Status
063A c7t50060E80033A9938d0 ssd111 000000 Online
                           ssd121 000001 Online
                           ssd76  000016 Online
                           ssd86  000017 Online
                           ssd41  000032 Online
                           ssd51  000033 Online
                           ssd6   000048 Offline(E)
                           ssd16  000049 Offline(E)
...
```

If you reconfigured the HDLM operating environment in any other way:

Execute the `format` command to make sure that the reconfigured environment is recognized by Solaris.

11. Move to the `/opt/DynamicLinkManager/ config` directory and then execute the following command.

    Specify a different value in *n* each time you reconfigure the environment.

```
# cp -ip dlmfdrv.unconf dlmfdrv.unconfbakn
```

12. Execute the `dlmsetconf` utility to create the HDLM driver configuration definition file (`/kernel/drv/dlmfdrv.conf`) and the HDLM driver unconfiguration definition file (`/opt/DynamicLinkManager/config/dlmfdrv.unconf`).

```
# /opt/DynamicLinkManager/bin/dlmsetconf
```

    The KAPL10242-I message appears when you execute the `dlmsetconf` utility. Type `y` in response.

13. Execute the following command to compare the HDLM driver unconfiguration definition file before and after execution of the `dlmsetconf` utility:

```
# diff  dlmfdrv.unconfbakn dlmfdrv.unconf
```

    Compare the old and new files to make sure that the reconfigured environment is recognized by HDLM. An example is shown below.

    Execution example after a host bus adapter is replaced:

    Check that all paths for the same type of HBA device have been deleted from the shaded portion, and that a path indicating the same type of HBA device has been added in the part after the shaded portion.

```
17,24c17,24
< ### Device=ssd41 target=65664 lun=0
< ### Path=/pci@8,700000/QLGC,qla@3,1/fp@0,0/ssd@w50060e80033a9938,0
< ### Device=ssd51 target=134012 lun=0
< ### Path=/pci@8,700000/QLGC,qla@3,1/fp@0,0/ssd@w50060e80033a9939,0
< ### Device=ssd6 target=65664 lun=0
< ### Path=/pci@8,700000/QLGC,qla@3/fp@0,0/ssd@w50060e80033a9938,0
< ### Device=ssd16 target=134012 lun=0
< ### Path=/pci@8,700000/QLGC,qla@3/fp@0,0/ssd@w50060e80033a9939,0
---
> ### Device=ssd111 target=65664 lun=0
> ### Path=/pci@8,700000/SUNW,emlxs@1,1/fp@0,0/ssd@w50060e80033a9938,0
> ### Device=ssd121 target=134012 lun=0
> ### Path=/pci@8,700000/SUNW,emlxs@1,1/fp@0,0/ssd@w50060e80033a9939,0
> ### Device=ssd76 target=65664 lun=0
> ### Path=/pci@8,700000/SUNW,emlxs@1/fp@0,0/ssd@w50060e80033a9938,0
> ### Device=ssd86 target=134012 lun=0
> ### Path=/pci@8,700000/SUNW,emlxs@1/fp@0,0/ssd@w50060e80033a9939,0
```

Execution example after a path is added by the addition of a host bus adapter port:

Check that a path for the same host bus adapter port has been added:

```
36a41,44
> ### Device=ssd74 target=65664 lun=2
> ###
Path=/pci@8,700000/SUNW,emlxs@1/fp@0,0/ssd@w50060e80033a9938,2
> ### Device=ssd84 target=134012 lun=2
> ###
Path=/pci@8,700000/SUNW,emlxs@1/fp@0,0/ssd@w50060e80033a9939,2
52a61,64
> ### Device=ssd73 target=65664 lun=3
> ###
Path=/pci@8,700000/SUNW,emlxs@1/fp@0,0/ssd@w50060e80033a9938,3
> ### Device=ssd83 target=134012 lun=3
> ###
Path=/pci@8,700000/SUNW,emlxs@1/fp@0,0/ssd@w50060e80033a9939,3
```

Execution example after a path is added by the addition of a channel adapter port:

Check that a path for the same channel adapter port has been added:

```
26a31,32
> ### Device=ssd119 target=134012 lun=2
> ###
Path=/pci@8,700000/SUNW,emlxs@1,1/fp@0,0/ssd@w50060e80033a9939,
2
28a35,36
> ### Device=ssd84 target=134012 lun=2
```

```
> ###
Path=/pci@8,700000/SUNW,emlxs@1/fp@0,0/ssd@w50060e80033a9939,2
34a43,44
> ### Device=ssd118 target=134012 lun=3
> ###
Path=/pci@8,700000/SUNW,emlxs@1,1/fp@0,0/ssd@w50060e80033a9939,
3
36a47,48
> ### Device=ssd83 target=134012 lun=3
> ###
Path=/pci@8,700000/SUNW,emlxs@1/fp@0,0/ssd@w50060e80033a9939,3
```

14. To remove an LU from HDLM management, use a text editor to edit the HDLM driver unconfiguration definition file (`/opt/DynamicLinkManager/config/dlmfdrv.unconf`).

    In the lines that define the LU that is to be removed from HDLM management, delete the hash mark (#) in front of the LDEV property line. Do not change anything else.

15. To remove an LU from HDLM management, apply the contents of the HDLM driver unconfiguration definition file to the HDLM driver configuration definition file.

    Use the following command to apply the settings in the HDLM driver unconfiguration definition file (`/opt/DynamicLinkManager/config/dlmfdrv.unconf`) to the HDLM driver configuration definition file (`/kernel/drv/dlmfdrv.conf`):

    ```
    # /opt/DynamicLinkManager/bin/dlmsetconf -u
    ```

16. When you use VxVM, from VxVM, disable the devices that you do not need to access.

    Specify settings as shown in 3.12.3.2 or 3.12.3.3.

17. Restart the host, with the reconfiguration option specified.

    Execute one of the following commands:

    - `# reboot -- -r`
    - `boot -r` (at the `ok` prompt)

18. Execute the `dlnkmgr` command's `view` operation specifying the `-lu` parameter, and check that the configuration has changed.

19. If you are using Sun Cluster, remove the device ID information of devices that do not exist. Execute the following command:

    ```
    # /usr/cluster/bin/scdidadm -C
    ```

20. If you are using Sun Cluster, assign device IDs. Execute the following command:

    ```
    # /usr/cluster/bin/scdidadm -r
    ```

21. If you are using Sun Cluster, notify the server of the new device ID information. Execute the following command:

    ```
    # /usr/cluster/bin/scdidadm -ui
    ```

22. If you are using Sun Cluster, update the Sun Cluster global device namespace. Execute the following command:

```
# /usr/cluster/bin/scgdevs
```

23. If you cleared the settings for a Quorum device in step 3, perform the following procedure:

    When using, as a Quorum device, an HDLM device of an LU newly recognized for the first time by HDLM driver, set up the Quorum device on any node in the cluster. The following shows an example of executing the command:

    Check the device ID of the HDLM device.

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -path
# /usr/cluster/bin/scdidadm -L
```

    Set up the Quorum device:

```
# /usr/cluster/bin/scconf -a -q globaldev=d4
```

    For cluster configurations with two nodes, execute the following command to clear the install mode for the cluster:

```
# /usr/cluster/bin/scconf -c -q reset
```

24. If you cleared the settings for a disk device group in step 3, perform registration for the disk device group.

    If you are using VxVM, see 3.12.5.If you are using SDS, see 3.13.2.2.

    If you are using SVM, see 3.14.2.2.

25. Recreate any metadevices you deleted in step 6.

    for details, see 3.13.2.1.

26. Recreate any volumes you deleted in step 6.

    For details, see 3.14.2.1.

27. If you unregistered VxVM disk groups from the resource groups in step 4, re-register the resource groups, referring to the VCS documentation.

28. If you recreated a file system, volume, or metadevice, restore the data backed up in step 1.

29. If you disabled the settings for automatic mounting in step 2, enter the settings again.

4.4.5 Specifying whether a logical unit is to be managed by HDLM (when not using Sun Cluster)

This section describes the procedures for modifying the configuration of an HDLM-managed device.

You must use the `dlmsetconf` utility for creating the HDLM driver configuration definition file to modify the configuration of an HDLM-managed device, therefore you should also see 7.5.

4.4.5.1 Changing an HDLM-managed device to a non-HDLM-managed device

To change an HDLM-managed device to a non-HDLM-managed device:

1. Check the LDEV information of the LU to be removed from HDLM management.

   Use the following command to check the LDEV information:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -drv
PathID HDevName Device LDEV
000000 c7t0d0   sd46   9970/9980.15001.066A
000001 c7t0d0   sd47   9970/9980.15001.066A
000002 c7t0d0   sd48   9970/9980.15001.066A
...
#

KAPL01001-I The HDLM command completed normally. Operation name = view, completion time =
yyyy/mm/dd hh:mm:ss
```

2. Edit the HDLM driver unconfiguration definition file (`/opt/DynamicLinkManager/config/dlmfdrv.unconf`).

   Referring to step 1 in 3.5.4, use a text editor to edit the HDLM driver unconfiguration definition file (`/opt/DynamicLinkManager/config/dlmfdrv.unconf`).

3. Apply the settings in the HDLM driver unconfiguration definition file (`/opt/DynamicLinkManager/config/dlmfdrv.unconf`) to the HDLM driver configuration file (`/kernel/drv/dlmfdrv.conf`).

   Execute following the `dlmsetconf` utility for creating the HDLM driver configuration definition file to apply the settings in the HDLM driver unconfiguration definition file (`/opt/DynamicLinkManager/config/dlmfdrv.unconf`) to the HDLM driver configuration definition file (`/kernel/drv/dlmfdrv.conf`):

```
# /opt/DynamicLinkManager/bin/dlmsetconf -u
```

4. Restart the host, with the reconfiguration option specified:

```
# reboot -- -r
```

4.4.5.2 Changing a non-HDLM-managed device to an HDLM-managed device

To change a non-HDLM-managed device to an HDLM-managed device:

1. Check the path information of the LU to be managed by HDLM.

   Use the `format` command to check the information for each path of the LU to be managed by HDLM. Figure 4.5 shows an example of executing the `format` command.

```
# format
Searching for disks...done
AVAILABLE DISK SELECTIONS:
     0. c1t0d0 <SUN36G cyl 24620 alt 2 hd 27 sec 107>
        /pci@8,600000/SUNW,qlc@4/fp@0,0/ssd@w21000004cf991b9b,0
     1. c1t1d0 <SUN36G cyl 24620 alt 2 hd 27 sec 107>
        /pci@8,600000/SUNW,qlc@4/fp@0,0/ssd@w21000004cf99152c,0
     2. c5t0d0 <HITACHI-OPEN-3-2103 cyl 3336 alt 2 hd 15 sec 96>
        /pci@8,700000/JNI,FCR@2/sd@0,0
     3. c5t1d0 <HITACHI-OPEN-3-2103 cyl 3336 alt 2 hd 15 sec 96>
        /pci@8,700000/JNI,FCR@2/sd@1,0
     4. c5t2d0 <HITACHI-OPEN-3-2103 cyl 3336 alt 2 hd 15 sec 96>
        /pci@8,700000/JNI,FCR@2/sd@2,0
     5. c6t0d1 <HITACHI-OPEN-3-2104 cyl 3336 alt 2 hd 15 sec 96>
        /pseudo/dlmndrv@1/dlmfdrv@0,1
     6. c6t0d2 <HITACHI-OPEN-3-2104 cyl 3336 alt 2 hd 15 sec 96>
        /pseudo/dlmndrv@1/dlmfdrv@0,2
     7. c6t0d3 <HITACHI-OPEN-3-2104 cyl 3336 alt 2 hd 15 sec 96>
        /pseudo/dlmndrv@1/dlmfdrv@0,3
Specify disk (enter its number):
```

Figure 4.5 Example of executing the format command

In this example, the shaded portion indicates the path information of the LU to be managed by HDLM.

2. Edit the HDLM driver unconfiguration definition file (`/opt/DynamicLinkManager/config/dlmfdrv.unconf`).

   Using a text editor, edit the definitions in the HDLM driver unconfiguration definition file (`/opt/DynamicLinkManager/config/dlmfdrv.unconf`) for the LU to be managed by HDLM.

   Based on the path information of the LU to be managed by HDLM checked in step 1 (shaded portion), find the correct LU in the HDLM driver unconfiguration definition file, and add a hash mark (#) in front of the `LDEV` property line corresponding to the LU. Adding the hash mark specifies that the LU indicated by that LDEV attribute is to be managed by HDLM.

```
####### T0 D0(9970/9980.15001.066A) #######
### Device=sd46 target=0 lun=0
### Path=/pci@8,700000/JNI,FCR@2/sd@0,0
### Device=sd47 target=1 lun=0
### Path=/pci@8,700000/JNI,FCR@2/sd@1,0
### Device=sd48 target=2 lun=0
### Path=/pci@8,700000/JNI,FCR@2/sd@2,0
LDEV=9970/9980.15001.066A;
```

se047110.tif

   Figure 4.6 and Figure 4.7 show examples of editing the HDLM driver unconfiguration definition file (`/opt/DynamicLinkManager/config/dlmfdrv.unconf`).

```
####### TO D0(9970/9980.15001.066A) #######
### Device=sd46 target=0 lun=0
### Path=/pci@8,700000/JNI,FCR@2/sd@0,0
### Device=sd47 target=1 lun=0
### Path=/pci@8,700000/JNI,FCR@2/sd@1,0
### Device=sd48 target=2 lun=0
### Path=/pci@8,700000/JNI,FCR@2/sd@2,0
LDEV=9970/9980.15001.066A;
```

Figure 4.6 The /opt/DynamicLinkManager/config/dlmfdrv.unconf file (before editing)

```
####### TO D0(9970/9980.15001.066A) #######
### Device=sd46 target=0 lun=0
### Path=/pci@8,700000/JNI,FCR@2/sd@0,0
### Device=sd47 target=1 lun=0
### Path=/pci@8,700000/JNI,FCR@2/sd@1,0
### Device=sd48 target=2 lun=0
### Path=/pci@8,700000/JNI,FCR@2/sd@2,0
#LDEV=9970/9980.15001.066A;
```

Figure 4.7 Example of editing the /opt/DynamicLinkManager/config/dlmfdrv.unconf file

The shaded portion indicate the text to be edited.

Note

Do not edit anything other than adding the hash mark (#) to the beginning of the LDEV property line.

3. Apply the contents that were edited in step 2 to the HDLM driver configuration definition file (/kernel/drv/dlmfdrv.conf).

   Execute following the dlmsetconf utility to apply the settings in the HDLM driver unconfiguration definition file (/opt/DynamicLinkManager/config/dlmfdrv.unconf) to the HDLM driver configuration definition file (/kernel/drv/dlmfdrv.conf):

```
# /opt/DynamicLinkManager/bin/dlmsetconf -u
```

4. Execute the dlmcfgmgr utility for managing the HDLM configuration, or restart the host, with the reconfiguration option specified.

   The following shows an example of executing the command:

   - When executing the dlmcfgmgr utility:

```
# /opt/DynamicLinkManager/bin/dlmcfgmgr -a
```

   - When restarting the host, with the reconfiguration option specified:

```
# reboot -- -r
```

4.4.6 Specifying whether a logical unit is to be managed by HDLM (when using Sun Cluster)

4.4.6.1 Changing an HDLM-managed device to a non-HDLM-managed device

Note

- When you remove an HDLM-managed LU from management, perform the remove operations on all nodes in the cluster, and make the configuration the same on each node.
- When you remove an HDLM-managed LU from management, restart the host.

To change an HDLM-managed device to a non-HDLM-managed device:

1. Log in to Solaris as a user with root privileges.
2. Check the devices to be changed to non-HDLM-managed devices.

```
# /opt/DynamicLinkManger/bin/dlnkmgr view -drv
PathID HDevName              Device LDEV
            :
000002 c5t50060E80033A9902d2 ssd27  9970/9980.15001.063C
            :
000005 c5t50060E80033A9902d2 ssd35  9970/9980.15001.063C
KAPL01001-I The HDLM command completed normally. Operation name =
view, completion time = yyyy/mm/dd hh:mm:ss
#
```

The shaded portion indicates the LDEV information about the LU that is to be removed from HDLM management.

3. If setup is being performed for Quorum devices and disk device groups in the LU that you want to remove from HDLM management, release the LU settings.

   When you unregister the SDS or SVM diskset, the diskset data will be deleted so first back up any necessary data. For details how to release settings for Quorum devices and VxVM disk groups, see 3.16.2.2 of 3.16.2.

   For details on how to release settings for SDS or SVM disksets, see 3.16.2.5 of 3.16.2 or 3.16.2.6 of 3.16.2.

4. Edit the HDLM driver unconfiguration definition file (`/opt/DynamicLinkManager/config/dlmfdrv.unconf`).

   Using a text editor, edit the definitions in the HDLM driver unconfiguration definition file (`/opt/DynamicLinkManager/config/dlmfdrv.unconf`) for the LU that you want to remove from HDLM management. Based on the LDEV information of the LU that you want to remove from HDLM management checked in step 2 (shaded portion), find the target LU in the HDLM driver unconfiguration definition file (`/opt/DynamicLinkManager/config/dlmfdrv.unconf`), and delete the hash mark (#) in front of the `LDEV` property line corresponding to the LU. Deleting the hash mark specifies that the LU indicated by that LDEV attribute is not to be managed by HDLM.

   Figure 4.8 and Figure 4.9 show examples of editing the HDLM driver unconfiguration definition file (`/opt/DynamicLinkManager/config/dlmfdrv.unconf`).

```
################## T69760 D2(9970/9980.15001.063C) ##################
### Device=ssd27 target=69760 lun=2
### Path=/pci@8,700000/SUNW,qlc@2/fp@0,0/ssd@w50060e80033a9902,2
### Device=ssd35 target=70054 lun=2
### Path=/pci@8,700000/SUNW,qlc@3/fp@0,0/ssd@w50060e80033a9938,2
#LDEV=9970/9980.15001.063C;
```

Figure 4.8 The /opt/DynamicLinkManager/config/dlmfdrv.unconf file (before editing)

```
################## T69760 D2(9970/9980.15001.063C) ##################
### Device=ssd27 target=69760 lun=2
### Path=/pci@8,700000/SUNW,qlc@2/fp@0,0/ssd@w50060e80033a9902,2
### Device=ssd35 target=70054 lun=2
### Path=/pci@8,700000/SUNW,qlc@3/fp@0,0/ssd@w50060e80033a9938,2
LDEV=9970/9980.15001.063C;
```

Figure 4.9 The /opt/DynamicLinkManager/config/dlmfdrv.unconf file (after editing)

The shaded portion indicates the text to be edited.

Note

Do not edit anything other than deleting the hash mark (#) from the beginning of the LDEV property line.

5. Apply the settings in the HDLM driver unconfiguration definition file (/opt/DynamicLinkManager/config/dlmfdrv.unconf) to the HDLM driver configuration definition file (/kernel/drv/dlmfdrv.conf).

   Execute the following command:

   ```
   # /opt/DynamicLinkManager/bin/dlmsetconf -u
   ```

6. If you are using VxVM, you need to release, from the VxVM disabled state, the device that you want to remove from HDLM management.

   Delete the target device from the /etc/vx/vxvm.exclude file.

   Figure 4.10 shows the results output by executing the dlmvxexclude utility for assisting creation of the VxVM configuration file and Figure 4.11 shows the coding in the /etc/vx/vxvm.exclude file.

```
# /opt/DynamicLinkManager/bin/dlmvxexclude
exclude_all 0
paths
c3t50060E80033A9902d0s2 /pci@8,700000/SUNW,qlc@2/fp@0,0/
ssd@w50060e80033a9902,0
c3t50060E80033A9902d1s2 /pci@8,700000/SUNW,qlc@2/fp@0,0/
ssd@w50060e80033a9902,1
c4t50060E80033A9938d0s2 /pci@8,700000/SUNW,qlc@3/fp@0,0/
ssd@w50060e80033a9938,0
c4t50060E80033A9938d1s2 /pci@8,700000/SUNW,qlc@3/fp@0,0/
ssd@w50060e80033a9938,1
c6t50060E80033A9938d0s2 /pseudo/dlmndrv@2/
dlmfdrv@w50060e80033a9938,0
c6t50060E80033A9938d1s2 /pseudo/dlmndrv@2/
dlmfdrv@w50060e80033a9938,1
#
controllers
#
product
#
pathgroups
#
KAPL10704-I dlmvxexclude completed successfully.
```

Figure 4.10 Result output by executing the dlmvxexclude utility

```
exclude_all 0
paths
c3t50060E80033A9902d0s2 /pci@8,700000/SUNW,qlc@2/fp@0,0/ssd@w50060e80033a9902,0
c3t50060E80033A9902d1s2 /pci@8,700000/SUNW,qlc@2/fp@0,0/ssd@w50060e80033a9902,1
c3t50060E80033A9902d2s2 /pci@8,700000/SUNW,qlc@2/fp@0,0/ssd@w50060e80033a9902,2
c4t50060E80033A9938d0s2 /pci@8,700000/SUNW,qlc@3/fp@0,0/ssd@w50060e80033a9938,0
c4t50060E80033A9938d1s2 /pci@8,700000/SUNW,qlc@3/fp@0,0/ssd@w50060e80033a9938,1
c4t50060E80033A9938d2s2 /pci@8,700000/SUNW,qlc@3/fp@0,0/ssd@w50060e80033a9938,2
c6t50060E80033A9938d0s2 /pseudo/dlmndrv@2/dlmfdrv@w50060e80033a9938,0
c6t50060E80033A9938d1s2 /pseudo/dlmndrv@2/dlmfdrv@w50060e80033a9938,1
c6t50060E80033A9938d2s2 /pseudo/dlmndrv@2/dlmfdrv@w50060e80033a9938,2
#
controllers
#
product
#
pathgroups
#
```

Figure 4.11 Coding in the /etc/vx/vxvm.exclude file

The shaded portion indicates the device not to be managed by HDLM and to be released from the VxVM disabled state.

For details on the `dlmvxexclude` utility, see 7.6 in Chapter 7.

7. Stop all nodes in the cluster.

   On the desired one node, execute the following command:

```
# /usr/cluster/bin/scshutdown -g0 -y
```

8. Start all nodes in the cluster simultaneously, in cluster mode.

At the ok prompt, execute the boot -r command.

Note

The following message might be output to the console window or syslog while the node is running:

```
Unable to open '/dev/rdsk/c5t50060E80033A9902d2s2'
```

This message is output when an HDLM-managed device has been changed to a non-HDLM-managed device. Perform the procedure for updating the device ID information (steps 12 to 15) so that this message is not output next time the node starts. For details, refer to the Sun Cluster documentation.

9. Execute the dlnkmgr command's view operation on all nodes in the cluster, and display the status of each program.

   The following shows an example of executing the command:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -sys
HDLM Version                : xx-xx
Service Pack Version        :
Load Balance                : on(rr)
Support Cluster             :
Elog Level                  : 3
Elog File Size (KB)         : 9900
Number Of Elog Files        : 2
Trace Level                 : 0
Trace File Size(KB)         : 1000
Number Of Trace Files       : 4
Path Health Checking        : on(30)
Auto Failback               : off
Intermittent Error Monitor  : off
HDLM Manager Ver      WakeupTime
Alive        xx-xx    yyyy/mm/dd hh:mm:ss
HDLM Alert Driver Ver      WakeupTime          ElogMem Size
Alive             xx-xx    yyyy/mm/dd hh:mm:ss
HDLM Driver  Ver      WakeupTime
Alive        xx-xx    yyyy/mm/dd hh:mm:ss
License Type Expiration
Permanent    -
KAPL01001-I The HDLM command completed normally. Operation name = view, completion time =
yyyy/mm/dd hh:mm:ss
#
```

   – Even when cluster software is used, it is not displayed in Support Cluster. Nevertheless, no problem exists with the functionality for cluster support.

10. Use the execution results from the view operation to check that the programs are operating normally.

    If HDLM Manager, HDLM Alert Driver, and HDLM Driver are all Alive, the programs are operating normally.

11. Check the LU that you removed from HDLM management.

    You can check the LU by using the format command. Figure 4.12 shows an example of executing the format command. The shaded portion represents an LU not managed by HDLM.

```
# format
Searching for disks...done

AVAILABLE DISK SELECTIONS:
       0. c1t9d0 <FUJITSU-MAP3735NP-0107 cyl 47994 alt 2 hd 4 sec 747>
          /pci@8,700000/scsi@6,1/sd@9,0
       1. c2t1d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424>
          /pci@8,600000/SUNW,qlc@4/fp@0,0/ssd@w21000004cfaf4b6e,0
       2. c2t2d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424>
          /pci@8,600000/SUNW,qlc@4/fp@0,0/ssd@w21000004cfb3a627,0
       3. c3t50060E80033A9902d2 <HITACHI-OPEN-3-SUN-2105 cyl 3336 alt 2 hd 15 sec 96>
          /pci@8,700000/SUNW,qlc@2/fp@0,0/ssd@w50060e80033a9902,2
       4. c4t50060E80033A9938d2 <HITACHI-OPEN-3-SUN-2105 cyl 3336 alt 2 hd 15 sec 96>
          /pci@8,700000/SUNW,qlc@3/fp@0,0/ssd@w50060e80033a9938,2
       5. c5t50060E80033A9902d0 <HITACHI-OPEN-3-SUN-2105 cyl 3336 alt 2 hd 15 sec 96>
          /pseudo/dlmndrv@1/dlmfdrv@w50060e80033a9902,0
       6. c5t50060E80033A9902d1 <HITACHI-OPEN-3-SUN-2105 cyl 3336 alt 2 hd 15 sec 96>
          /pseudo/dlmndrv@1/dlmfdrv@w50060e80033a9902,1
Specify disk (enter its number):
```

Figure 4.12 Example of executing the format command

12. Delete the device ID information for devices that do not exist.

    Execute the following command:

    ```
    # /usr/cluster/bin/scdidadm -C
    ```

13. Allocate the device ID.

    Execute the following command:

    ```
    # /usr/cluster/bin/scdidadm -r
    ```

    Note

    The following message might be output when the `scdidadm -r` command is executed:

    ```
    Unable to open '/dev/rdsk/c5t50060E80033A9902d2s2'
    ```

    This message is output when a logical device file for a non-HDLM-managed device exists. Delete the unnecessary logical device file. For details, refer to the Sun Cluster documentation.

14. Notify the node of the most recent device ID information.

    Execute the following command:

    ```
    # /usr/cluster/bin/scdidadm -ui
    ```

15. Update the Sun Cluster global device namespace.

    Execute the following command:

    ```
    # /usr/cluster/bin/scgdevs
    ```

16. When using a non-HDLM-managed device, as a Quorum device, set up the Quorum device in any node on the cluster.

    The following shows examples of executing the commands:

    Check the device ID of the device not to be managed by HDLM:

```
# /usr/cluster/bin/scdidadm -L
```

    Set up the Quorum device:

```
# /usr/cluster/bin/scconf -a -q globaldev=d4
```

    Execute the following command to clear the install mode of a cluster when it consists of two nodes:

```
# /usr/cluster/bin/scconf -c -q reset
```

17. Perform registration for the disk device group, create the file system, or set up the global mount as necessary.

    For details, refer to the Sun Cluster documentation.

4.4.6.2 Changing a non-HDLM-managed device to an HDLM-managed device (when the node must be restarted)

Note

On all nodes in the cluster, have HDLM driver recognize the LUs to be managed, and make the configuration the same on each node.

To change a non-HDLM-managed device to an HDLM-managed device (when the node needs to be restarted):

1. Log in to Solaris as a user with root privileges.
2. Check the LU to be managed by HDLM.

   You can check the LU by using the `format` command. Figure 4.13 shows an example of executing the `format` command. The shaded portion represents the LU to be managed by HDLM.

```
# format
Searching for disks...done

AVAILABLE DISK SELECTIONS:
       0. c1t9d0 <FUJITSU-MAP3735NP-0107 cyl 47994 alt 2 hd 4 sec 747>
          /pci@8,700000/scsi@6,1/sd@9,0
       1. c2t1d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424>
          /pci@8,600000/SUNW,qlc@4/fp@0,0/ssd@w21000004cfaf4b6e,0
       2. c2t2d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424>
          /pci@8,600000/SUNW,qlc@4/fp@0,0/ssd@w21000004cfb3a627,0
       3. c3t50060E80033A9902d2 <HITACHI-OPEN-3-SUN-2105 cyl 3336 alt 2 hd 15 sec 96>
          /pci@8,700000/SUNW,qlc@2/fp@0,0/ssd@w50060e80033a9902,2
       4. c4t50060E80033A9938d2 <HITACHI-OPEN-3-SUN-2105 cyl 3336 alt 2 hd 15 sec 96>
          /pci@8,700000/SUNW,qlc@3/fp@0,0/ssd@w50060e80033a9938,2
       5. c5t50060E80033A9902d0 <HITACHI-OPEN-3-SUN-2105 cyl 3336 alt 2 hd 15 sec 96>
          /pseudo/dlmndrv@1/dlmfdrv@w50060e80033a9902,0
       6. c5t50060E80033A9902d1 <HITACHI-OPEN-3-SUN-2105 cyl 3336 alt 2 hd 15 sec 96>
          /pseudo/dlmndrv@1/dlmfdrv@w50060e80033a9902,1
Specify disk (enter its number):
```

Figure 4.13 Example of executing the format command

3. If setup is being performed for Quorum devices and disk device groups in the LU to be managed by HDLM, release the LU settings.

   When you unregister the SDS or SVM diskset, the diskset data will be deleted so first back up any necessary data. For details on how to release settings for Quorum devices and VxVM disk groups, see 3.16.2.2 of 3.16.2.

   For details on how to release SDS or SVM disksets, see 3.16.2.5 of 3.16.2 or 3.16.2.6 of 3.16.2.

4. Make sure that registration is performed for the key used in disk reservation for the LU to be managed by HDLM. Execute the following command on any node:

   ```
   # /usr/cluster/lib/sc/scsi -c inkeys -d HDLM
   the-name-of-the-logical-device-file-of-slice-2-of-the-LU-to-be-managed-by-HDLM
   ```

   The following shows an example of executing the command:

   ```
   # /usr/cluster/lib/sc/scsi -c inkeys -d /dev/rdsk/c3t50060E80033A9902d2s2
   ```

   If the key is displayed, perform the steps beginning with step 5. Even when you release the setup for Quorum devices and disk device groups, since the key used for disk reservation will remain registered, you will need to restart the node in non-cluster mode and unregister the key.

   If the key is not displayed, perform the steps beginning with step 9.

5. Simultaneously stop all nodes in the cluster.

   Execute the following command on any node:

   ```
   # /usr/cluster/bin/scshutdown -g0 -y
   ```

6. Start all nodes in the cluster in non-cluster mode.

   At the `ok` prompt, execute the `boot -x` command.

7. Unregister the key for the LU to be managed by HDLM.

   Execute the following command on any node:

```
# /usr/cluster/lib/sc/scsi -c scrub -d HDLM
name-of-the-logical-device-file-of-slice-2-of-the-LU-to-be-managed-by HDLM
```

   The following shows an example of executing the command:

```
# /usr/cluster/lib/sc/scsi -c scrub -d /dev/rdsk/c3t50060E80033A9902d2s2
```

8. Make sure that the key for the LU to be managed by HDLM has been unregistered.

   Perform step 4 again on any node, and check that the key is not displayed.

9. Edit the HDLM driver unconfiguration definition file (`/opt/DynamicLinkManager/config/dlmfdrv.unconf`).

   Using a text editor, edit the definitions in the HDLM driver unconfiguration definition file (`/opt/DynamicLinkManager/config/dlmfdrv.unconf`) for the LU to be managed by HDLM. Based on the path information of the LU to be managed by HDLM checked in step 2 (shaded portion), find the correct LU in the HDLM driver unconfiguration definition file (`/opt/DynamicLinkManager/config/dlmfdrv.unconf`), and add a hash mark (#) in front of the `LDEV` property line corresponding to the LU. Adding the hash mark specifies that the LU indicated by that LDEV attribute is to be managed by HDLM.

   Figure 4.14 and Figure 4.15 show examples of editing the HDLM driver unconfiguration definition file (`/opt/DynamicLinkManager/config/dlmfdrv.unconf`).

```
################### T69760 D2(9970/9980.15001.063C) ###################
### Device=ssd27 target=69760 lun=2
### Path=/pci@8,700000/SUNW,qlc@2/fp@0,0/ssd@w50060e80033a9902,2
### Device=ssd35 target=70054 lun=2
### Path=/pci@8,700000/SUNW,qlc@3/fp@0,0/ssd@w50060e80033a9938,2
LDEV=9970/9980.15001.063C;
```

Figure 4.14 The /opt/DynamicLinkManager/config/dlmfdrv.unconf file (before editing)

```
################### T69760 D2(9970/9980.15001.063C) ###################
### Device=ssd27 target=69760 lun=2
### Path=/pci@8,700000/SUNW,qlc@2/fp@0,0/ssd@w50060e80033a9902,2
### Device=ssd35 target=70054 lun=2
### Path=/pci@8,700000/SUNW,qlc@3/fp@0,0/ssd@w50060e80033a9938,2
#LDEV=9970/9980.15001.063C;
```

Figure 4.15 The /opt/DynamicLinkManager/config/dlmfdrv.unconf file (after editing)

   The shaded portion indicates the text to be edited.

   Note

   Do not make any modifications other than adding or deleting the hash mark (#) in front of the `LDEV` property line corresponding to the LU.

10. Apply the settings in the HDLM driver unconfiguration definition file (/opt/DynamicLinkManager/config/dlmfdrv.unconf) to the HDLM driver configuration definition file (/kernel/drv/dlmfdrv.conf):

    Execute the following command:

```
# /opt/DynamicLinkManager/bin/dlmsetconf -u
```

11. If you are using VxVM, you need to disable, from VxVM, any devices (such as sd or ssd devices) other than the devices corresponding to the added LU.

    Specify settings as shown in 3.12.3.2 of 3.12.3 or 3.12.3.3 of 3.12.3.

12. Stop all nodes in the cluster.

    If the key is displayed in step 4, and the node is running in non-cluster mode in step 6, execute the following command on all nodes:

```
# shutdown -g0 -y -i0
```

    When running nodes in cluster mode, execute the following command on any node:

```
# scshutdown -g0 -y
```

13. Start all nodes in the cluster simultaneously, in cluster mode.

    At the ok prompt, execute the boot -r command.

    Note

    The following message might be output to the console window or syslog while the node is running:

```
Could not read symbolic link for:
/dev/rdsk/c3t50060E80033A9902d2s2 path not loaded
```

    This message is output when HDLM deletes a logical device file for an sd or ssd device managed by HDLM. Perform the procedure for updating the device ID information (steps 17 to 20) so that this message is not output next time the node starts. For details, refer to the Sun Cluster documentation.

14. Execute the dlnkmgr command's view operation on all nodes in the cluster, and display the status of each program.

    The following shows an example of executing the command:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -sys
HDLM Version               : xx-xx
Service Pack Version       :
Load Balance               : on(rr)
Support Cluster            :
Elog Level                 : 3
Elog File Size (KB)        : 9900
Number Of Elog Files       : 2
Trace Level                : 0
Trace File Size(KB)        : 1000
Number Of Trace Files      : 4
```

```
Path Health Checking         : on(30)
Auto Failback               : off
Intermittent Error Monitor   : off
HDLM Manager Ver      WakeupTime
Alive        xx-xx    yyyy/mm/dd hh:mm:ss
HDLM Alert Driver Ver      WakeupTime          ElogMem Size
Alive             xx-xx    yyyy/mm/dd hh:mm:ss
HDLM Driver  Ver      WakeupTime
Alive        xx-xx    yyyy/mm/dd hh:mm:ss
License Type Expiration
Permanent    -
KAPL01001-I The HDLM command completed normally. Operation name = view, completion time =
yyyy/mm/dd hh:mm:ss
#
```

– Even when cluster software is used, it is not displayed in `Support Cluster`. Nevertheless, no problem exists with the functionality for cluster support.

15. Use the execution results from the `view` operation to check that the programs are operating normally.

    If `HDLM Manager`, `HDLM Alert Driver`, and `HDLM Driver` are all `Alive`, the programs are operating normally.

16. Check the status of the path for accessing the LU to be managed by HDLM.

    The following shows an example of executing the command. The shaded portion is the path accessing the LU to be managed by HDLM.

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -path
Paths:000006 OnlinePaths:000006
PathStatus   IO-Count    IO-Errors
Online       0           0

PathID PathName                        DskName
  iLU           ChaPort Status    Type IO-Count  IO-Errors DNum HDevName
000000 0000.0000.0000000000011080.0000 HITACHI .OPEN-3      -SUN.15001
  063A          4j      Online    Own         0          0    - c5t50060E80033A9902d0
000001 0000.0000.0000000000011080.0001 HITACHI .OPEN-3      -SUN.15001
  063B          4j      Online    Own         0          0    - c5t50060E80033A9902d1
000002 0000.0000.0000000000011080.0002 HITACHI .OPEN-3      -SUN.15001
  063C          4j      Online    Own         0          0    - c5t50060E80033A9902d2
000003 0001.0000.00000000000111A6.0000 HITACHI .OPEN-3      -SUN.15001
  063A          1C      Online    Own         0          0    - c5t50060E80033A9902d0
000004 0001.0000.00000000000111A6.0001 HITACHI .OPEN-3      -SUN.15001
  063B          1C      Online    Own         0          0    - c5t50060E80033A9902d1
000005 0001.0000.00000000000111A6.0002 HITACHI .OPEN-3      -SUN.15001
  063C          1C      Online    Own         0          0    - c5t50060E80033A9902d2
KAPL01001-I The HDLM command completed normally. Operation name = view, completion time
= yyyy/mm/dd hh:mm:ss
#
```

17. Delete the device ID information for devices that do not exist.

    Execute the following command:

```
# /usr/cluster/bin/scdidadm -C
```

18. Allocate the device ID.

    Execute the following command:

```
# /usr/cluster/bin/scdidadm -r
```

19. Notify the node of the most recent device ID information.

    Execute the following command:

```
# /usr/cluster/bin/scdidadm -ui
```

20. Update the Sun Cluster global device namespace.

    Execute the following command:

```
# /usr/cluster/bin/scgdevs
```

21. When using an HDLM-managed device, as a Quorum device, set the Quorum device on any node in the cluster.

    The following shows examples of executing the commands:

    Check the device ID of the HDLM device:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -path
# /usr/cluster/bin/scdidadm -L
```

    Set up the Quorum device:

```
# /usr/cluster/bin/scconf -a -q globaldev=d4
```

    Execute the following command to clear the install mode of a cluster when it consists of two nodes:

```
# /usr/cluster/bin/scconf -c -q reset
```

22. Execute the following command to view the HDevName and iLU columns:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -path
```

    Check the correspondence between `HDevName` (the name of the logical device file of the HDLM device) and `iLU` (the internal LU number) from the execution results of the command, compare the correspondence with the correspondence that took place before the HDLM installation, and then re-register HDLM devices with the programs that will use them.

    If you want to use volume management software to create a file system on an HDLM-managed device, skip to step 24.

23. Perform registration for the disk device group.

    If you are using VxVM, see 3.12.5.

    If you are using SDS, see 3.13.2.2 of 3.13.2.

    If you are using SVM, see 3.14.2.2 of 3.14.2.

24. Create the file system and set up the global mount as necessary.

    For details, refer to the Sun Cluster documentation.

4.4.6.3 Changing a non-HDLM-managed device to an HDLM-managed device (for dynamic reconfiguration)

Note

- Dynamic reconfiguration cannot be performed when using VxVM because the host must be restarted in order to use (in VxVM) an LU that was changed from a non-HDLM-managed device to an HDLM-managed device. Perform the procedure 4.4.6.2.
- For details on the precautions for the `dlmcfgmgr` utility for managing the HDLM configuration, see 7.3.

To change a non-HDLM-managed device to an HDLM-managed device:

1. Log in to Solaris as a user with root privileges.
2. Check the LU to be managed by HDLM.

   You can check the LU by using the `format` command. Figure 4.16 shows an example of executing the `format` command. The shaded portion represents the LU to be managed by HDLM.

```
# format
Searching for disks...done

AVAILABLE DISK SELECTIONS:
       0. c1t9d0 <FUJITSU-MAP3735NP-0107 cyl 47994 alt 2 hd 4 sec 747>
          /pci@8,700000/scsi@6,1/sd@9,0
       1. c2t1d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424>
          /pci@8,600000/SUNW,qlc@4/fp@0,0/ssd@w21000004cfaf4b6e,0
       2. c2t2d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424>
          /pci@8,600000/SUNW,qlc@4/fp@0,0/ssd@w21000004cfb3a627,0
       3. c3t50060E80033A9902d2 <HITACHI-OPEN-3-SUN-2105 cyl 3336 alt 2 hd 15 sec 96>
          /pci@8,700000/SUNW,qlc@2/fp@0,0/ssd@w50060e80033a9902,2
       4. c4t50060E80033A9938d2 <HITACHI-OPEN-3-SUN-2105 cyl 3336 alt 2 hd 15 sec 96>
          /pci@8,700000/SUNW,qlc@3/fp@0,0/ssd@w50060e80033a9938,2
       5. c5t50060E80033A9902d0 <HITACHI-OPEN-3-SUN-2105 cyl 3336 alt 2 hd 15 sec 96>
          /pseudo/dlmndrv@1/dlmfdrv@w50060e80033a9902,0
       6. c5t50060E80033A9902d1 <HITACHI-OPEN-3-SUN-2105 cyl 3336 alt 2 hd 15 sec 96>
          /pseudo/dlmndrv@1/dlmfdrv@w50060e80033a9902,1
Specify disk (enter its number):
```

Figure 4.16 Example of executing the format command

3. If setup is performed for Quorum devices and disk device groups in the LU to be managed by HDLM, release the LU settings.

   When you unregister the SDS or SVM diskset, the diskset data will be deleted so first back up any necessary data. For details on how to release settings for Quorum devices and VxVM disk groups, see 3.16.2.2 of 3.16.2.

For details on how to release SDS or SVM disksets, see 3.16.2.5 of 3.16.2 or 3.16.2.6 of 3.16.2.

4. Make sure that registration is performed for the key used in disk reservation for the LU to be managed by HDLM.

   Execute the following command on any node:

```
# /usr/cluster/lib/sc/scsi -c inkeys -d
the-name-of-the-logical-device-file-of-slice-2-of-the-LU-to-be-managed-by-HDML
```

   The following shows an example of executing the command:

```
# /usr/cluster/lib/sc/scsi -c inkeys -d /dev/rdsk/c3t50060E80033A9902d2s2
```

   If the key is displayed, place the LU under HDLM management from step 5 of 4.4.6.2. Even when you release the setup for Quorum devices and disk device groups, since the key used for disk reservation will remain registered, you will need to restart the node in non-cluster mode and unregister the key.

   If the key is not displayed, perform the steps beginning with step 5.

5. Edit the HDLM driver unconfiguration definition file (`/opt/DynamicLinkManager/config/dlmfdrv.unconf`).

   Using a text editor, edit the definitions in the HDLM driver unconfiguration definition file (`/opt/DynamicLinkManager/config/dlmfdrv.unconf`) for the LU to be managed by HDLM. Based on the path information of the LU to be managed by HDLM checked in step 2 (shaded portion), find the correct LU in the HDLM driver unconfiguration definition file (`/opt/DynamicLinkManager/config/dlmfdrv.unconf`), and add a hash mark (#) in front of the `LDEV` property line corresponding to the LU. Adding the hash mark specifies that the LU indicated by that LDEV attribute is to be managed by HDLM.

   Figure 4.17 and Figure 4.18 show examples of editing the HDLM driver unconfiguration definition file (`/opt/DynamicLinkManager/config/dlmfdrv.unconf`).

```
################### T69760 D2(9970/9980.15001.063C) ###################
### Device=ssd27 target=69760 lun=2
### Path=/pci@8,700000/SUNW,qlc@2/fp@0,0/ssd@w50060e80033a9902,2
### Device=ssd35 target=70054 lun=2
### Path=/pci@8,700000/SUNW,qlc@3/fp@0,0/ssd@w50060e80033a9938,2
LDEV=9970/9980.15001.063C;
```

Figure 4.17 The/opt/DynamicLinkManager/config/dlmfdrv.unconf file (before editing)

```
################### T69760 D2(9970/9980.15001.063C) ###################
### Device=ssd27 target=69760 lun=2
### Path=/pci@8,700000/SUNW,qlc@2/fp@0,0/ssd@w50060e80033a9902,2
### Device=ssd35 target=70054 lun=2
### Path=/pci@8,700000/SUNW,qlc@3/fp@0,0/ssd@w50060e80033a9938,2
#LDEV=9970/9980.15001.063C;
```

Figure 4.18 The/opt/DynamicLinkManager/config/dlmfdrv.unconf file (after editing)

The shaded portion indicates indicate the text to be edited.

Note

Do not make any modifications other than adding or deleting the hash mark (#) in front of the LDEV property line corresponding to the LU.

6. Apply the settings in the HDLM driver unconfiguration definition file (/opt/DynamicLinkManager/config/dlmfdrv.unconf) to the HDLM driver configuration definition file (/kernel/drv/dlmfdrv.conf):

   Execute the following command:

```
# /opt/DynamicLinkManager/bin/dlmsetconf -u
```

7. Execute the dlmcfgmgr utility shown below to make the HDLM driver recognize the LU to be managed by HDLM:

```
# /opt/DynamicLinkManager/bin/dlmcfgmgr -a
```

8. Check the status of the path accessing the LU to be managed by HDLM.

   The following shows an example of executing the command. The shaded portion is the path to access the LU to be managed by HDLM.

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -path
Paths:000006 OnlinePaths:000006
PathStatus   IO-Count    IO-Errors
Online       0            0

PathID PathName                          DskName
  iLU           ChaPort Status   Type IO-Count  IO-Errors  DNum HDevName
000000 0000.0000.0000000000011080.0000 HITACHI .OPEN-3      -SUN.15001
  063A          4j      Online   Own         0          0     - c5t50060E80033A9902d0
000001 0000.0000.0000000000011080.0001 HITACHI .OPEN-3      -SUN.15001
  063B          4j      Online   Own         0          0     - c5t50060E80033A9902d1
000002 0001.0000.00000000000111A6.0000 HITACHI .OPEN-3      -SUN.15001
  063A          1C      Online   Own         0          0     - c5t50060E80033A9902d0
000003 0001.0000.00000000000111A6.0001 HITACHI .OPEN-3      -SUN.15001
  063B          1C      Online   Own         0          0     - c5t50060E80033A9902d1
000004 0000.0000.0000000000011080.0002 HITACHI .OPEN-3      -SUN.15001
  063C          4j      Online   Own         0          0     - c5t50060E80033A9902d2
000005 0001.0000.00000000000111A6.0002 HITACHI .OPEN-3      -SUN.15001
  063C          1C      Online   Own         0          0     - c5t50060E80033A9902d2
KAPL01001-I The HDLM command completed normally. Operation name = view, completion time
= yyyy/mm/dd hh:mm:ss
```

9. Delete the device ID information for devices that do not exist.

   Execute the following command:

```
# /usr/cluster/bin/scdidadm -C
```

10. Allocate the device ID.

    Execute the following command:

```
# /usr/cluster/bin/scdidadm -r
```

11. Notify the node of the most recent device ID information.

    Execute the following command:

```
# /usr/cluster/bin/scdidadm -ui
```

12. Update the Sun Cluster global device namespace.

    Execute the following command:

```
# /usr/cluster/bin/scgdevs
```

13. When using an HDLM-managed device, as a Quorum device, set the Quorum device on any node in the cluster.

    The following shows examples of executing the commands:

    Check the device ID of the HDLM device:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -path
# /usr/cluster/bin/scdidadm -L
```

    Set up the Quorum device:

```
# /usr/cluster/bin/scconf -a -q globaldev=d4
```

    Execute the following command to clear the install mode of a cluster when it consists of two nodes:

```
# /usr/cluster/bin/scconf -c -q reset
```

14. Execute the following command to view the HDevName and iLU columns:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -path
```

    Check the correspondence between `HDevName` (the name of the logical device file of the HDLM device) and `iLU` (the internal LU number) from the execution results of the command, compare the correspondence with the correspondence that took place before the HDLM installation, and then re-register HDLM devices with the programs that will use them.

    If you want to use volume management software to create a file system on an HDLM-managed device, skip to step 16.

15. Perform registration for the disk device group.

    If you are using VxVM, see 3.12.5.

    If you are using SDS, see 3.13.2.2 of 3.13.2.

    If you are using SVM, see 3.14.2.2 of 3.14.2.

16. Create the file system and set up the global mount as necessary.
    For details, refer to the Sun Cluster documentation.

Chapter 5 Troubleshooting

This chapter describes how to check HDLM error information, and then how to take action if an error occurs in HDLM. The actions you should take are described separately for path errors, HDLM program errors, and for other types of errors.

- 5.1 Information collected by using the DLMgetras utility for collecting HDLM error information
- 5.2 Checking Error Information in Messages
- 5.3 Actions Taken for a Path Error
- 5.4 Actions to Take for a Path Error in a Boot Disk Environment
- 5.5 Actions Taken for a Program Error
- 5.6 Actions Taken for Other Errors

5.1 Information collected by using the DLMgetras utility for collecting HDLM error information

Immediately after an error occurs, execute the `DLMgetras` utility to collect HDLM error information, because restarting the machine may delete existing error information collected by the `DLMgetras` utility. For details on the `DLMgetras` utility and the error information it collects, see 7.2.

5.2 CheckingError Information in Messages

When you want to configure the system so that HDLM messages are output to syslog, specify `user` for the name of the system function defined in the `/etc/syslog.conf` file. In the following example, the system function name is `user`, and messages at the Information level or higher are output to the `/etc/syslog.conf` file:

```
user.info           /tmp/syslog.user.log
```

You can check path errors by referring to the KAPL08*xxx* messages that are output to syslog.

To obtain detailed information about the failed path, check the execution results of the `view` operation as indicated by the error message.

For details on the `view` operation, see 6.7.

The following is an example of a message:

```
KAPL08022-E Error in path occurred. ErrorCode = aa....aa,PathID = bb...bb,PathName =
cc...cc.dd...dd.ee...ee.ff...ff,DNum = gg...gg,HDevName = hh...hh
```

The elements of the message are explained below.

`ErrorCode`

The error number generated when Solaris detected the path error.

PathID

The ID assigned to a path. This ID is called the *AutoPATH_ID*. AutoPATH_IDs are re-assigned every time the host is restarted. When you want to add a new LU without restarting the host, AutoPATH_IDs will be re-assigned to each path of the LU when you execute the dlmcfgmgr utility for managing the HDLM configuration.

The path ID is also the same as the PathID displayed by the dlnkmgr command's view operation. For details on the view operation, see 6.7.

PathName

The path name, which indicates a physical path. When you modify the system configuration or replace a hardware item, you should check the path names to identify the physical paths that will be affected by the change.

A path name consists of the following four elements, separated by periods:

- Host port number (hexadecimal)
- Bus number (hexadecimal)
- Target ID (hexadecimal)
- Host LU number (hexadecimal)

The path name is also the same as the PathName displayed by the dlnkmgr command's view operation. For details on the path name, see 6.7.

DNum

A Dev number.

A hyphen (-) is displayed.

This is the same as the DNum displayed by the view operation. For details on the view operation, see 6.7.

HDevName

The name of the host device.

The logical device file name of the HDLM device which the path accesses, minus the slice number, is displayed in the format c*U*t*X*d*Y* (for example, c6t5d0).

This is the same as the HDevName displayed by the view operation. For details on the view operation, see 6.7.

5.3 Actions Taken for a Path Error

When a path error is detected, HDLM performs failover for the path and outputs the KAPL08022-E message. This message indicates that an error has occurred in the components that make up the path. Figure 5.1 indicates these components.

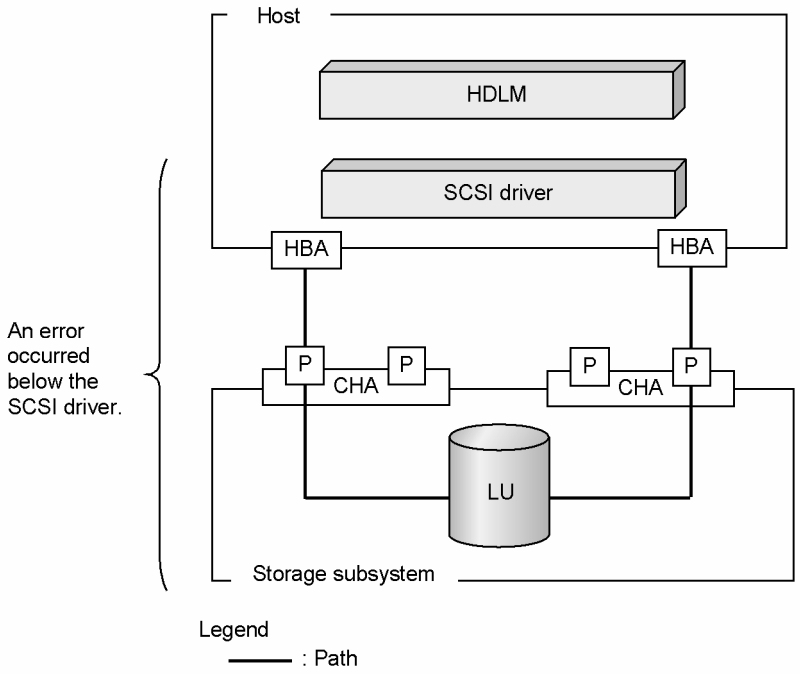

Figure 5.1 Error Location When the KAPL08022-E Message is Output

Figure 5.2 shows the troubleshooting procedure when the KAPL08022-E message is output.

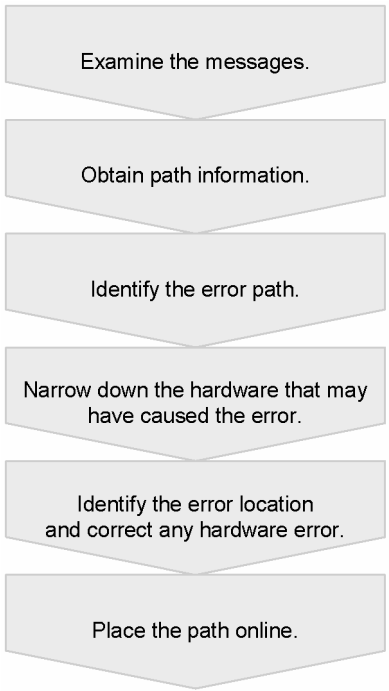

Figure 5.2 Troubleshooting Procedure When a Path Error Occurs

If an error occurs in a boot disk environment on a path connected to the boot disk, perform the procedure described in 5.4.

The following shows the procedure for using the dlnkmgr command to handle a path error.

5.3.1 Examine the Messages

Examine the message that is output to syslog in the host by using applications or tools for monitoring messages. If the KAPL08022-E message is output, view the message to check the path in which the error occurs. For details on the message, see 5.2.

5.3.2 Obtain Path Information

Obtain path information to narrow down the hardware in which the error occurred.

Execute the following command:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -path -iem > pathinfo.txt
```

pathinfo.txt is the redirection-output file name. Use a file name that matches your environment.

5.3.3 Identify the Error Path

Check the obtained path information to find the error path. In the Status column, the error path has the status Offline(E) or Online(E).

5.3.4 Narrow Down the Hardware That May Have Caused the Error

Check the DskName, iLU, and ChaPort columns for the error path to narrow down the hardware that may be the cause of the error. To physically identify the hardware corresponding to DskName, iLU, and ChaPort, use the information provided by the storage-subsystem management program.

5.3.5 Identify the Error Location and Correct any Hardware Error

Use the Solaris and hardware management tools to identify the error location, and then take appropriate, corrective action. For hardware maintenance, contact your hardware vendor or maintenance company if there is a maintenance contract.

5.3.6 Place the Path Online

After the path recovers from the error, use the dlnkmgr command's online operation to place online the path that went offline due to the error. For details on the online operation, see 6.5. Execute the following command:

```
# /opt/DynamicLinkManager/bin/dlnkmgr online
```

Executing this command places all the offline paths online.

If any path cannot be placed online due to an error, the KAPL01039-W message appears. To ignore such paths and to continue processing, type `y`. To cancel the processing, type `n`.

Recheck the status of the paths that cannot be placed online, and take appropriate, corrective action.

Note

Even if you execute the `dlnkmgr` command's `online` operation (with the `-pathid` parameter) for the paths in the `Offline(E)` status, some paths may not be placed online. In this case, use the`dlnkmgr` command's `offline` operation to change those paths from the `Offline(E)` status to the `Offline(C)` status. Then, execute the `online` operation. The system automatically checks whether the target paths are available, puts available paths into the `Online` status, and puts unavailable paths into the `Offline(E)` status. The paths in the `Offline(E)` status have not recovered from an error yet. Take corrective action for the error, and then execute the `online` operation again. For details on the `online` operation, see 6.5. For details on the `offline` operation, see 6.4.

5.4 Actions to Take for a Path Error in a Boot Disk Environment

This section describes the action to take if an error occurs in a boot disk environment on a path connecting to the boot disk.

5.4.1 Path Errors During Boot Processing

The procedure for investigating the cause of an error, and the appropriate action to take, will differ according to the stage of the boot processing during which the path error occurred. The following describes error investigation methods at each stage of boot processing.

5.4.1.1 When a Path Error Occurs at the Initial Stage of Boot Processing

The initial stage of boot processing is the period before the `root` file system's mount device (one of the paths to the boot disk) specified as the boot device in the OBP is switched (remounted) to the device specified in the `rootdev` parameter. At this stage, accessing the boot disk does not involve the HDLM driver: The boot disk is accessed via the path specified as the boot device in the OBP. If an error occurs on this path at this stage, boot processing will be cancelled. In this case, investigate the error and take appropriate action, referring to the following information:

- The message output to the console by the OBP
- The message output to the console by Solaris
- The message output to the console by the HBA driver

In Solaris 8 or 9, if an error occurs only on the path specified as the boot device in the OBP, you can boot the system by specifying, in the OBP, a boot device name with a different path. In Solaris 10, reboot after recovering from the path error.

5.4.1.2 When a Path Error Occurs After the HDLM Driver Starts Path Processing

At completion of the initial stage of boot processing, the HDLM driver starts accessing the boot disk using multiple paths. If a path error occurs at this stage, recovery action will depend on whether or not the HDLM manager has started.

When a path error occurs before the HDLM manager starts

If boot processing is cancelled due to errors on all paths to the boot disk

HDLM does not output the KAPL08022-E message on detection of a path error. Investigate the error and take appropriate action, referring to the following information:

- The message output to the console by Solaris
- The message output to the console by the HBA driver

If an error occurs on one or more of the paths to the boot disk

HDLM outputs the KAPL08022-E message to the console, syslog, and error log after the HDLM manager has started. Investigate the error and take appropriate action after the HDLM manager has started, referring to the following information:

- The KAPL08022-E message output by HDLM to the console, syslog, and error log
- The message output to the console by Solaris
- The message output to the console by the HBA driver

When a path error occurs after the HDLM manager starts

On detection of a path error, HDLM immediately outputs the KAPL08022-E message to the console, syslog, and error log. Therefore, even if boot processing is cancelled at this stage due to errors on all paths to the boot disk, investigate the error and take appropriate action, referring to the following information:

- The KAPL08022-E message output by HDLM to the console, syslog, and error log
- The message output to the console and syslog by Solaris
- The message output to the console and syslog by the HBA driver

5.4.2 Path Errors After Boot Processing Completes

Take the same action as for a local boot disk environment. For details, see 5.3.

5.5 Actions Taken for a Program Error

The following describes the troubleshooting procedure for handling errors that occur in an HDLM program. Figure 5.3 shows the troubleshooting procedure.

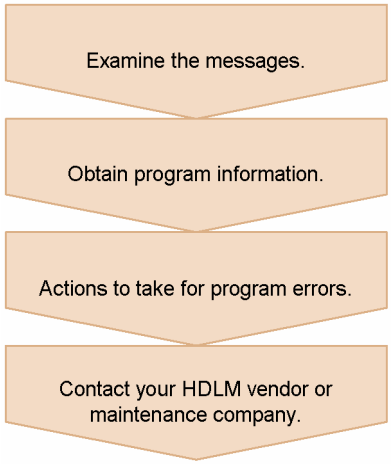

Figure 5.3 Troubleshooting Procedure When a Program Error Occurs

The following shows the procedure for handling a program error by using the `dlnkmgr` command.

5.5.1 Examine the Messages

Examine the message that is output to syslog in the host. If an error occurs in an HDLM program, a message other than KAPL08*xxx* is output to syslog. Check the output message. message. Messages with error level `E` (Error) or higher require corrective action.

5.5.2 Obtain Program Information

Obtain the information to be reported to your HDLM vendor or maintenance company (if there is a maintenance contract for HDLM).

Use the `DLMgetras` utility for collecting HDLM error information to collect the error information. For details on the `DLMgetras` utility and the information it collects, see 7.2.

Some of the information collected by the `DLMgetras` utility might be cleared when the system is restarted. When an error occurs, as soon as possible execute the `DLMgetras` utility.

5.5.3 Actions to Take for Program Errors

Take the action recommended for the message in Chapter 8.

If the error occurs again after you take the corrective action, use the `dlnkmgr` command's `view` operation to check the status of the HDLM program, and then take the appropriate corrective action. For details on the `view` operation, see 6.7.

Execute the following command:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -sys
```

If the KAPL01012-E message appears as a result of executing the command

The following shows the KAPL01012-E message:

```
KAPL01012-E Could not connect the HDLM manager. Operation name = view
```

If this message appears, start the HDLM manager.

For details on how to start the HDLM manager, see 4.3.1.

If the KAPL01013-E message appears as a result of executing the command

The following shows the KAPL01013-E message:

```
KAPL01013-E An error occurred in internal processing of the HDLM command. Operation name
= view, details = aa...aa
```

aa...aa indicates character string. If this message appears, restart the host.

If the same error re-occurs after you take action, go to the following section *(4) Contact your HDLM vendor or maintenance company*.

5.5.4 Contact Your HDLM Vendor or Maintenance Company

If the error cannot be resolved, contact your HDLM vendor or maintenance company if there is a maintenance contract for HDLM, and report the information collected by the `DLMgetras` utility for collecting HDLM error information.

5.6 Actions Taken for Other Errors

When the cause of an error may be related to HDLM but is neither a path error nor an HDLM program error, execute the `DLMgetras` utility for collecting HDLM error information, and then report the collected information to the HDLM vendor or maintenance company if there is a maintenance contract for HDLM. For details on the `DLMgetras` utility and the information it collects, see 7.2.

Chapter 6 Command Reference

This chapter describes the HDLM command (dlnkmgr) and its operations.

- 6.1 Overview of the HDLM Command (dlnkmgr)
- 6.2 clear (Clears the Path Statistics to the Initial Value)
- 6.3 help (Displays the Operation Format)
- 6.4 offline (Places a Path or Paths Offline)
- 6.5 online (Places a Path or Paths Online)
- 6.6 set (Sets the Operating Environment)
- 6.7 view (Displays Information)
- 6.8 monitor (Displays I/O information at a specified interval)

6.1 Overview of the HDLM Command (dlnkmgr)

This section describes how to specify the HDLM command dlnkmgr and its subcommands (called *operations* in HDLM). Note that messages might refer to this command and the operations as the *HDLM command*.

Command format

Enter the command using the following format:

```
dlnkmgr operation-name [parameter [parameter-value]]
```

dlnkmgr

The command name.

operation-name

The type of operation entered after dlnkmgr.

parameter

A value required for an operation.

parameter-value

A value required for a parameter.

Operations of the dlnkmgr command

Table 6.1 shows the operations of the dlnkmgr command and their functions.

Table 6.1 Operations of the dlnkmgr Command

Operation	Functions
clear	Initializes(0) the statistics (I/O count and I/O errors) of all paths managed by the HDLM system. For details, see 6.2.

Operation	Functions
help	Displays the format of the operation used for HDLM. For details, see 6.3.
offline	Places offline an online path or paths. For details, see 6.4.
online	Places online an offline path or paths. For details, see 6.5.
set	Sets the HDLM operating environment. For details, see 6.6.
view	Displays HDLM program information, path information, LU information, HBA port information, CHA port information, and corresponding information about each HDLM device, sd or ssd device, and LDEV. For details, see 6.7.
monitor	Displays I/O information at a specified interval. For details, see 6.8.

Note

- Execute the command as a user with root privileges.
- To specify a value containing a space in a parameter, enclose the entire value in double quotation marks (").

6.2 clear (Clears the Path Statistics to the Initial Value)

The dlnkmgr command's clear operation clears the statistics (I/O count and I/O errors) of all paths that are managed by the HDLM system to the initial value.

6.2.1 Format

6.2.1.1 To set the path statistics to 0

```
/opt/DynamicLinkManager/bin/dlnkmgr clear -pdst [-s]
```

6.2.1.2 To display the format of the clear operation

```
/opt/DynamicLinkManager/bin/dlnkmgr clear -help
```

6.2.2 Parameters

6.2.2.1 To set the path statistics to 0

-pdst

Clears statistics (I/O count and I/O errors) of all paths managed by HDLM to the initial value (0).

Example

```
# /opt/DynamicLinkManager/bin/dlnkmgr clear -pdst
```

```
KAPL01049-I Would you like to execute the operation? Operation name = clear [y/n]:y
KAPL01001-I The HDLM command completed normally. Operation name = clear, completion time
= yyyy/mm/dd hh:mm:ss
#
```

-s

Executes the command without displaying a message asking for confirmation from the user. Specify this parameter if you want to skip the response to the confirmation message: for example, when you want to execute the command in a shell script or batch file.

Example

```
# /opt/DynamicLinkManager/bin/dlnkmgr clear -pdst -s
KAPL01001-I The HDLM command completed normally. Operation name = clear, completion time
= yyyy/mm/dd hh:mm:ss
#
```

6.2.2.2 To display the format of the clear operation

-help

Displays the format of the `clear` operation.

Example

```
# /opt/DynamicLinkManager/bin/dlnkmgr clear -help
clear:
  Format
    dlnkmgr clear -pdst [-s]
KAPL01001-I The HDLM command completed normally. Operation name = clear, completion time
= yyyy/mm/dd hh:mm:ss
#
```

6.3 help (Displays the Operation Format)

The `dlnkmgr` command's `help` operation displays the list of operations available for the `dlnkmgr` command, or the format of individual operations.

6.3.1 Format

```
/opt/DynamicLinkManager/bin/dlnkmgr help
   [operation-name] [operation-name] ...
```

6.3.2 Parameter

operation-name

Specify the name of the HDLM command operation whose format you want to know.

You can specify multiple operation names. If you specify multiple operation names, the formats will be displayed in the order you specified operation names. The following shows the operation names that you can specify:

- clear

- help
- online
- offline
- set
- view
- monitor

If you do not specify any operation names, the help operation displays the names of all operations available for the dlnkmgr command.

Examples

Example 1

The following example shows how to display the names of all the operations available for the dlnkmgr command.

```
# /opt/DynamicLinkManager/bin/dlnkmgr help
dlnkmgr:
  Format
    dlnkmgr { clear | help | offline | online | set | view | monitor }
KAPL01001-I The HDLM command completed normally. Operation name = help, completion time =
yyyy/mm/dd hh:mm:ss
#
```

Example 2

The following example shows how to display the formats of multiple operations.

```
# /opt/DynamicLinkManager/bin/dlnkmgr help online offline help
online:
  Format
    dlnkmgr online [-path]
                    { -hba HBAPortNumber.BusNumber | -hbaid HBA_ID } [-s]
    dlnkmgr online [-path] { -cha -pathid AutoPATH_ID | -chaid CHA_ID } [-s]
    dlnkmgr online [-path] [-pathid AutoPATH_ID] [-s]
  Valid value
    AutoPATH_ID     { 000000 - 999999 }(Decimal)
    HBA_ID          { 00000 - 99999 }(Decimal)
    CHA_ID          { 00000 - 99999 }(Decimal)
offline:
  Format
    dlnkmgr offline [-path]
                     { -hba HBAPortNumber.BusNumber | -hbaid HBA_ID } [-s]
    dlnkmgr offline [-path] { -cha -pathid AutoPATH_ID | -chaid CHA_ID } [-s]
    dlnkmgr offline [-path] -pathid AutoPATH_ID [-s]
  Valid value
    AutoPATH_ID     { 000000 - 999999 }(Decimal)
    HBA_ID          { 00000 - 99999 }(Decimal)
    CHA_ID          { 00000 - 99999 }(Decimal)
help:
  Format
    dlnkmgr help { clear | offline | online | set | view | monitor }
KAPL01001-I The HDLM command completed normally. Operation name = help, completion time =
yyyy/mm/dd hh:mm:ss
#
```

Example 3

The following example shows how to display the operation names that can be specified for the `help` operation

```
# /opt/DynamicLinkManager/bin/dlnkmgr help help
help:
  Format
    dlnkmgr help { clear | offline | online | set | view | monitor }
KAPL01001-I The HDLM command completed normally. Operation name = help, completion time =
yyyy/mm/dd hh:mm:ss
#
```

6.4 offline (Places a Path or Paths Offline)

The `dlnkmgr` command's `offline` operation places online paths offline. Specify the paths you want to place offline by specifying an HBA port or channel adapter port to which the target paths are connected, or by specifying a single path.

The last path accessing each LU cannot be placed offline.

Placing too many paths offline may prevent path switching if an error occurs. Before placing a path offline, use the `view` operation to check how many online paths remain. For details on the `view` operation, see 6.7.

6.4.1 Format

6.4.1.1 To place the path offline

```
/opt/DynamicLinkManager/bin/dlnkmgr offline
   [-path]
   {-hba host-port-number.bus-number
   |-hbaid HBA-port-ID
   |-cha -pathid AutoPATH_ID
   |-chaid CHA-port-ID
   |-pathid AutoPATH_ID}
   [-s]
```

6.4.1.2 To display the format of the offline operation

```
/opt/DynamicLinkManager/bin/dlnkmgr offline -help
```

6.4.2 Parameters

6.4.2.1 To place the path offline

-path

Indicates that the target of the operation is a path or paths managed by HDLM.

This parameter is optional because the offline operation is only effective on paths.

Make sure, however, that you specify the target path in a subsequent parameter: -hba, -hbaid, -cha, -chaid or -pathid.

-hba *host-port-number.bus-number*

Use this parameter to place offline, at one time, all paths that pass through a specific HBA port. The command will place offline all paths connected to the HBA port that has the specified host port number and bus number.

Specify the host port number and bus number of the target HBA port: the numbers are found in the PathName field displayed using the view operation. Enter a period between these two parameter values. For details on the view operation, see 6.7.

Example

The following example shows how to place offline all paths connected to the HBA port whose host port number is 0001 and bus number is 0000.

```
# /opt/DynamicLinkManager/bin/dlnkmgr offline -hba 0001.0000
KAPL01055-I All the paths which pass the specified HBA will be changed to the Offline(C)
status. Is this OK? [y/n]:y
KAPL01056-I If you are sure that there would be no problem when all the paths which pass
the specified HBA are placed in the Offline(C) status, enter y.  Otherwise, enter n.
[y/n]:y
KAPL01061-I 3 path(s) were successfully placed offline(C). 0 path(s) could not be placed
offline(C). Operation name = offline
#
```

-hbaid *HBA-port-ID*

Use this parameter to place offline, at one time, all paths that pass through a specific HBA port. The command will place offline all paths connected to the HBA port that has the specified *HBA-port-ID*. To display the HBA port IDs, execute the following view operation:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -hba
```

For details about how to execute the view operation to display the HBA port IDs, see

For details, see 6.7.2.4 in 6.7.2.

Example

The following example shows how to place offline all paths connected to the HBA port whose *HBA-port-ID* is 00001.

```
# /opt/DynamicLinkManager/bin/dlnkmgr offline -hbaid 00001
KAPL01102-I All the paths which pass the specified HBA port will be changed to the
Offline(C) status. Is this OK? [y/n]:y
KAPL01103-I If you are sure that there would be no problem when all the paths which pass
the specified HBA port are placed in the Offline(C) status, enter y.  Otherwise, enter
n. [y/n]:y
KAPL01061-I 15 path(s) were successfully placed offline(C). 0 path(s) could not be placed
offline(C). Operation name = offline
```

```
#
```

-cha -pathid *AutoPATH_ID*

Use this parameter to place offline, at one time, all paths that pass through a specific channel adapter port. The command will place offline all paths that pass through the channel adapter port to which the path with the specified *AutoPATH_ID* is connected.

Specify the current AutoPATH_ID of the target path, which is displayed by using the `view` operation. For details on the `view` operation, see 6.7. Leading zeros can be omitted (`000001` and `1` indicate the same *AutoPATH_ID*) . Therefore, when the target AutoPATH_ID is `000000`, enter `000000` or `0` for the parameter value.

AutoPATH_IDs are re-assigned every time the host is restarted. When you want to add a new LU without restarting the host, AutoPATH_IDs are re-assigned to each path of the LU when you execute the `dlmcfgmgr` utility for managing the HDLM configuration. Always make sure that you execute the `view` operation to find the current AutoPATH_ID of the target path, before executing the `offline` operation.

Example

The following example shows how to place offline all paths connected to the channel adapter port `0A`. In this example, a path whose AutoPATH_ID is `000001` is connected to the target channel adapter port.

```
# /opt/DynamicLinkManager/bin/dlnkmgr offline -cha -pathid 000001
KAPL01055-I All the paths which pass the specified CHA port will be changed to the
Offline(C) status. Is this OK? [y/n]:y
KAPL01056-I If you are sure that there would be no problem when all the paths which pass
the specified CHA are placed in the Offline(C) status, enter y.  Otherwise, enter n.
[y/n]:y
KAPL01061-I 2 path(s) were successfully placed offline(C). 0 path(s) could not be placed
offline(C). Operation name = offline
#
```

-chaid *CHA-port-ID*

Use this parameter to place offline, at one time, all paths that pass through a specific CHA port. The command will place offline all paths connected to the CHA port that has the specified CH*A-port-ID*. To display the CHA port IDs, execute the following `view` operation:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -cha
```

For details about how to execute the `view` operation to display the CHA port IDs, see

6.7.2.5 in 6.7.2.

Example

The following example shows how to place offline all paths connected to the CHA port whose CH*A-port-ID* is `00001`.

```
# /opt/DynamicLinkManager/bin/dlnkmgr offline -chaid 00001
KAPL01102-I All the paths which pass the specified CHA port will be changed to the
Offline(C) status. Is this OK? [y/n]:y
KAPL01103-I If you are sure that there would be no problem when all the paths which pass
the specified HBA port are placed in the Offline(C) status, enter y.  Otherwise, enter
n. [y/n]:y
KAPL01061-I 15 path(s) were successfully placed offline(C). 0 path(s) could not be placed
offline(C). Operation name = offline
#
```

-pathid *AutoPATH_ID*

Use this parameter to place a single path offline.

Specify the current AutoPATH_ID of the target path, which is displayed by using the `view` operation. For details on the `view` operation, see 6.7. Leading zeros can be omitted (`000001` and `1` indicate the same *AutoPATH_ID*) . Therefore, when the target AutoPATH_ID is `000000`, enter `000000` or `0` for the parameter value.

AutoPATH_IDs are re-assigned every time the host is restarted. When you want to add a new LU without restarting the host, AutoPATH_IDs are re-assigned to each path of the LU when you execute the `dlmcfgmgr` utility. Always make sure that you execute the `view` operation to find the current AutoPATH_ID of the target path, before executing the `offline` operation.

`-s`

Executes the command without displaying the message asking for confirmation of command execution from the user. Specify this parameter if you want to skip the response to the confirmation message: for example, when you want to execute the command in a shell script or batch file.

Example

The following example shows how to place a path, whose AutoPATH_ID is `000001`, offline without asking for confirmation of command execution from the user.

```
# /opt/DynamicLinkManager/bin/dlnkmgr offline -pathid 1 -s
KAPL01061-I 1 path(s) were successfully placed offline. 0 path(s) could not be placed
offline. Operation name = offline
#
```

6.4.2.2 To display the format of the offline operation

`-help`

Displays the format of the `offline` operation.

Example

The following example shows how to display the format of the `offline` operation.

```
# /opt/DynamicLinkManager/bin/dlnkmgr offline -help
offline:
  Format
    dlnkmgr offline [-path]
                     { -hba HBAPortNumber.BusNumber | -hbaid HBA_ID } [-s]
    dlnkmgr offline [-path] { -cha -pathid AutoPATH_ID | -chaid CHA_ID } [-s]
    dlnkmgr offline [-path] -pathid AutoPATH_ID [-s]
  Valid value
    AutoPATH_ID     { 000000 - 999999 }(Decimal)
    HBA_ID          { 00000 - 99999 }(Decimal)
    CHA_ID          { 00000 - 99999 }(Decimal)
KAPL01001-I The HDLM command completed normally. Operation name = offline, completion
time = yyyy/mm/dd hh:mm:ss
#
```

Reference

Executing the `view` operation together with standard UNIX commands enables you to filter the path information listed for a specific HBA port or channel adapter port. For details on the `view` operation, see 6.7.

We recommend that you execute the following command and verify the information on the target paths before you execute the `offline` operation to place offline all paths connected to a specific HBA port or channel adapter port.

Example 1

The following example shows how to filter and display the information on all paths that pass through the HBA port whose host port number is `0004` and bus number is `0000`.

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -path | grep 0004.0000
```

The above commands will display information on all the paths that pass through the specified HBA port.

Example 2

The following example shows how to filter and display the information on all paths that pass through the channel adapter port `0A` of the Thunder 9500V Series.

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -path -stname | grep 9500V | grep 0A
```

The above commands will display information pertaining to only those paths that pass through the specified channel adapter port.

6.5 online (Places a Path or Paths Online)

The `dlnkmgr` command's `online` operation places offline paths online. You can specify the paths you want to place online by specifying an HBA port or channel adapter port to which the target paths are connected, or by specifying a single path.

6.5.1 Format

6.5.1.1 To place the path online

`/opt/DynamicLinkManager/bin/dlnkmgr online`
`[-path]`
`[-hba` ***host-port-number.bus-number***
`|-hbaid` ***HBA-port-ID***
`|-cha -pathid` ***AutoPATH_ID***
`|-chaid` ***CHA-port-ID***
`|-pathid` ***AutoPATH_ID***`]`
`[-s]`

6.5.1.2 To display the format of the online operation

```
/opt/DynamicLinkManager/bin/dlnkmgr online -help
```

6.5.2 Parameters

6.5.2.1 To place the path online

`-path`

Indicates that the target of the operation is a path managed by HDLM.

This parameter is optional because the `online` operation is only effective on paths.

You can specify the target path in a subsequent parameter: `-hba`, `-hbaid`, `-cha`, `-chaid`, or `-pathid`. If you do not specify any of these parameters, the command places all the offline paths online. If there is a path that cannot be placed online, a message asking whether you would like to continue processing will be displayed. To ignore the offline path that cannot be placed online and to continue processing, enter `y`. To stop the processing, enter `n`.

`-hba` *host-port-number.bus-number*

Use this parameter to place online, at one time, all paths that pass through a specific HBA port. The command will place online all paths connected to the HBA port that has the specified host port number and bus number.

Specify the host port number and bus number of the target HBA port: the numbers are found in the `PathName` field displayed using the `view` operation. Enter a period between these two parameter values. For details on the `view` operation, see 6.7.

Example

The following example shows how to place online all paths connected to the HBA port whose host port number is `0001` and bus number is `0000`.

When the confirmation message is displayed, the user enters `y` to continue, or `n` to cancel the operation.

```
# /opt/DynamicLinkManager/bin/dlnkmgr online -hba 0001.0000
KAPL01057-I All the paths which pass the specified HBA will be changed to the Online
status. Is this OK? [y/n]:y
KAPL01061-I 3 path(s) were successfully placed online. 0 path(s) could not be placed
online. Operation name = online
#
```

`-hbaid` *HBA-port-ID*

Use this parameter to place online, at one time, all paths that pass through a specific HBA port. The command will place online all paths connected to the HBA port that has the specified *HBA-port-ID*. To display the HBA port IDs, execute the following `view` operation:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -hba
```

For details about how to execute the `view` operation to display the HBA port IDs, see 6.7.2.4 in 6.7.2.

Example

The following example shows how to place online all paths connected to the HBA port whose *HBA-port-ID* is `00001`.

```
# /opt/DynamicLinkManager/bin/dlnkmgr online -hbaid 00001
KAPL01104-I All the paths which pass the specified HBA port will be changed to the
Online(C) status. Is this OK? [y/n]:y
KAPL01061-I 15 path(s) were successfully placed online. 0 path(s) could not be placed
online. Operation name = online
#
```

`-cha -pathid` *AutoPATH_ID*

Use this parameter to place online, at one time, all paths that pass through a specific channel adapter port. The command will place online all paths that pass through the channel adapter port to which the path with the specified *AutoPATH_ID* is connected.

Specify the current AutoPATH_ID of the target path, which is displayed by using the `view` operation. For details on the `view` operation, see 6.7. Leading zeros can be omitted (`000001` and `1` indicate the same *AutoPATH_ID*) . Therefore, when the target AutoPATH_ID is `000000`, enter `000000` or `0` for the parameter value.

AutoPATH_IDs are re-assigned every time the host is restarted. When you want to add a new LU without restarting the host, AutoPATH_IDs are re-assigned to each path of the LU when you execute the `dlmcfgmgr` utility for managing the HDLM configuration. Always make sure that you execute the `view` operation to find the current AutoPATH_ID of the target path, before executing the `online` operation.

Example

The following example shows how to place online the paths connected to the channel adapter port `0A`. In this example, a path whose AutoPATH_ID is `000002` is connected to the target channel adapter port.

```
# /opt/DynamicLinkManager/bin/dlnkmgr online -cha -pathid 000002
KAPL01057-I All the paths which pass the specified CHA port will be changed to the Online
status. Is this OK? [y/n]:y
KAPL01061-I 2 path(s) were successfully placed online. 0 path(s) could not be placed
online. Operation name = online
#
```

`-chaid` *CHA-port-ID*

Use this parameter to place online, at one time, all paths that pass through a specific CHA port. The command will place online all paths connected to the CHA port that has the specified CH*A-port-ID*. To display the CHA port IDs, execute the following `view` operation:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -cha
```

For details, see 6.7.2.5 in 6.7.2.

Example

The following example shows how to place online all paths connected to the CHA port whose CH*A-port-ID* is `00001`.

```
# /opt/DynamicLinkManager/bin/dlnkmgr offline -chaid 00001
KAPL01104-I All the paths which pass the specified CHA port will be changed to the
Online(C) status. Is this OK? [y/n]:y
KAPL01061-I 15 path(s) were successfully placed online. 0 path(s) could not be placed
online. Operation name = online
#
```

`-pathid` *AutoPATH_ID*

Use this parameter to place a single path online.

Specify the current AutoPATH_ID of the target path, which is displayed by using the `view` operation. For details on the `view` operation, see 6.7. Leading zeros can be omitted (`000001` and `1` indicate the same *AutoPATH_ID*) . Therefore, when the target AutoPATH_ID is `000000`, enter `000000` or `0` for the parameter value.

AutoPATH_IDs are re-assigned every time the host is restarted. When you want to add a new LU without restarting the host, AutoPATH_IDs are re-assigned to each path of the LU when you execute the `dlmcfgmgr` utility. Always make sure that you execute the `view` operation to find the current AutoPATH_ID of the target path, before executing the `online` operation.

`-s`

Executes the command without displaying the message asking for confirmation of command execution from the user. Specify this parameter if you want to skip the response to the confirmation message: for example, when you want to execute the command in a shell script or batch file.

Example

The following example shows how to place a path, whose AutoPATH_ID is `000002`, online without asking for confirmation of command execution from the user.

```
# /opt/DynamicLinkManager/bin/dlnkmgr online -pathid 2 -s
KAPL01061-I 1 path(s) were successfully placed online. 0 path(s) could not be placed
online. Operation name = online
#
```

6.5.2.2 To display the format of the online operation

`-help`

Displays the format of the `online` operation.

Example

The following example shows how to display the format of the `online` operation.

```
# /opt/DynamicLinkManager/bin/dlnkmgr online -help
online:
  Format
    dlnkmgr online [-path]
                   { -hba HBAPortNumber.BusNumber | -hbaid HBA_ID } [-s]
    dlnkmgr online [-path] { -cha -pathid AutoPATH_ID | -chaid CHA_ID } [-s]
    dlnkmgr online [-path] [-pathid AutoPATH_ID] [-s]
  Valid value
    AutoPATH_ID     { 000000 - 999999 }(Decimal)
    HBA_ID          { 00000 - 99999 }(Decimal)
    CHA_ID          { 00000 - 99999 }(Decimal)
KAPL01001-I The HDLM command completed normally. Operation name = online, completion
time = yyyy/mm/dd hh:mm:ss
#
```

Reference

Executing the `view` operation together with standard UNIX commands enables you to filter the path information listed for a specific HBA port or channel adapter port. For details on the `view` operation, see 6.7.

We recommend that you execute the following command and verify the information on the target paths before you execute the `online` operation to place online all paths connected to a specific HBA port or channel adapter port.

Example 1

The following example shows how to filter and display the information on all paths that pass through the HBA port whose host port number is `0004` and bus number is `0000`.

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -path | grep 0004.0000
```

The above commands will display information on all the paths that pass through the specified HBA port.

Example 2

The following example shows how to filter and display the information on all paths that pass through the channel adapter port `0A` of the Thunder 9500V Series.

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -path -stname | grep 9500V | grep 0A
```

The above commands will display information pertaining to only those paths that pass through the specified channel adapter port.

6.6 set (Sets the Operating Environment)

The `dlnkmgr` command's `set` operation sets the HDLM operating environment.

6.6.1 Format

6.6.1.1 To set up the HDLM operating environment

```
/opt/DynamicLinkManager/bin/dlnkmgr set
   {-lb {on [-lbtype {rr|exrr}]|off}
   |-ellv error-log-collection-level
   |-elfs error-log-file-size
   |-elfn number-of-error-log-files
   |-systflv trace-level
   |-systfs trace-file-size
   |-systfn number-of-trace-files
   |-pchk {on [-intvl checking-interval]|off}
   |-afb {on [-intvl checking-interval]|off}
   |-iem { on [-intvl error-monitoring-interval]
   [-iemnum number-of-times-error-is-to-occur] | off }
```

```
|-lic}
[-s]
```

6.6.1.2 To display the format of the set operation

```
/opt/DynamicLinkManager/bin/dlnkmgr set -help
```

6.6.2 Parameters

6.6.2.1 To set up the HDLM operating environment

Table 6.2 shows the default and recommended values for each setting. If you change the set value by using the `set` operation, the new value takes effect immediately.

Table 6.2 Default values and recommended values

Item name	Default value	Recommended value
Load balancing	`on` Algorithm is round robin	`on` The recommended algorithm depends on the operating environment.
Error log collection level	`3`: Collects error information at the Information level or higher.	`3`: Collects error information at the Information level or higher.
Error log file size	`9900` (kilobytes)	`9900` (kilobytes)
Number of error log files	`2`	`2`
Trace level	`0`: Do not output any trace.	`0`: Do not output any trace.
Trace file size	`1000` (kilobytes)	`1000` (kilobytes)
Number of trace files	`4`	`4`
Path health checking	`on` (Interval is 30 minutes.)	`on` (Interval is 30 minutes.) The recommended checking interval depends on the operating environment.
Automatic failback	`off`	`off`
Intermittent Error Monitor	`off`	`off`

`-lb {on [-lbtype {rr|exrr}]|off}`

Enables or disables load balancing.

`on`: Enabled

`off`: Disabled

`-lbtype {rr|exrr}`

Specify the algorithm to be used for load balancing.

`rr`: Round robin

exrr: Extended round robin

The type of algorithm specified by the -lbtype parameter remains stored in the system, even when you disable load balancing function by specifying -lb off. Therefore, when you re-enable load balancing without specifying an algorithm, load balancing will be executed according to the setting stored in the system.

-ellv *error-log-collection-level*

Specify the level of error information you want to collect for an error log.

The following shows the log files in which an error log collection level can be set:

HDLM manager logs:

dlmmgr*n*.log (*n* indicates a file number from 1 to 16)

Table 6.3 shows the values of the error log collection level.

Table 6.3 Values of the Error Log Collection Level

Value	Description
0	Collects no error log.
1	Collects error information at the Error level or higher.
2	Collects error information at the Warning level or higher.
3	Collects error information at the Information level or higher.
4	Collects error information at the Information level or higher (including maintenance information).

The amount of log information to be output increases as the value of the setting increases.

As the amount of log information to be output increases, the amount of time before existing information is overwritten becomes shorter.

Example

```
# /opt/DynamicLinkManager/bin/dlnkmgr set -ellv 1
KAPL01049-I Would you like to execute the operation? Operation name = set [y/n]: y
KAPL01001-I The HDLM command completed normally. Operation name = set, completion time
= yyyy/mm/dd hh:mm:ss
#
```

-elfs *error-log-file-size*

Specify a value from 100 to 2000000 (in kilobytes) for the size of the error log files. The target log files for which a file size can be set are HDLM manager logs (dlmmgr*n*.log (*n* indicates a file number from 1 to 16)) The specified value is applied for HDLM manager logs.

By specifying both the log file size and the number of log files, you can collect up to 32,000,000 kilobytes (approximately 30GB) of error logs in total.

If the sizes of all log files in a log file group reach the specified value, the oldest file of the corresponding log file group is overwritten by the new log information

-elfn *number-of error-log-files*

Specify the number of error log files. Specify a value from 2 to 16.

By specifying both the log file size and the number of log files, you can collect up to 32,000,000 kilobytes (approximately 30GB) of error logs in total.

The target log files for which the number of files can be set are HDLM manager logs (`dlmmgr`*n*`.log` (*n* indicates a file number from `1` to `16`)) only.

`-systflv` *trace-level*

Specify the trace output level.

The trace files for which trace levels can be set are `hdlmtr`*n*`.log` (*n* indicates a file number from `1` to `64`).

Table 6.4 shows the values of the trace level. When an error occurs, we recommend that you set the value of the trace level to `1` or higher and re-collect log information.

Table 6.4 Trace Level Values

Value	Description
`0`	Does not output any trace.
`1`	Only outputs error information.
`2`	Outputs a summary of program operation.
`3`	Outputs details of program operation.
`4`	Outputs all information.

The amount of log information to be output increases as the value of the setting increases.

As the amount of log information to be output increases, the amount of time before existing information is overwritten becomes shorter.

`-systfs` *trace-file-size*

Specify the size of the trace file in kilobytes. Specify a value from `100` to `16000`.

By specifying both the trace file size and the number of trace files, you can collect up to 1024000 kilobytes of trace files in total.

If a value smaller than the set value is specified, the execution confirmation message KAPL01097-W is displayed and the trace file is deleted for the time being.

The trace files for which a file size can be set are `hdlmtr`*n*`.log` (*n* indicates a file number from `1` to `64`). The trace files are fixed in length. Thus, even if the amount of written trace information is less than the setting file size, the size of each output trace file is always fixed. When trace data is written to all trace files, new trace data overwrites old trace data (the oldest file is overwritten first).

`-systfn` *number-of-trace-files*

Specify the number of trace files. Specify a value from `2` to `64`.

By specifying both the trace file size and the number of trace files, you can collect up to 1024000 kilobytes of trace files in total.

If a value smaller than the set value is specified, the execution confirmation message KAPL01097-W is displayed and the trace file is deleted for the time being.

The trace files for which the number of files can be set are `hdlmtr`*`n`*`.log` (*n* indicates a file number from `1` to `64`).

`-pchk {on [-intvl` ***checking-interval***`]|off}`

Enables or disables path health checking.

`on`: Enabled

`off`: Disabled

For a standby host, or a host connected to the Freedom Storage 5700E, Freedom Storage 5800, Thunder 9200, or Thunder 9500V Series, or TagmaStore AMS/WMS series storage subsystem, we recommend that you activate path health checking to enable detection of errors in paths where I/Os do not occur.

When you specify `on`, specify the checking interval of path health checking by specifying the parameter immediately following `on`. If you do not specify a checking interval, path health checking is executed in the following interval:

- When the checking interval has not been specified before:

  Every 30 minutes (default setting)

- When the checking interval has been specified before:

  The interval specified in the last time

The explanation for the following sub-parameter describes how to specify the checking interval.

`-intvl` ***checking-interval***

Specify the checking interval between path health checks (in minutes). Specify a value from 1 to 1440 minutes depending on the user environment.

When you change the checking interval, the new setting takes effect immediately. When the checking interval is shortened and the checking interval after the change (from the end of the previous path health check) has already elapsed, the path health check will start immediately.

The path health check interval setting remains stored in the system even if you disable the function by changing the setting of path health checking to `off`. Therefore, when you re-enable path health checking and do not specify the interval, path health checking is executed at the interval stored in the system.

`-afb {on [-intvl` ***checking-interval***`]|off}`

Enables or disables automatic failback.

`on`: Enabled

`off`: Disabled

Enabling automatic failback may automatically place online those paths that were intentionally placed offline (for example, paths placed offline for maintenance work).

When intermittent errors occur in paths or storage subsystems, path status alternates between the online and offline status frequently, so I/O performance might deteriorate.

Automatic failback is executed in the following paths:

- Path where an error occurred and for which the KAPL08022-E message was displayed

– Path where an error occurred at the startup of the HDLM manager.

To prevent an intermittent error from deteriorating I/O performance, we recommend that you also enable intermittent error monitoring when enabling automatic failback. Intermittent error monitoring is specifiable only when automatic failback is enabled.

For details on the relationship between automatic failback and intermittent error monitoring, see Table 6.5.

When you specify `on`, specify the checking interval of path status by specifying the parameter immediately following `on`. If you do not specify a checking interval, path status is checked in the following interval:

– When the checking interval has not been specified before:

  Every minute (default setting)

– When the checking interval has been specified before:

  The interval specified in the last time

The explanation for the following sub-parameter describes how to specify the interval between path status checks.

`-intvl` *checking-interval*

Specify the interval between path status checks (in minutes). Specify a value from `1` to `1440` minutes. The default is `1`. Specify the interval appropriate for your operation environment.

If intermittent error monitoring is `on` and the number of times that the error is to occur is set to a value of 2 or more, the following condition must be satisfied:

error-monitoring-interval >= checking-interval-for-automatic-failback x *number-of-times-error-is-to-occur-during-intermittent-error-monitoring*

If this condition is not satisfied, the KAPL01080-W message is output and an error occurs. In such a case, change any of the following settings: the checking interval for automatic failback, intermittent error monitoring interval, or the number of times that the error is to occur.

When you set the number of times that the error is to occur to `1`, the above condition does not need to be satisfied.

When you changed the error monitor interval during monitoring an intermittent error, the new setting takes effect immediately. When the checking interval is shortened and the checking interval after the change (from the end of the previous path status check) has already elapsed, the path status check will start immediately.

This setting remains stored in the system, even if you disable the function by changing the setting of automatic failback to `off`. Therefore, when you re-enable automatic failback and do not change the interval, path status checks are executed at the interval stored in the system.

`-iem { on [-intvl` *error-monitoring-interval*`] [-iemnum` *number-of-times-error-is-to-occur*`] | off }`

Enables or disables intermittent error monitoring.

`on`: Enabled

`off`: Disabled

Intermittent error monitoring can be set when automatic failback is set to `on`.

When you use automatic failback, we recommend that you set intermittent error monitoring to `on` to prevent an intermittent error from reducing I/O performance.

If `on` is specified, in the subsequent parameters, specify the intermittent error monitoring interval and the number of times that the error is to occur. The system assumes that an intermittent error is occurring if the specified number of times that the error is to occur is reached during the monitoring interval (from the time that the monitoring interval starts, until the specified interval ends). A path that is assumed to have an intermittent error is excluded from automatic failback. Intermittent error monitoring is performed on individual paths. Intermittent error monitoring starts when a path is recovered from the error by using automatic failback.

If you omit the intermittent error monitoring interval or the number of times that the error is to occur, each setting is specified as follows:

- When the intermittent error monitoring interval or the number of times that the error is to occur has not been specified before:

  The intermittent error monitoring interval is set to 30 minutes, and the number of times that the error is to occur is set to 3.

- When the intermittent error monitoring interval or the number of times that the error is to occur has been specified before:

  The values specified the last time are set.

When a value of 2 or more is specified in number of times, the following condition must be satisfied:

```
error-monitoring-interval >= checking-interval-for-automatic-failback x
number-of-times-error-is-to-occur-during-intermittent-error-monitoring
```

If this condition is not satisfied, the KAPL01080-W message is output and an error occurs. In such a case, change any of the following settings: the checking interval for automatic failback, intermittent error monitoring interval, or the number of times that the error is to occur.

When you set the number of times that the error is to occur to `1`, the above condition does not need to be satisfied.

The following shows the sub-parameters to specify the error monitoring interval and the number of times that the error is to occur (in order for the system to determine that an intermittent error is occurring):

`-intvl` *error-monitoring-interval*

Specify the monitoring interval for an intermittent error (in minutes). Use a value from `1` to `1440`. The default is `30`.

During intermittent error monitoring, if changes are made in the settings of the intermittent error monitoring interval, the error count and the elapsed time measured since monitoring starts are set to 0, and monitoring starts by using the new settings.

Outside the duration of intermittent error monitoring, if changes are made in the settings of the intermittent error monitoring interval, the new settings take effect after the next time automatic failback succeeds. Outside the duration of intermittent error monitoring, the number of errors that determine an intermittent error is not counted and this value is not changed.

The monitoring interval specified in this parameter is stored even though specifying `-iem off` disables intermittent error monitoring. Therefore, when you re-enable intermittent error monitoring and the monitoring interval is not specified, error monitoring will be executed for the stored monitoring interval.

`-iemnum` *number-of-times-error-is-to-occur*

Specify the number of times the error is to occur. Use a value from `1` to `99`. The default is `3`.

During intermittent error monitoring, if you change the number of times that the error is to occur, the error count and the elapsed time measured since monitoring starts are reset to 0, and monitoring starts by using the new settings.

Outside the duration of intermittent error monitoring, if you change the number of times that the error is to occur, the new settings take effect after the next time automatic failback succeeds. Outside the duration of intermittent error monitoring, the number of errors that determine an intermittent error is not counted and this value is not changed.

The number of times that the error is to occur specified in this parameter is stored even though specifying `-iem off` disables intermittent error monitoring. Therefore, when you re-enable intermittent error monitoring and the number of times is not specified, the error monitoring will be executed using the stored number of times.

When the `set -iem on` operation is executed during error monitoring, even though you do not change the conditions for intermittent error, the number of errors and the time elapsed since the error monitoring starts are reset to 0. Then intermittent error monitoring resumes with the changed settings.

If you set automatic failback to `off` while intermittent error monitoring is `on`, intermittent error monitoring becomes disabled. Note, however, that if you use the `view -sys` operation to display the HDLM function configuration, `Intermittent Error Monitor` shows `on`. When automatic failback is returned to `on`, intermittent error monitoring becomes enabled.

The executable operations for automatic failback and intermittent error monitoring depend on the setting status for those functions. Table 6.5 shows the relationship between the setting status for automatic failback and intermittent error monitoring and the executable operations for those functions.

Table 6.5 Relationship Between the Setting Status for Automatic Failback and Intermittent Error Monitoring and the Executable Operations

Setting		Executable operation	Result of operation
AFB	IEM		
on	on	Set AFB to on.	The operations of AFB and IEM do not change.
		Change the AFB setting.	AFB operates using new settings.#1
		Set AFB to off.	▪ AFB and IEM are disabled. ▪ The error count, elapsed monitoring time, and information about paths not subject to automatic failback are reset.
		Set IEM to on.	▪ When a path is being monitored (in the period of conditional intermittent error monitoring), the value of the error count and the elapsed monitoring time are reset to 0, and then intermittent error monitoring restarts. ▪ When a path is not being monitored (outside the period of conditional intermittent error monitoring), nothing changes.
		Change the IEM settings.	▪ When a path is being monitored (in the period of conditional intermittent error monitoring), the value of the error count and the elapsed monitoring time are reset to 0, and then intermittent error monitoring restarts according to the conditions for intermittent error after change.#1 ▪ When a path is not being monitored (outside the period of conditional intermittent error monitoring), the IEM settings will take effect when the path is recovered from the error status by performing automatic failback.
		Set IEM to off.	▪ IEM is disabled. ▪ The error count, elapsed monitoring time, and information about paths not subject to automatic failback are reset.
	off	Set AFB to on.	The operations of AFB and IEM do not change.
		Change the AFB setting.	AFB operates using new settings.
		Set AFB to off.	AFB is disabled.
		Set IEM to on.	IEM is enabled.#1
off	on#2	Set AFB to on.	AFB and IEM are enabled.#1
		Set AFB to off.	The operations of AFB and IEM do not change.
	off	Set AFB to on.	AFB is enabled.
		Set AFB to off.	The operations of AFB and IEM do not change.

Legend:

AFB: Automatic failback

IEM: Intermittent error monitoring

#1

When this condition is not satisfied, the KAPL01080-W message is output and an error occurs. The status of the intermittent error monitoring does not change.

#2

Because automatic failback is `off`, intermittent error monitoring is disabled.

Example

The following example shows how to monitor for intermittent errors.

```
# /opt/DynamicLinkManager/bin/dlnkmgr set -iem on -intvl 20 -iemnum 2
KAPL01049-I Would you like to execute the operation? Operation name = set [y/n]: y
KAPL01001-I The HDLM command completed normally. Operation name = set, completion time
= yyyy/mm/dd hh:mm:ss
#
```

`-lic`

Specify this option when a license is updated. The HDLM license is provided by a license key or a license key file. A license key file is a file that stores an HDLM license key.

If you use a license key file:

Store the license key file named `hdlm_license` directly under `/var/tmp`, and then execute the `set -lic` operation. A message confirming that the license key has been registered is displayed, depending on the license key type described in the license key file. When a temporary license key or emergency license key has been registered, the expiration period is displayed (KAPL01071-I, KAPL01072-I).

If you do not use a license key file:

When the `set -lic` operation is executed, a message (KAPL01068-I) asking the user to enter a license key appears. Enter the license key. A message confirming that the license key has been registered is displayed, depending on the license key type described in the license key file. When a temporary license key or emergency license key has been registered, the expiration period is displayed (KAPL01071-I, KAPL01072-I).

Table 6.6 lists the license key types.

Table 6.6 License Key Types

Type	Description
Permanent license key	Permanent licenses are valid for using HDLM in an ongoing basis.
Temporary license key#	The temporary license key is used temporarily, for example, when a user performs product evaluations. Temporary licenses are valid for only 120 days after installation. You cannot reuse a temporary license key.
Emergency license key	The emergency license key is used temporarily, for example, when issuing a permanent license key that is delayed for some reasons. Emergency licenses are valid for 30 days after they are entered. You cannot reuse an emergency license key.

#

A temporary license key cannot be installed by using the `dlnkmgr` command's `set` operation.

Example 1

The following example shows how to update the license key (when the license key file exists).

```
# /opt/DynamicLinkManager/bin/dlnkmgr set -lic
KAPL01049-I Would you like to execute the operation? Operation name = set [y/n]: y
KAPL01071-I A permanent license was installed.
KAPL01001-I The HDLM command completed normally. Operation name = set, completion time
= yyyy/mm/dd hh:mm:ss
#
```

Example 2

The following example shows how to update the license key (when the license key file does not exist).

```
# /opt/DynamicLinkManager/bin/dlnkmgr set -lic
KAPL01049-I Would you like to execute the operation? Operation name = set [y/n]: y
KAPL01083-I There is no license key file. File name = /var/tmp/hdlm_license
KAPL01068-I Enter a license key:*************
KAPL01071-I A permanent license was installed.
KAPL01001-I The HDLM command completed normally. Operation name = set, completion time
= yyyy/mm/dd hh:mm:ss
#
```

`-s`

Executes the command without displaying the message asking for confirmation of command execution from the user. Specify this parameter if you want to skip the response to the confirmation message. For example, when you want to execute the command in a shell script or batch file.

6.6.2.2 To display the format of the set operation

`-help`

Displays the format of the `set` operation.

Example

6.7 view (Displays Information)

The `dlnkmgr` command's `view` operation displays HDLM program information, path information, LU information, HBA port information, CHA port information,and corresponding information about an HDLM device, sd or ssd device, and LDEV.

6.7.1 Format

6.7.1.1 To display program information

```
/opt/DynamicLinkManager/bin/dlnkmgr view -sys
   [-sfunc|-msrv|-adrv|-pdrv|-lic]
   [-t]
```

6.7.1.2 To display path information

To display path information

```
/opt/DynamicLinkManager/bin/dlnkmgr view -path
   [-hdev host-device-name]
   [-stname]
   [-iem]
   [-srt {pn|lu|cp}]
   [-t]
```

To display path information (by selecting a display item)

```
/opt/DynamicLinkManager/bin/dlnkmgr view -path -item
   [pn] [dn] [lu] [cp] [type] [ic] [ie] [dnu] [hd] [iep]
   [-hdev host-device-name]
   [-stname]
   [-srt {pn|lu|cp}]
   [-t]
```

To display a summary of path information

```
/opt/DynamicLinkManager/bin/dlnkmgr view -path -c
   [-stname]
   [-srt {lu|cp}]
   [-t]
```

6.7.1.3 To display LU information

To display LU information

```
/opt/DynamicLinkManager/bin/dlnkmgr view -lu
```

```
   [-hdev host-device-name | -pathid AutoPATH_ID]
   [-t]
```

To display LU information (by adding items to be displayed)

```
/opt/DynamicLinkManager/bin/dlnkmgr view -lu -item
   [ [slpr] [pn] [cp] [clpr] [type] [ic] [ie] [dnu] [iep]
   [pd] | all ]
   [-hdev host-device-name | -pathid AutoPATH_ID]
   [-t]
```

To display a summary of LU information

```
/opt/DynamicLinkManager/bin/dlnkmgr view -lu -c [-t]
```

To display a summary of LU information (by adding items to be displayed)

```
/opt/DynamicLinkManager/bin/dlnkmgr view -lu -c -item
   [slpr]
   [-t]
```

6.7.1.4 To display HBA port information

```
/opt/DynamicLinkManager/bin/dlnkmgr view -hba [-srt pb] [-c] [-t]
```

6.7.1.5 To display CHA port information

```
/opt/DynamicLinkManager/bin/dlnkmgr view -cha [-srt cp] [-t]
```

6.7.1.6 To display corresponding information about an HDLM device, sd or ssd device, and LDEV

```
/opt/DynamicLinkManager/bin/dlnkmgr view -drv [-t]
```

6.7.1.7 To display the format of the view operation

```
/opt/DynamicLinkManager/bin/dlnkmgr view -help
```

6.7.2 Parameters

This section describes the parameters for the `view` operation, in the following order:

6.7.2.1

6.7.2.2

6.7.2.3

6.7.2.4

6.7.2.5

6.7.2.6

6.7.2.7

6.7.2.1 To display program information

```
-sys [-sfunc|-msrv|-adrv|-pdrv|-lic]
```

Displays the HDLM program information.

Use one of the subsequent parameters (following `-sys`) to specify the program information that you want to display. If you do not specify a subsequent parameter, the command displays all program information.

Table 6.7 describes the specifiable parameters, displayed information, displayed items, and a corresponding description.

`-t`

Does not display the title for each information item.

Table 6.7 Displayed program information

Parameter and program information	Item	Description
`-sfunc` Information about the HDLM function settings	`HDLM Version`	HDLM version number.
	`Service Pack Version`	HDLM service pack version number. This item is blank if a service pack is not installed.
	`Load Balance`	Setting for load balancing ▪ Setting status `on`: Enabled `off`: Disabled ▪ Algorithm When the setting status of load balancing is `on`, the type of algorithm used for load balancing is displayed within the parentheses, ( ), following `on`. `rr`: Round robin `extended rr`: Extended round robin
	`Support Cluster`	Blank#

Parameter and program information	Item	Description
	Elog Level	Error log collection level ▪ 0: Collects no error information. ▪ 1: Collects error information at the Error level or higher. ▪ 2: Collects error information at the Warning level or higher. ▪ 3: Collects error information at the Information level or higher. ▪ 4: Collects error information at the Information level or higher (including maintenance information).
	Elog File Size (KB)	Size of the error log file in kilobytes
	Number Of Elog Files	Number of error log files
	Trace Level	Trace output level ▪ 0: Does not output any trace. ▪ 1: Only outputs error information. ▪ 2: Outputs a summary of program operation. ▪ 3: Outputs details of program operation. ▪ 4: Outputs all information.
	Trace File Size(KB)	Trace file size in kilobytes.
	Number Of Trace Files	Number of trace files.
	Path Health Checking	▪ Setting for path health checking on: Enabled off: Disabled ▪ Checking interval When the setting of the path health checking is on, the checking interval of path health checking is displayed within the parentheses, ( ), after on. The time is in minutes.
	Auto Failback	▪ Setting for automatic failback on: Enabled off: Disabled ▪ Checking interval When the setting of the automatic failback is on, the checking interval of automatic failback is displayed within the parentheses, ( ), following on. The time is in minutes.

Parameter and program information	Item	Description
	Intermittent Error Monitor	▪ Setting for intermittent error monitoring on: Enabled off: Disabled When automatic failback is off, intermittent error monitoring is disabled although Intermittent Error Monitor shows on. When automatic failback becomes on, intermittent error monitoring becomes enabled. ▪ Intermittent error monitoring interval and number of times that the error is to occur When intermittent error monitoring is set to on, the specified intermittent error monitoring interval and number of times that the error is to occur are displayed in parentheses, (), following on. The format is *number-of-times-error-is-to-occur*/*monitoring-interval*. The time is in minutes.
-msrv Information about the HDLM manager	HDLM Manager	Status of the HDLM manager Alive: Normal Dead: Stopped
	Ver	Version number of the HDLM manager
	WakeupTime	Startup time of the HDLM manager
-adrv Information about the HDLM alert driver	HDLM Alert Driver	Status of the HDLM alert driver Alive: Normal Dead: Stopped
	Ver	Version number of the HDLM alert driver
	WakeupTime	Startup time of the HDLM alert driver
	ElogMem Size	Size of error log memory for the HDLM alert driver in kilobytes
-pdrv Information about the HDLM driver	HDLM Driver	Status of the HDLM driver Alive: Normal Dead: Stopped
	Ver	Version number of the HDLM driver
	WakeupTime	Startup time of the HDLM driver
-lic Information about the HDLM license	License Type	License type ▪ Permanent: permanent license ▪ Temporary: temporary license ▪ Emergency: emergency license
	Expiration	License expiration When using a permanent license: - When using a temporary license or emergency license: The license expiration period is displayed in the following format: *yyyy/mm/dd* (*n*days after). When the view -sys -lic operation is executed, (*n*days after) appears if there are *n* days left until the license period expires. ▪ When there are 100 days left until the license period (2006/08/21) expires 2006/08/21(100days after)

#

When you use cluster software, the settings of the cluster support function and the kinds of cluster servers are not displayed. However, the cluster support function will operate normally.

Examples

Example 1

The following example shows how to display information about the HDLM function settings.

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -sys -sfunc
HDLM Version                  : xx-xx
Service Pack Version          :
Load Balance                  : on(rr)
Support Cluster               :
Elog Level                    : 3
Elog File Size(KB)            : 9900
Number Of Elog Files          : 2
Trace Level                   : 0
Trace File Size(KB)           : 1000
Number Of Trace Files         : 4
Path Health Checking          : on(30)
Auto Failback                 : off
Intermittent Error Monitor    : off
KAPL01001-I The HDLM command completed normally. Operation name = view, completion time =
yyyy/mm/dd hh:mm:ss
#
```

Example 2

The following example shows how to display information about the HDLM manager.

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -sys -msrv
HDLM Manager Ver      WakeupTime
Alive        xx-xx    yyyy/mm/dd hh:mm:ss
KAPL01001-I The HDLM command completed normally. Operation name = view, completion time =
yyyy/mm/dd hh:mm:ss
#
```

Example 3

The following example shows how to display information about the HDLM alert driver.

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -sys -adrv
HDLM Alert Driver Ver       WakeupTime          ElogMem Size
Alive             xx-xx     yyyy/mm/dd hh:mm:ss 4096
KAPL01001-I The HDLM command completed normally. Operation name = view, completion time =
yyyy/mm/dd hh:mm:ss
#
```

Example 4

The following example shows how to display information about the HDLM driver.

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -sys -pdrv
HDLM Driver Ver       WakeupTime
Alive       xx-xx     yyyy/mm/dd hh:mm:ss
KAPL01001-I The HDLM command completed normally. Operation name = view, completion time =
yyyy/mm/dd hh:mm:ss
#
```

Example 5

The following example shows how to display information about the HDLM license.

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -sys -lic
License Type Expiration
Permanent    -
KAPL01001-I The HDLM command completed normally. Operation name = view, completion time =
yyyy/mm/dd hh:mm:ss
#
```

6.7.2.2 To display path information

When displaying path information, if the -item parameter or the -c parameter is specified at the same time as the -path parameter, you can select the items to display and display a summary of path information. This section describes each parameter and the displayed path information.

To display path information

-path

When you specify the -path parameter and do not specify either the -c or -item parameter, the command displays information about the paths managed by HDLM without abbreviating or selecting items.

In the subsequent sub-parameters (following -path), you can filter the paths to be listed (-hdev) and sort the list (-srt). When you omit both parameters, the command displays information for all the paths in order of increasing AutoPATH_IDs.

For details on what is displayed in each item, see Table 6.9.

*AutoPATH_ID*s displayed by the -path parameter depend on the sequence in which HDLM detects the paths when a host is started. Because of this, make sure that you use the path name (PathName) to identify a path.

The following shows the subsequent sub-parameters (following - path):

-hdev *host-device-name*

Filters the information only for the paths accessing the specified host device.

Specify the name of the logical device file (minus the slice number, in c*U*t*X*d*Y* format) for the HDLM device to indicate the desired host device.

The *host-device-name* string is case-sensitive.

-stname

Use this parameter to display the model ID of the storage subsystem in the product ID element of the DskName field. When this parameter is omitted, the command displays the product ID or emulation type of the storage subsystem instead.

For more information pertaining to the storage subsystem and corresponding product ID, see Table 6.11.

-iem

Use this parameter to add IEP to path information and display information about intermittent errors.

-srt {pn|lu|cp}

Use this parameter to sort the list of path information in ascending order, according to the specified sorting keys.

The sorting keys are as follows: the first sorting key is the name of the storage subsystem (DskName), the second sorting key is the value specified by the -srt parameter, and the third sorting key is AutoPATH_ID.

The available parameter values to specify the second sorting key are:

- pn: Path name
- lu: LU number of the storage subsystem
- cp: Port number of the channel adapter

When the -srt parameter is omitted, the path information is listed in order of ascending AutoPATH_IDs.

-t

Does not display the title for each information item.

Example

The following example shows how to display information about the paths that access a host device.

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -path  -hdev c6t5d0
Paths:000002 OnlinePaths:000002
PathStatus   IO-Count    IO-Errors
Online       15          0

PathID PathName                          DskName
iLU              ChaPort Status     Type IO-Count IO-Errors DNum
HDevName
000000 0000.0000.0000000000000005.0000 HITACHI .OPEN-3          .15001
0670             1H      Online     Own        12         0    -
c6t5d0
000001 0001.0000.0000000000000006.0000 HITACHI .OPEN-3          .15001
0670             2H      Online     Own         3         0    -
c6t5d0
KAPL01001-I The HDLM command completed normally. Operation name = view,
completion time = yyyy/mm/dd hh:mm:ss
#
```

To display path information (by selecting a display item)

-path -item

When you specify the -path parameter together with the -item parameter, the command only displays the items specified by the value of the -item parameter.

If you specify no value for the -item parameter, only the PathID and the Status fields are displayed.

Table 6.8 shows the items that can be selected by using the -item parameter and the parameter values that can be specified after the -item parameter.

Table 6.8 Items that can be selected by the -path -item parameter and parameter values that can be specified after the -path -item parameter

Selectable items	Parameter values that can be specified after the -path -item parameter
PathID#	None
Status#	None
PathName	pn
DskName	dn
iLU	lu
ChaPort	cp
Type	type
IO-Count	ic
IO-Errors	ie
DNum	dnu
HDevName	hd
IEP	iep

#

Because both PathID and Status are always displayed, they do not require any parameters to be specified.

In the subsequent sub-parameters (following -path -item), you can filter the paths to be listed (-hdev) and sort the list (-srt). When you omit both parameters, the command displays information for all the paths in order of increasing AutoPATH_IDs.

The following shows the subsequent sub-parameters (following -path - item):

-hdev *host-device-name*

Filters the information only for the paths accessing the specified host device.

Specify the name of the logical device file (minus the slice number, in c*U*t*X*d*Y* format) for the HDLM device to indicate the desired host device.

The *host-device-name* string is case-sensitive.

When you specify this parameter, HDevName is displayed by default. Therefore, it is not necessary to specify hd for the -item parameter.

-stname

Use this parameter to display the model ID of the storage subsystem in the product ID element of the DskName field. When this parameter is omitted, the command displays the product ID or emulation type of the storage subsystem instead.

For more information pertaining to the storage subsystem and corresponding product ID, see Table 6.11.

When you use this parameter, DskName is displayed by default. Therefore, it is not necessary to specify dn for the -item parameter.

-srt {pn|lu|cp}

Use this parameter to sort the list of path information in ascending order, according to the specified sorting keys.

The sorting keys are as follows: the first sorting key is the name of the storage subsystem (DskName), the second sorting key is the value specified by the -srt parameter, and the third sorting key is AutoPATH_ID.

The available parameter values to specify the second sorting key are:

- pn: Path name
- lu: LU number of the storage subsystem
- cp: Port number of the channel adapter

When the -srt parameter is omitted, the path information is listed in order of ascending AutoPATH_IDs.

When you use this parameter, the items used for the sorting keys (DskName, AutoPATH_ID, and the item specified by this parameter) are displayed by default. Therefore, it is not necessary to specify these items for the -item parameter.

-t

Does not display the title for each information item.

Example

In the following example, IO-Count is selected as the display item and the path information is sorted in ascending order of the LUs.

To display a summary of path information

-path -c

When you specify the -path parameter together with the -c parameter, only the PathID, DskName, iLU, CP, Status, and Type fields are displayed from the path information managed by HDLM. The displayed contents are shortened so that each instance of path information is displayed on a single line.

For details on what is displayed in each item, see Table 6.9.

When you use the -c parameter, the number of characters that can be displayed in the product ID element of the DskName field is limited to 10. Therefore, when there are 11 or more characters in the product ID, the 8th and following characters are abbreviated to an ellipsis (...).

The following shows the subsequent sub-parameters (following -path -c):

-stname

Use this parameter to display the model ID of the storage subsystem in the product ID element of the DskName field. When this parameter is omitted, the command displays the product ID or emulation type of the storage subsystem instead.

For more information pertaining to the storage subsystem and corresponding product ID, see Table 6.11.

-srt {lu|cp}

Use this parameter to sort the list of path information in ascending order, according to the specified sorting keys.

The sorting keys are as follows: the first sorting key is the name of the storage subsystem (DskName), the second sorting key is the value specified by the -srt parameter, and the third sorting key is AutoPATH_ID.

The available parameter values to specify the second sorting key are:

- lu: LU number of the storage subsystem
- cp: Port number of the channel adapter

When the -srt parameter is omitted, the path information is listed in order of ascending AutoPATH_IDs.

-t

Does not display the title for each information item.

Example

The following example shows how to display a summary of the path information in order of iLUs.

Displayed path information

Table 6.9 describes the displayed path information. The following explains the table headings:

- No summary displayed: The user specifies the -path parameter or -path -item parameter.
- Summary displayed: The user specifies the -path -c parameter.

Table 6.9 Displayed path information

Displayed Item		Description
No summary displayed	**Summary displayed**	
Paths		Total number of displayed paths, indicated by a decimal number.
OnlinePaths		Number of available paths in the displayed paths, indicated by a decimal number. When the value of Paths equals the value of OnlinePaths, all paths are online. If the value of OnlinePaths is less than that of Paths, some paths might have an error status, in which case you should check the status of individual paths and take appropriate action for any paths that have an error status.
PathStatus		Status of the displayed paths. The displayed status indicates the following:

Displayed Item		
		▪ Online: All paths are available. ▪ Reduced: Some paths are not available. Reduced means that some paths might have an error status, in which case you should check the status of individual paths and take appropriate action for any paths that have an error status.
IO-Count		Total I/O count for the displayed paths, indicated by a decimal number. The maximum value that can be displayed is 2^{32} - 1 (4294967295). If the total I/O count reaches the maximum value, it is reset, and the count is re-started from 0.
IO-Errors		Total I/O error count for the displayed paths, indicated by a decimal number. The maximum value that can be displayed is 2^{32} - 1 (4294967295). If the total I/O error count reaches the maximum value, it is reset, and the count is re-started from 0.
PathID		AutoPATH_ID indicated by a decimal number. AutoPATH_ID is assigned when the host is restarted. When a new LU was added and the host has not been restarted, AutoPATH_ID is re-assigned to each path of the LU when you execute the dlmcfgmgr utility for managing the HDLM configuration.
PathName#	-	The path name, which indicates a physical path. When you modify the system configuration or replace a hardware item, you should check the path names to identify the physical path that will be affected by the change. Path name consists of the following four elements, separated by periods: ▪ Host port number (hexadecimal number) ▪ Bus number (hexadecimal number) ▪ Target ID (hexadecimal number) ▪ Host LU number (hexadecimal number) For details on each element of the path name and its representation in Solaris, see Table 6.10.
DskName#	DskName	Storage subsystem name, which identifies the storage subsystem that is accessed by a path. A storage subsystem name consists of the following three elements, separated by periods: ▪ Vendor ID: The name of the storage subsystem vendor (for example, HITACHI). ▪ Product ID: Indicates the storage subsystem product ID, emulation type, or model ID name (for example, OPEN-3). For more details on the storage subsystem and corresponding product ID, see Table 6.11. ▪ Serial number: The serial number of the storage subsystem (for example, 15001). You can identify an actual storage subsystem by referencing the above information from the storage subsystem management program.
iLU#	iLU	LU number of the storage subsystem. This number combined with the storage subsystem name (shown in DskName) identifies the LU that is accessed by a path. ▪ For the Freedom Storage 7700E, Lightning 9900 Series, Lightning 9900V Series, and TagmaStore USP, indicated by a hexadecimal number. The first two characters of iLU are the CU number, and the last two characters are the internal LU number within the CU. ▪ For the Freedom Storage 5700E, Freedom Storage 5800, Thunder 9200, Thunder 9500V Series, and TagmaStore AMS/WMS series, indicated by a decimal number. The entire value of iLU is the internal LU number within the storage subsystem. You can identify an actual LU by referencing iLU from the storage subsystem management program. ▪ For Universal Storage Platform V, indicated by a hexadecimal number. The first two characters of iLU are the number of the logical DKC (Disk Controller), the middle two numbers are the CU number, and the last two characters are the internal LU number

Displayed Item		
		within the CU.
ChaPort#	CP	Port number of the channel adapter, which identifies the CHA port that is mounted on the storage subsystem. You can identify an actual CHA port by referencing this number from the storage subsystem management program.
Status		Status of the path ▪ Online: Online ▪ Offline(C): Placed offline by using a command ▪ Offline(E): Offline due to an error ▪ Online(E): Failure has occurred (If none of the paths accessing one LU have an Online status, one of those paths is changed to the Online(E) status.) Paths that are Offline(E) or Online(E) require corrective action. Take appropriate action by referring to 5.3
Type#	Type	Attribute of the path ▪ Own: Owner path ▪ Non: Non-owner path When connecting to Freedom Storage 7700E, Lightning 9900 Series, Lightning 9900V Series, TagmaStore USP, or Universal Storage Platform V all paths are owner paths.
IO-Count#	--	Total I/O count for the path, indicated by a decimal number. The maximum value that can be displayed is 2^{32} - 1 (4294967295). If the total I/O count reaches the maximum value, it is reset, and the count is re-started from 0. To reset the IO-Count value to 0, execute the dlnkmgr command's clear operation. Executing the clear operation also resets the number of I/O errors (IO-Errors) to 0. For details on the clear operation, see 6.2.
IO-Errors#	--	Total I/O error count for the path, indicated by a decimal number. The maximum value that can be displayed is 2^{32} - 1 (4294967295). If the total I/O error count reaches the maximum value, it is reset, and the count is re-started from 0. To reset the IO-Errors value to0, execute the dlnkmgr command's clear operation. Executing the clear operation also clears the number of I/O operations (IO-Count) to 0. For details on the clear operation, see 6.2.
DNum#	--	A hyphen (-) is displayed.
HDevName#	--	Host device name. The name of the logical device file (minus the slice number, in c*U*t*X*d*Y* format) for the HDLM device *U*: Controller number reserved by HDLM. *X*: Target ID or WWN (World Wide Name) of the sd or ssd device for the HDLM device *Y*: LUN of the sd or ssd device for the HDLM device
IEP#	--	Information about the intermittent error. This item is displayed only when you specify -iem together the -path parameter. One of the following values is displayed for each path: ▪ - Indicates that intermittent error monitoring is disabled or the monitoring time for an intermittent error is out of range (the path status is Online(E) or Offline(E)). ▪ A value of at least 0 Indicates the number of errors that occurred during intermittent error monitoring (the path status is Online(E), Offline(E), or Online). ▪ *

Displayed Item		
		Indicates that an intermittent error occurred (automatic failback does not check the path) (the path status is `Online`, `Onlin(E)`, `Offline(E)`, or `Offline`(C)).

Legend:

--: Not displayed

#

The path information is displayed only when a value is specified for the `-path -item` parameter.

Table 6.10 Elements of a path name

Element	Solaris representation	Corresponding information
Host port number (example: 0000)	Host port number	A unique number in a host: HDLM assigns this unique number to each HBA port. When two HBAs that have two ports are installed on the host, a value from 0000 to 0003 is assigned to each port.
Bus number (example: 0000)	0 (fixed value)	(Not applicable)
Target ID (example: 000000000000001A)	Target ID	The `target` clause of the `/kernel/drv/sd.conf` file.
Host LU number (example: 0005)	Logical unit ID or LUN	The `lun` clause of the `/kernel/drv/sd.conf` file.

Table 6.11 Product ID displayed by the view -path operation

Model names of storage subsystems	Product ID	
	Without the `-stname` parameter	With the `-stname` parameter (Displays the following for the model ID)
Freedom Storage 5700E	product identifier#	5700E
Freedom Storage 5800		5800
Thunder 9200		9200
Thunder 9500V Series		9500V
TagmaStore AMS		AMS
TagmaStore WMS		WMS
Freedom Storage 7700E	Emulation type#	7700E
Lightning 9900 Series		9910/9960
Lightning 9900V Series		9970/9980
▪ TagmaStore NSC55 ▪ Hitachi TagmaStore Universal Storage Platform 100 ▪ Hitachi TagmaStore Universal Storage Platform 600 ▪ Hitachi TagmaStore Universal Storage Platform 1100		USP
SVS		SVS
Hitachi Universal Storage Platform V		USP_V
XP48		XP48
XP128		XP128
XP512		XP512
XP1024		XP1024
XP10000		XP10000
XP12000		XP12000
XP24000		XP24000

#

When a summary is displayed by specifying the `-path -c` parameter and there are more than 11 characters in the summary string, the 8th and following characters are abbreviated to an ellipsis (. . .).

6.7.2.3 To display LU information

When displaying LU information, if the -item parameter, -c parameter, or the -c -item parameter is specified at the same time as the -lu parameter, you can add and display items and display a summary of LU information. This section describes each parameter and the LU information and displayed items.

To display LU information

-lu

When neither the -c nor -item parameter is specified with the -lu parameter, the information about the LU recognized by HDLM is displayed. The sorting key is iLU and its configuration information is displayed for each LU.

By using the subsequent parameter (-hdev or -pathid), you can filter the LU information to be displayed. If you do not specify the -hdev or -pathid parameter, the information about all LUs recognized by HDLM is displayed.

For details on the contents of each display item, see Table 6.14.

The following shows the subsequent parameters:

-hdev *host-device-name*

Filters the information only for the LUs corresponding to the specified host device.

Specify the name of the logical device file (minus the slice number, in c*U*t*X*d*Y* format) for the HDLM device to indicate the desired host device.

The *host-device-name* string is case-sensitive.

-pathid *AutoPATH_ID*

Use this parameter to display only the information about the LU to which the path with the specified *AutoPATH_ID* is connected.

-t

Does not display the title for each information item.

Example

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -lu
Product      : 9500V
SerialNumber  : 0051
LUs          : 3

iLU  HDevName Device PathID Status
0001 c6t2d0   sd61   000003 Online
              sd64   000009 Online
0002 c6t2d1   sd62   000004 Online
              sd65   000010 Online
0003 c6t2d2   sd63   000005 Online
              sd66   000011 Online

Product      : 9970/9980
SerialNumber  : 15001
LUs          : 3

iLU  HDevName Device PathID Status
0670 c6t5d0   sd74   000000 Online
              sd80   000006 Offline(E)
```

```
0671 c6t5d1   sd75   000001 Online
              sd81   000007 Offline(E)
0672 c6t5d2   sd76   000002 Online
              sd82   000008 Offline(C)
KAPL01001-I The HDLM command completed normally. Operation name = view, completion time
= yyyy/mm/dd hh:mm:ss
#
```

To display LU information (by adding items to be displayed)

`-lu -item`

The items specified with the `-item` parameter are added to the items to be displayed by the `-lu` parameter and displayed.

When the value of the `-item` parameter is omitted or `all` is specified, all the items that can be displayed are displayed.

If no value or `all` is specified for the `-item` parameter, all the items that can be added are displayed.

Table 6.12 shows the items that can be added by using the `-item` parameter and the parameters that can be specified after the `-item` parameter.

Table 6.12 Items that can be added by using the `-lu -item` parameter and parameters that can be specified after the `-lu -item` parameter

Item that can be added	Subsequent parameter
SLPR	slpr
PathName	pn
ChaPort	cp
CLPR	clpr
Type	type
IO-Count	ic
IO-Errors	ie
DNum	dnu
IEP	iep
Physical Device	pd
All items are displayed	all

In the subsequent sub-parameters (`-hdev` or `-pathid`), you can filter the LU information to be displayed. When you omit both parameters, the command displays the information about all the LUs recognized by HDLM.

For details on the contents of each displayed item, see Table 6.14.

The following shows the subsequent sub-parameters:

`-hdev` *host-device-name*

Filters the information only for the LUs corresponding to the specified host device.

Specify the name of the logical device file (minus the slice number, in c*U*t*X*d*Y* format) for the HDLM device to indicate the desired host device.

The *host-device-name* string is case-sensitive.

-pathid *AutoPATH_ID*

Use this parameter to display only the information about the LU to which the path with the specified *AutoPATH_ID* is connected.

-t

Does not display the title for each information item.

Example

The following example shows how to add SLPR, PathName, ChaPort, CLPR, Type, IO-Count, IO-Errors, DNum, IEP, and Physical Device to the displayed items for LU information.

To display a summary of LU information

-lu -c

When the -c parameter is specified with the -lu parameter, a summary of LU configuration information is displayed on a line. The total number of paths recognized by HDLM and the number of online paths are displayed for each LU.

You cannot specify the -c parameter together with the -hdev or -pathid parameter.

For details on the contents of each display item, see Table 6.14.

-t

Does not display the title for each information item.

Example

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -lu -c
Product       S/N   LUs iLU  HDevName Paths  OnlinePaths
9500V         0051    3 0001 c6t2d0       2            2
                        0002 c6t2d1       2            2
                        0003 c6t2d2       2            2
9970/9980     15001   3 0670 c6t5d0       2            1
                        0671 c6t5d1       2            1
                        0672 c6t5d2       2            1
KAPL01001-I The HDLM command completed normally. Operation name = view, completion time
= yyyy/mm/dd hh:mm:ss
#
```

To display a summary of LU information (by adding items to be displayed)

-lu -c -item

The items specified with the -item parameter are added to the items to be displayed by the -lu -c parameter and displayed.

If no value is specified for the -item parameter, all the items that can be added are displayed. For the contents of each displayed item, see Table 6.14.

Table 6.13 shows the items that can be added by using the -item parameter and the parameters that can be specified after the -item parameter.

Table 6.13 Items that can be added by using the -lu -c -item parameter and parameters that can be specified after the -lu -c -item parameter

Item that can be added	Subsequent parameter
SLPR	slpr

The following show the subsequent sub-parameter:

-t

Does not display the title for each information item.

Example

The following example describes how to add and display a summary of SLPR.

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -lu -c -item
Product   S/N     LUs iLU  SLPR HDevName Paths  OnlinePaths
USP       0014050   3 0001    1 c6t2d0        2           2
                      0002    1 c6t2d1        2           2
                      0003    1 c6t2d2        2           2
9970/9980 15001     3 0670    - c6t5d0        2           1
                      0671    - c6t5d1        2           1
                      0672    - c6t5d2        2           1
KAPL01001-I The HDLM command completed normally. Operation name = view,
completion time = yyyy/mm/dd hh:mm:ss
#
```

Displayed LU information

Table 6.14 describes the displayed LU information. The following explains the table headings:

- No summary displayed: The user specifies the -lu parameter or -lu -item parameter.
- Summary displayed: The user specifies the -lu -c parameter or -lu -c -item parameter.

Table 6.14 Displayed LU information

Displayed item		Description
No summary displayed	**Summary displayed**	
Product		Model ID of the storage subsystem
Serial number	S/N	Serial number of the storage subsystem
LUs		Total number of LUs managed by HDLM among the LUs in the storage subsystem
iLU		LU number of the storage subsystem. This number combined with the storage subsystem name (shown in DskName) identifies the LU that is accessed by a path. ▪ For the Freedom Storage 7700E, Lightning 9900 Series, Lightning 9900V Series, and TagmaStore USP, indicated by a hexadecimal number. The first two characters of iLU are the CU number, and the last two characters are the internal LU number within the CU. ▪ For the Freedom Storage 5700E, Freedom Storage 5800, Thunder 9200 Series, Thunder 9500V Series, and TagmaStore AMS/WMS series, indicated by a decimal number. The entire value of iLU is the internal LU number within the storage subsystem. You can identify an actual LU by referencing iLU from the storage subsystem management program. ▪ For Universal Storage Platform V, indicated by a hexadecimal number. The first two characters of iLU are the number of the logical DKC (Disk Controller), the middle two numbers are the CU number, and the last two characters are the internal LU number within the CU.
SLPR[#1]	SLPR[#2]	The number of the SLPR to which the LU belongs, indicated by a decimal number from 0 to 31. A hyphen (–) is displayed if the storage logical partition function for the storage subsystem for the target LU is not supported.
HDevName	HDevName	Host device name. The name of the logical device file (minus the slice number, in c*U*t*X*d*Y* format) for the HDLM device *U*: Controller number reserved by HDLM. *X*: Target ID or WWN (World Wide Name) of the sd or ssd device for the HDLM device *Y*: LUN of the sd or ssd device for the HDLM device
Device	--	Type of SCSI device (sd or ssd) and corresponding instance number
PathID	--	AutoPATH_ID indicated by a decimal number. AutoPATH_ID is assigned when the host is restarted. When a new LU was added and the host has not been restarted, AutoPATH_ID is re-assigned to each path of the LU when you execute the dlmcfgmgr utility for managing the HDLM configuration.

Displayed item		
PathName#1	--	The path name, which indicates a path. When you modify the system configuration or replace a hardware item, you should check the path names to identify the path that will be affected by the change. Path name consists of the following four elements, separated by periods: ▪ Host port number (hexadecimal number) ▪ Bus number (hexadecimal number) ▪ Target ID (hexadecimal number) ▪ Host LU number (hexadecimal number) For details on each element of the path name and its representation in Solaris, see Table 6.10
ChaPort#1	--	Port number of the channel adapter, which identifies the CHA port that is mounted on the storage subsystem. You can identify an actual CHA port by referencing this number from the storage subsystem management program.
CLPR#1	--	The number of the CLPR to which the CHA port belongs, indicated by a decimal number from 0 to 31. Note that a hyphen (-) is displayed if the following items are subject to display: ▪ CHA ports in the storage subsystem that do not support cache logical partition function ▪ Paths connected to the Snapshot Image of the Copy-on-write Snapshot of the TagmaStore AMS/WMS series
Status	--	Status of the path ▪ Online: Online ▪ Offline(C): Placed offline by using a command ▪ Offline(E): Offline due to an error ▪ Online(E): Failure has occurred (If none of the paths accessing one LU have an Online status, one of those paths is changed to the Online(E) status.) Paths that are Offline(E) or Online(E) require corrective action. Take appropriate action by referring to 5.3
Type#1	--	Attribute of the path ▪ Own: Owner path ▪ Non: Non-owner path When connected to the Freedom Storage 7700E, Lightning 9900 Series, Lightning 9900V Series, TagmaStore USP, or Universal Storage Platform V, all paths are owner paths.
IO-Count#1	--	Total I/O count for the path, indicated by a decimal number. The maximum value that can be displayed is 2^{32} - 1 (4294967295). If the total I/O count reaches the maximum value, it is reset, and the count is re-started from 0. To reset the IO-Count value to 0, execute the dlnkmgr command's clear operation. Executing the clear operation also resets the number of I/O errors (IO-Errors) to0. For details on the clear operation, see 6.2.

Displayed item		
IO-Errors[#1]	--	Total I/O error count for the path, indicated by a decimal number. The maximum value that can be displayed is 2^{32} - 1 (4294967295). If the total I/O error count reaches the maximum value, it is reset, and the count is re-started from 0. To reset the IO-Errors value to 0, execute the dlnkmgr command's clear operation. Executing the clear operation also clears the number of I/O operations (IO-Count) to 0. For details on the clear operation, see 6.2.
DNum[#1]	--	A hyphen (-) is displayed.
IEP[#1]	--	The displayed paths are assumed to be in an intermittent error status and checked whether those paths are to be operated for automatic failback. One of the following values is displayed for each path: ▪ -: Indicates that intermittent error monitoring is disabled or the monitoring time for an intermittent error is out of range. ▪ A value of at least 0: Indicates the number of errors that occurred during intermittent error monitoring. *: Indicates that an intermittent error occurred (automatic failback does not check the path).
Physical Device[#1]	--	SCSI device configuration information. The names of the physical device files under the /devices directory are displayed.
--	Paths	Total number of the paths recognized by HDLM for the LU to be displayed, indicated by a decimal number
--	OnlinePaths	Number of available paths among the paths recognized by HDLM for the LU to be displayed, indicated by a decimal number. When the value of Paths is equal to the value of OnlinePaths, all paths are online. If the value of OnlinePaths is less than that of Paths, some paths might have an error status, in which case you should check the status of individual paths and take appropriate action for any paths that have an error status.

Legend:

--: Not displayed

#1

This information is displayed when one of the following conditions exist:

- The user selected the item to be displayed by using the -lu -item parameter.

- all was specified.

- No value was specified for the parameter.

#2

This information is displayed when one of the following conditions exist:

- The user selected the item to be displayed by using the -lu -c -item parameter.

- No value was specified for the parameter.

6.7.2.4 To display HBA port information

-hba

Use this parameter to display information about the paths that pass through each HBA port. HDLM displays the unique HBA port ID assigned to each HBA port. This HBA port ID is used in monitor operations, and can also be used in Offline and Online operations. Table 6.15 lists and describes the displayed items.

Table 6.15 Items displayed as HBA port information

Item	Description
HbaID	HBA port ID
Port.Bus	Port (host port number in hexadecimal) and Bus (bus number in hexadecimal), separated by a period
IO-Count	Total I/O count for the displayed paths, indicated by a decimal number. The maximum value that can be displayed is 2^{32} - 1 (4294967295). If the total I/O count reaches the maximum value, it is reset, and the count is re-started from 0.
IO-Errors	Total I/O error count for the displayed paths, indicated by a decimal number. The maximum value that can be displayed is 2^{32} - 1 (4294967295). If the total I/O error count reaches the maximum value, it is reset, and the count is re-started from 0.
Paths	Total number of displayed paths, indicated by a decimal number.
OnlinePaths	Number of available paths in the displayed paths, indicated by a decimal number. When the value of Paths equals the value of OnlinePaths, all paths are online. If the value of OnlinePaths is less than that of Paths, some paths might have an error status, in which case you should check the status of individual paths and take appropriate action for any paths that have an error status.
Physical Device	HBA port configuration information. The names of the physical device files under the /devices directory are displayed.

Example

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -hba
HbaID Port.Bus  IO-Count   IO-Errors  Paths  OnlinePaths
Physical Device
00000 0000.0000      71520          0     20          20
/devices/pci@8,70000/JNI,FCE@2
00001 0002.0000        425        100     20          10
/devices/pci@8,70000/JNI,FCE@3
00002 0001.0000       1425          0     20          20
/devices/pci@8,70000/JNI,FCE@4
KAPL01001-I The HDLM command completed normally. Operation name = view, completion time
= yyyy/mm/dd hh:mm:ss
#
```

The following shows the subsequent sub-parameters:

-srt pb

Sorts the HBA port information by the host port numbers (Port) and bus numbers (Bus).

-c

Does not display the SCSI device configuration information (Physical Device).

-t

Does not display the title for each information item.

6.7.2.5 To display CHA port information

`-cha`

Use this parameter to display information about the paths that pass through CHA port. HDLM displays the unique CHA port ID assigned to each CHA port. This CHA port ID is used in `monitor` operations, and can also be used in `Offline` and `Online` operations. Table 6.16 lists and describes the displayed items.

Table 6.16 Items displayed as CHA port information

Item	Description
`ChaID`	CHA port ID
`Product`	Model ID of the storage subsystem
`S/N`	Serial number of the storage subsystem
`ChaPort`	Port number of the channel adapter, which identifies the CHA port that is mounted on the storage subsystem. You can identify an actual CHA port by referencing this number from the storage subsystem management program.
`IO-Count`	Total I/O count for the displayed paths, indicated by a decimal number. The maximum value that can be displayed is 2^{32} - 1 (4294967295). If the total I/O count reaches the maximum value, it is reset, and the count is re-started from 0.
`IO-Errors`	Total I/O error count for the displayed paths, indicated by a decimal number. The maximum value that can be displayed is 2^{32} - 1 (4294967295). If the total I/O error count reaches the maximum value, it is reset, and the count is re-started from 0.
`Paths`	Total number of displayed paths, indicated by a decimal number.
`OnlinePaths`	Number of available paths in the displayed paths, indicated by a decimal number. When the value of `Paths` equals the value of `OnlinePaths`, all paths are online. If the value of `OnlinePaths` is less than that of `Paths`, some paths might have an error status, in which case you should check the status of individual paths and take appropriate action for any paths that have an error status.

Example

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -cha
ChaID Product          S/N              ChaPort IO-Count   IO-Errors  Paths  OnlinePaths
00000 SANRISE_USP      1100             7A           777           0     10           10
00001 SANRISE_USP      1000             7B           100           0     20           20
00002 SANRISE_USP      1100             8A             0           0     10           10
00003 SANRISE_USP      1000             8B           333          77     20           10
KAPL01001-I The HDLM command completed normally. Operation name = view, completion time =
yyyy/mm/dd  hh:mm:ss
#
```

The following shows the subsequent sub-parameters:

`-srt cp`

Sorts the CHA port information by the storage subsystem model IDs (`Product`), storage subsystem serial numbers (`S/N`), and CHA port numbers (`ChaPort`).

-t

Does not display the title for each information item.

6.7.2.6 To display corresponding information about an HDLM device, sd or ssd device, and LDEV

-drv

Displays the information about the PathID, the HDLM device, the sd or ssd device for the HDLM device, and the LDEV (the LU in the storage subsystem).

For details on the contents of each display item, see Table 6.17.

-t

Does not display the title for each information item.

Table 6.17 To display corresponding information about an HDLM device, sd or ssd device, and LDEV

Item	Description
PathID	AutoPATH_ID, indicated by a decimal number. AutoPATH_ID is assigned when the host is restarted. When a new LU was added and the host has not been restarted, AutoPATH_ID is re-assigned to each path of the LU when you execute the dlmcfgmgr utility for managing the HDLM configuration.
HDevName	Host device name. The name of the logical device file (minus the slice number, in c*U*t*X*d*Y* format) for the HDLM device *U*: Controller number reserved by HDLM. *X*: Target ID or WWN (World Wide Name) of the sd or ssd device for the HDLM device *Y*: LUN of the sd or ssd device for the HDLM device
Device	Type of SCSI device (sd or ssd) and corresponding instance number
LDEV	The model ID, serial number, and iLU number for the storage subsystem, separated by periods. Information about LDEV can identify an actual LU that is managed by HDLM.

Example

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -drv
PathID HDevName Device LDEV
000000 c6t5d0   sd74   9970/9980.15001.0670
000001 c6t5d1   sd75   9970/9980.15001.0671
000002 c6t5d2   sd76   9970/9980.15001.0672
000003 c6t5d3   sd77   9970/9980.15001.0673
000004 c6t5d4   sd78   9970/9980.15001.0674
000005 c6t5d5   sd79   9970/9980.15001.0675
000006 c6t5d0   sd80   9970/9980.15001.0670
000007 c6t5d1   sd81   9970/9980.15001.0671
000008 c6t5d2   sd82   9970/9980.15001.0672
000009 c6t5d3   sd83   9970/9980.15001.0673
000010 c6t5d4   sd84   9970/9980.15001.0674
000011 c6t5d5   sd85   9970/9980.15001.0675
KAPL01001-I The HDLM command completed normally. Operation name = view, completion time
= yyyy/mm/dd hh:mm:ss
#
```

6.7.2.7 To display view operation format

`-help`

Displays the `view` operation format

Example

6.8 monitor (Displays I/O information at a specified interval)

The `dlnkmgr` command's `monitor` operation displays I/O information for each HBA or CHA port at a specified interval. HBA ports are specified by their HBA port ID, and CHA ports are specified by their CHA port ID. You can display HBA and CHA port IDs by executing the following `view` operation:

- HBA port IDs

```
/opt/DynamicLinkManager/bin/dlnkmgr view -hba
```

For details, see 6.7.2.4 in 6.7.2.

- CHA port IDs

```
/opt/DynamicLinkManager/bin/dlnkmgr view -cha
```

For details, see 6.7.2.5 in 6.7.2.

If you perform dynamic reconfiguration of the HDLM device during the `monitor` operation, the KAPL01078-W message is displayed and the `monitor` operation stops.

6.8.1 Format

6.8.1.1 To display I/O information for each HBA port

`/opt/DynamicLinkManager/bin/dlnkmgr monitor -hbaid` ***HBA-port-ID*** `[-intvl` ***interval***`] [-count` ***count***`] [-t]`

6.8.1.2 To display I/O information for each CHA port

`/opt/DynamicLinkManager/bin/dlnkmgr monitor -chaid` ***CHA-port-ID*** `[-intvl` ***interval***`] [-count` ***count***`] [-t]`

6.8.1.3 To display the monitor operation format

```
/opt/DynamicLinkManager/bin/dlnkmgr monitor -help
```

6.8.2 Parameters

Note

If the number of paths to be displayed by the `monitor` operation is 500 or more, the execution of the `monitor` operation might be delayed as follows:

– The time displayed in `StartTime` might be later than the time the `monitor` operation is actually executed.

When the number of paths is approximately 2000, the time delay will be about seven seconds, and when it is 4000, 20 seconds.

– If you specify a shorter time value than the recommended value listed in Table 6.18 for the `-intvl` parameter, the time interval of the `monitor` operation might be several seconds to 20 seconds late for the value specified in the `-intvl` parameter.

If the number of paths is 500 or more, specify an appropriate value for the `-intvl` parameter based on the values listed in Table 6.18. However, when any of the following operations is performed concurrently with the `monitor` operation, even if you specify a value listed in the following table, the time interval of the `monitor` operation might be a few seconds late.

- The `view` operation
- Path health checking
- Automatic failback

Table 6.18 Recommended values for the `-intvl` parameter when there are 500 or more paths

Number of paths (approximate)	Recommended value specified for the -intvl parameter (seconds)
500	3 or more
1000	5 or more
2000	10 or more
4000	30 or more

6.8.2.1 To display I/O information for each HBA port

`-hbaid` ***HBA-port-ID***

Specifies the HBA port ID of the HBA port for which I/O information is to be displayed. The subsequent parameters specify the display interval and count. The format for the subsequent parameters are as follows:

`-intvl` ***interval***

Specifies the interval at which I/O information is to be displayed, in the range from 1 to 60 (seconds). When this parameter is omitted, the default (3 seconds) is applied.

`-count` ***count***

Specifies the number of times I/O information is to be displayed, in the range from 1 to 999. When this parameter is omitted, the default (10) is applied.

-t

Does not display the title for each information item.

Table 6.19 Items displayed as I/O information for each HBA port

Item	Description
HbaID	HBA port ID
Port.Bus	Port (host port number in hexadecimal) and Bus (bus number in hexadecimal) separated by a period
Physical Device	HBA port configuration information. The names of the physical device files under the /devices directory are displayed.
StartTime	Date and time in which the command began executing, in the format *yyyy*/*mm*/*dd* *hh*:*mm*:*ss*
ElapsedTime	Time elapsed since command execution began, in the format *hh*:*mm*:*ss*
IO-Count	Total I/O count for the displayed paths during a specified interval indicated by a whole number (indicated by a decimal number). The maximum value that can be displayed is 2^{32} - 1 (4294967295). If the total I/O count reaches the maximum value, it is reset, and the count is re-started from 0.
IO-Errors	Total I/O error count for the displayed paths during a specified interval, indicated by a whole number (indicated by a decimal number). The maximum value that can be displayed is 2^{32} - 1 (4294967295). If the total I/O error count reaches the maximum value, it is reset, and the count is re-started from 0.

Example

```
# /opt/DynamicLinkManager/bin/dlnkmgr monitor -hbaid 0 -intvl 30 -count 3
HbaID          : 00000
Port.Bus       : 0001.0000
Physical Device : /devices/pci@8,70000/JNI,FCE@2
StartTime      : yyyy/mm/dd hh:mm:ss

ElapsedTime IO-Count   IO-Errors
00:00:30          223         10
00:01:00         1140          1
00:01:30          557          0
KAPL01001-I The HDLM command completed normally. Operation name = monitor, completion
time = yyyy/mm/dd hh:mm:ss
```

6.8.2.2 To display I/O information for each CHA port

-chaid *CHA-port-ID*

Specifies the CHA port ID of the CHA port for which I/O information is to be displayed. The subsequent parameters specify the display interval and count. The format for the subsequent parameters are as follows:

-intvl *interval*

Specifies the interval at which I/O information is to be displayed, in the range from 1 to 60 (seconds). When this parameter is omitted, the default (3 seconds) is applied.

-count *count*

Specifies the number of times I/O information is to be displayed, in the range from 1 to 999. When this parameter is omitted, the default (10) applied.

-t

Does not display the title for each information item.

Table 6.20 Items displayed as I/O information for each CHA port

Item	Description
ChaID	CHA port ID
Product	Model ID of the storage subsystem
SerialNumber	Serial number of the storage subsystem
ChaPort	Port number of the channel adapter, which identifies the CHA port that is mounted on the storage subsystem. You can identify an actual CHA port by referencing this number from the storage subsystem management program.
StartTime	Date and time in which the command began executing, in the format *yyyy*/*mm*/*dd* *hh*:*mm*:*ss*
ElapsedTime	Time elapsed since command execution began, in the format *hh*:*mm*:*ss*
IO-Count	The total I/O count for the displayed paths during a specified interval, indicated by a whole number (indicated by a decimal number). The maximum value that can be displayed is 2^{32} - 1 (4294967295). If the total I/O count reaches the maximum value, it is reset, and the count is re-started from 0.
IO-Errors	The total I/O error count for the displayed paths during a specified interval, indicated by a whole number (indicated by a decimal number). The maximum value that can be displayed is 2^{32} - 1 (4294967295). If the total I/O error count reaches the maximum value, it is reset, and the count is re-started from 0.

Example

```
# /opt/DynamicLinkManager/bin/dlnkmgr monitor -chaid 1 -intvl 30 -count 3
ChaID         : 00001
Product       : USP
SerialNumber  : 1100
ChaPort       : 1R
StartTime     : yyyy/mm/dd hh:mm:ss

ElapsedTime IO-Count   IO-Errors
00:00:30         1154         0
00:01:00          205        22
00:01:30           54        10
KAPL01001-I The HDLM command completed normally. Operation name = monitor, completion
time = yyyy/mm/dd hh:mm:ss
```

6.8.2.3 To display monitor operation format

-help

Displays the monitor operation format.

Example

```
# /opt/DynamicLinkManager/bin/dlnkmgr monitor -help
monitor:
  Format
    dlnkmgr monitor -hbaid HBA_ID [-intvl Interval-Time]
                                        [-count Number-Of-Times] [-t]
    dlnkmgr monitor -chaid CHA_ID [-intvl Interval-Time]
                                        [-count Number-Of-Times] [-t]
```

```
  Valid value
    HBA_ID             { 00000 - 99999 }(Decimal)
    CHA_ID             { 00000 - 99999 }(Decimal)
    Interval-Time      { 1 - 60 }(Seconds)        (Default Value 3)
    Number-Of-Times    { 1 - 999 }(Times)         (Default Value 10)
KAPL01001-I The HDLM command completed normally. Operation name = monitor, completion
time = yyyy/mm/dd hh:mm:ss
#
```

Chapter 7 Utility Reference

This chapter explains the utilities used by HDLM.

- 7.1 Overview of the utilities
- 7.2 DLMgetras (utility for collecting HDLM error information)
- 7.3 dlmcfgmgr (utility for managing the HDLM configuration)
- 7.4 dlmlisthdev (utility for assisting HDLM transitions)
- 7.5 dlmsetconf (utility for creating the HDLM driver configuration definition file)
- 7.6 dlmvxexclude (utility for assisting creation of the VxVM configuration file)

7.1 Overview of the utilities

HDLM provides the following utilities:

- The `DLMgetras` utility for collecting HDLM error information

  When an error occurs, this utility collects the files that contain information to be submitted to your HDLM vendor or maintenance company. For details on the `DLMgetras` utility, see 7.2.

- The `dlmcfgmgr`utility for managing the HDLM configuration

  This utility makes the HDLM driver recognize the contents of the HDLM driver configuration definition file (`/kernel/drv/dlmfdrv.conf`) without restarting the host. For details on the `dlmcfgmgr` utility, see 7.3.

- The `dlmlisthdev` utility for assisting HDLM transitions

  This utility creates a list that shows the correspondence between the logical device file names of the HDLM devices used in HDLM version 04-00 or 04-01 and the logical device file names of the HDLM devices newly created in HDLM version 5.9. For details on the `dlmlisthdev` utility, see 7.4.

- The `dlmsetconf` utility for creating the HDLM driver configuration definition file

  This utility creates the HDLM driver configuration definition file (`/kernel/drv/dlmfdrv.conf`). For details on the `dlmsetconf` utility, see 7.5.

- The `dlmvxexclude` utility for assisting creation of the VxVM configuration file

  This utility outputs the file that has the same format as the VxVM management target exclusion file that matches the definition in the HDLM driver configuration definition file (`/kernel/drv/dlmfdrv.conf`) to the specified directory. For details on the `dlmvxexclude` utility, see 7.6.

Note

The utilities must be executed by a user with root privileges.

7.2 DLMgetras (utility for collecting HDLM error information)

This utility collects information that is needed to analyze HDLM errors that have occurred: information such as error logs, integrated trace files, trace files, definition files, core files, system crash dump files, and libraries. The collected information is archived in a file and saved to the directory that you specified. The following files are output:

- *Name-of-the-directory-containing-a-system-crash-dump-file*/*system-crash-dump-file*

  This file is a symbolic link file to a system crash dump file.

- `hbsa.tar.Z`

  This file contains compressed error information of the HiCommand products other than HDLM. This file is output only when using the HiCommand product that is the target for collecting error information.

- `getras.tar.Z`

  This file contains compressed HDLM information and system information.

For details on the information that is stored in each file, see 7.2.3.

When you want to collect information other than that in 7.2.3, define the information to collect in a file for defining the information to be collected. Information defined in a file for defining the information to be collected is compressed into `getras.tar.Z`.

When the system is restarted, a part of the information that was collected by the `DLMgetras` utility will be cleared. If an error occurs, immediately execute this utility.

7.2.1 Format

`/opt/DynamicLinkManager/bin/DLMgetras`
`{`*directory-to-which-collected-information-is-output*

`[-f` *name-of-a-file-for-defining-the-information-to-be-collected*`]|-h}`

You can also use lower-case characters (`dlmgetras`) as follows:

`/opt/DynamicLinkManager/bin/dlmgetras`
`{`*directory-to-which-collected-information-is-output*

`[-f` *name-of-a-file-for-defining-the-information-to-be-collected*`]|-h}`

7.2.2 Parameters

directory-to-which-collected-information-is-output

Specify the output directory for the information that is to be collected by the `DLMgetras` utility for collecting HDLM error information. The collected information is compiled into two files shown in *Function* and output in the specified directory. You cannot specify multi-byte characters.

-f *name-of-a-file-for-defining-the -information-to-be-collected*

Use this parameter when you want to specify certain directories or files to be collected. In this parameter, specify the name of a file for defining the information to be collected, which defines the files and directories you want to collect.

Use an absolute path to specify the name of a file for defining the information to be collected. You cannot specify multi-byte characters.

Figure 7.1 shows an example of the coding in this file.

Figure 7.1 Coding example of a file for defining the information to be collected

Rules for coding a file for defining the information to be collected

- Use an absolute path to specify a directory or file whose information is to be collected. If you use a relative path to specify a directory or file, that directory or file will be searched for only within the directory in which the DLMgetras utility was executed, and the files found will be collected.
- Do not specify a directory that contains the directory to which the collected information is output. If you specify this directory, the DLMgetras utility will run indefinitely.
- Lines beginning with a hash mark (#) are handled as comment lines.
- If a hash mark (#) is encountered anywhere other than at the beginning of a line, it is assumed to be part of the path name.
- Only one file or directory can be specified per line.
- The root directory (/) cannot be specified.
- When a directory is specified, the DLMgetras utility collects all the files in that directory, including files contained in the directory's subdirectories. If no files are found in a specified directory, the utility does not perform file collection for that directory and does not create a directory for it in the destination directory.
- Set up the specified file or directory so that it can be read by users with root privileges. The DLMgetras utility can only obtain information for a file or directory that can be read.

-h

Displays the format of the DLMgetras utility.

Note

- Because the DLMgetras utility first stores error information in the specified directory to which collected information is output before compressing, ensure that information collection areas allocated are of adequate size.
- If the specified directory to which collected information is output already exists, an overwrite confirmation message is displayed. Responding by entering y instructs the DLMgetras utility to overwrite all of the existing files. Responding by entering n (or anything other than y) instructs the utility to terminate without executing.

  In the latter case, you can either rename the existing directory before re-executing the utility, or you can execute the utility with a different directory name specified.
- If you create an archive for the system crash dump collected by the DLMgetras utility, execute the following:

```
# cd directory-to-which-collected-information-is-output
# tar cvfh archive-file-name ./var
# compress archive-file-name
```

7.2.3 List of collected error information

The following illustrates the error information collected by executing the DLMgetras utility, which is explained separately in each output file.

- *Name-of-the-directory-containing-a-system-crash-dump-file/name-of-the-system-crash-dump-file*

  This is a symbolic link file to a system crash dump file.

  Set all files in the directory specified by the DUMPADM_SAVDIR parameter that is in the /etc/dumpadm.conf file.
- hbsa.tar.Z

  This file contains compressed error information of the HiCommand products other than HDLM. This file is output only when using the HiCommand product that is the target for collecting error information.
- getras.tar.Z

  This file contains HDLM information and system information.

  When you execute the DLMgetras utility, specifying the file for defining information to be collected, the getras.tar.Z file contains the information defined in the file for defining information to be collected.

  Table 7.1 lists the detailed information stored in the getras.tar.Z file.

  Table 7.1 lists the collected information with or without specifying the file for defining information to be collected.

Table 7.1 Information stored in the getras.tar.Z file

Output directory#1	Files	Explanation
var/opt/DynamicLinkManager/log	dlmguiinst.log	Log file for scripts of HDLM remote access interface

Output directory[#1]	Files	Explanation
	dlmmgr[1-16].log	HDLM Manager log file (including the driver log data)
	dlmsetconf1.log	Execution log file of the dlmsetconf utility for creating the HDLM driver configuration definition file. Information obtained during command execution and the details of errors that occurred are obtained.
	dlmsetconf2.log	Previous generation log file of dlmsetconf1.log.
	dlmwebagent[1-*N*].log[#2]	HDLM remote access interface log file
	dlmcfgmgr1.log dlmcfgmgr2.log	Execution log file of the dlmcfgmgr utility which collects the details of error occurred during the dlmcfgmgr utility for managing the HDLM configuration execution
	hdlmtr[1-64].log	Trace file
var/opt/DynamicLinkManager/log/mmap	hdlmtr.mm	Trace management file
opt/DynamicLinkManager/config	dlmfdrv.unconf	HDLM driver unconfiguration definition file
	dlmfdrv.unconf.bak	Previous generation of the HDLM driver unconfiguration definition file. dlmfdrv.unconf.bak is created if the dlmfdrv.unconf file already exists when you attempt to create the dlmfdrv.unconf file.
	dlmfdrv.unconf_oldbackup[#3]	Backup of the existing dlmfdrv.unconf file, taken when you perform an upgrade installation from HDLM 5.6.1 or earlier.
	dlmfdrv.unconf.bak_oldbackup[#3]	Backup of the existing dlmfdrv.unconf.bak file, taken when you perform an upgrade installation from HDLM 5.6.1 or earlier.
	dlmwebagent.properties	HDLM remote access interface setup file
	dlmmgr.xml	File for setting up HDLM functionality
opt/hitachi/HNTRLib/mmap	hntr*host-name*.mm	Memory mapped file created when HDLM 04-00 or earlier is used
opt/hitachi/HNTRLib2/mmap	hntr2*host-name*.mm	Memory mapped file
opt/hitachi/HNTRLib/spool	hntr[1-16].log	Integrated trace file created when HDLM 04-00 or earlier is used
integrated-trace-file-output-directory-specified-in-the-Hitachi-Network-Objectplaza-Trace-Library-utility (Default: var/opt/hitachi/HNTRLib2/spool)	*integrated-trace-file-prefix-specified-in-the-Hitachi-Network-Objectplaza-Trace-Library-utility*2[1-16].log[#4] (default: hntr)	Integrated trace file (Default: /var/opt/hitachi/HNTRLib2/spool/hntr2[1-16].log)
kernel/drv	dlmadrv.conf	HDLM alert driver configuration definition file
	dlmfdrv.conf	HDLM driver configuration definition file
	dlmfdrv.conf.bak	Previous generation of the HDLM driver configuration definition file. dlmfdrv.conf.bak is created if the dlmfdrv.conf file already exists when you attempt to create the dlmfdrv.conf file.

Output directory[#1]	Files	Explanation
	dlmfdrv.conf_oldbackup[#3]	Backup of the existing dlmfdrv.conf file, taken when you perform an upgrade installation from HDLM 5.6.1 or earlier.
	dlmfdrv.conf.bak_oldbackup[#3]	Backup of the existing dlmfdrv.conf.bak file, taken when you perform an upgrade installation from HDLM 5.6.1 or earlier.
	dlmndrv.conf	HDLM nexus driver configuration definition file
	emlxs.conf fca-pci.conf, fcaw.conf, fjpfca.conf, fjfcif.conf, fp.conf, jnic.conf, jnic146x.conf, jfca.conf, lpfc.conf, lpfs.conf, qla2200.conf, qla2300.conf, or qlc.conf	HBA driver setup file (file names depend on the installed HBA driver)
	md.conf	Upper limit information on metadevice, metaset
	scsi_vhci.conf	Sun StorageTek Traffic Manager Software setup and configuration definition file
	sd.conf	sd (SCSI driver) configuration definition file
	ssd.conf	ssd (SCSI driver) configuration definition file
etc	coreadm.conf	Core file configuration definition file
	devlink.tab	Logical device file creation definition file
	dumpadm.conf	OS crash dump configuration definition file
	mnttab	File system mount information
	syslog.conf	syslog output directory definition
	system	Kernel patch file
	vfstab	Startup-time file system auto-mount information
	driver_aliases driver_classes inittab minor_perm name_to_major path_to_inst release	Solaris settings and configuration information files

Output directory#1	Files	Explanation
etc/cfg/fp	fabric_WWN_map	
etc/default	init	
etc/dfs	dfstab	Command settings file for execution at system boot time
	fstypes	Settings file for default file system type (NFS), and other file system types
syslog-output-directory-name	*syslog-name*	syslog files defined in /etc/syslog.conf.
getrasinfo	(Not applicable)	Directory in which command execution results are stored
	cfg.txt	Link file information for the HDLM nexus driver
	dev.txt	List of logical device file names
	devices.txt	List of physical device file names
	df-k.txt	Mount information
	df-o.txt	i node utilization status
	dlmgetomtrace.dmp	In-memory trace information)
	dlmmgr-cha.txt	Output results of dlnkmgr view -cha
	dlmmgr-hba.txt	Output results of dlnkmgr view -hba
	dlmmgr-lu.txt	HDLM LU information
	dlmmgr-lu-all.txt	HDLM LU information (including the number of times an intermittent error occurred)
	dlmmgr-path.txt	HDLM path information
	dlmmgr-path-iem.txt	HDLM path information (including the number of times an intermittent error occurred)
	dlmmgr-sys.txt	HDLM system settings
	dlmmgr-drv.txt	Display the correspondences for each HDLM device, sd or ssd device and the LDEV information
	dmesg.txt	System diagnostic message
	dumpadm.txt	Dump device configuration (output results of the dumpadm command)
	env.txt	Environment variable information
	format.txt	Output results from the format command
	isainfo-b.txt	OS type (32 bit/64 bit)
	modinfo.txt	Information on loaded drivers
	pagesize.txt	Memory page size
	pkg_lang.txt	Product information of a package
	pkginfo-l.txt	List of package information

Output directory[#1]	Files	Explanation
	prtconf-pv.txt	Hardware configuration information of a host (output results of the prtconf -pv command)
	prtconf-vP.txt	List of driver information
	prtconfMemory.txt	Physical memory size
	ps-ef.txt	Process information
	psrinfo-v.txt	CPU information
	showrev-p.txt	List of installed patches
	svcs-a.txt	Output result of the svcs -a command (Solaris 10)
	svcs-l.txt	Output result of svcs -l hdlmload, svcs -l hdlm-manager, svcs -l hdlm-initpath and svcs -l hdlm-onlinepath (Solaris 10)
	swap-l.txt	Swap region information
	swap-s.txt	Swapping utilization status
	sysdef.txt	Kernel parameter value
	ulimit-a.txt	Limits on system resources (data segments, stack segments, file descriptors) that are available to processes
	uname-a.txt	OS version
	what.txt	Output of what command
VxVM[#5]	(Not applicable)	VxVM information collection directory
	VXVM_pkginfo.txt	VxVM package information
	VXVM_vxdmp.txt	/dev/vx/dmp directory list
	VXVM_vxrdmp.txt	/dev/vx/rdmp directory list
	VXVM_vxdsk.txt	/dev/vx/dsk directory list
	VXVM_vxrdsk.txt	/dev/vx/rdsk directory list
	VXVM_vxdisklist.txt	VxVM disk list
	VXVM_vxdisklist-s.txt	Output results of vxdisk -s list
	VXVM_vxdiskpath.txt	Output results of vxdisk path (this data is collected only for VxVM 4.0 or later.)
	VXVM_vxdglist.txt	List of disk groups
	VXVM_vxprint.txt	vxprint output results
etc/vx	All directories and files in the /etc/vx directory	VxVM configuration information
var/sadm/system/admin	CLUSTER	Solaris configuration information
/var/svc/log	system-device-hdlmload:default.log system-device-hdlmload:default.log.[0-9]	Log of the hdlmload service (Solaris 10)

Output directory[#1]	Files	Explanation
	system-hdlm-initpath:default.log system-hdlm-manager:default.log.[0-9]	Log of the hdlm-initpath service (Solaris 10)
	system-hdlm-manager:default.log system-hdlm-initpath:default.log.[0-9]	Log of the HDLM Manager service (Solaris 10)
	system-hdlm-onlinepath:default.log system-hdlm-onlinepath:default.log.[0-9]	Log of the hdlm-onlinepath service (Solaris 10)
var/vxvm[#5]	vxconfigd.log.txt	Log of error messages The output destination is either /var/vxvm/vxconfigd.log or /var/adm/syslog/syslog.log, depending on the debug option setting for the user.
var/adm/syslog	syslog.log.txt	Log of error messages The output destination is either /var/vxvm/vxconfigd.log or /var/adm/syslog/syslog.log, depending on the debug option setting for the user.
VCS[#5]	(Not applicable)	VCS information collection directory
	VCS_pkginfo.txt	VCS package information
	VCS_hastatus-summary.txt	VCS status information
	VCS_haclus-display.txt	VCS cluster information
	VCS_hasys-list.txt	System information in the VCS cluster
	VCS_hagrp-display.txt	Service group information
	VCS_hares-display.txt	Release information
var/VRTSvcs/log[#5]	engine_[A, ...].log	VCS engine log
SDS[#5]	(Not applicable)	SDS information collection directory
	SDS_pkginfo.txt	SDS package information
	SDS_metadb.txt	Status database information
	SDS_metastat.txt	Metadevice information
	SDS_metaset.txt	Diskset information
	SDS_metastat-s.txt	metastat -s *diskset-name* information
etc/lvm[#5]	md.tab	SDS or SVM input parameter file
	md.cf	SDS or SVM configuration backup file
	mddb.cf	Information on SDS or SVM status database copy storage locations
	mdlogd.cf	Log collected from SDS or SVM

Output directory[#1]	Files	Explanation
SVM[#5]	(Not applicable)	SVM information collection directory
	SVM_pkginfo.txt	SVM package information
	SVM_metadb.txt	Status database information
	SVM_metastat.txt	Volume information
	SVM_metaset.txt	Diskset information
	SVM_metastat-s.txt	metastat -s *diskset-name* information
etc/opt/SUNWmd[#5]	md.tab	SDS input parameter file
	md.cf	SDS configuration backup file
	mddb.cf	Information on SDS status database copy storage locations
SC3.0[#5]	(Not applicable)	Sun Cluster information collection directory
	SC3.0_pkginfo.txt	Package information
	SC3.0_scstat-q.txt	Cluster status snapshot
	SC3.0_scconf-p.txt	Cluster configuration information
	SC3.0_scdidadm-L.txt	Mapping information between HDLM drivers and device IDs
	SC3.0_dev_did_dsk.txt	List of directories and files in the /dev/did/dsk directory
	SC3.0_dev_did_rdsk.txt	List of directories and files in the /dev/did/rdsk directory
etc/cluster[#5]	All directories and files in the /etc/cluster directory	Cluster configuration information for Sun Cluster

#1

An output directory is created in the specified directory when you extract the getras.tar.Z file.

#2

The value *N* depends on the setting in the dlmwebagent.properties file. The default value is 8.

#3

The file might not be created depending on the device managed by HDLM.

#4

In the actual file name, a file number is appended to *Trace-file-prefix specified-in-the-Hitachi-Network-Objectplaza-Trace-Library-utility*2. For example, the default will be hntr21.log to hntr216.log. Note that the 2 following the integrated trace file prefix does not represent a file number.

#5

This information is collected only if the software for which information is to be collected is installed.

7.3 dlmcfgmgr (utility for managing the HDLM configuration)

This utility makes the HDLM driver recognize the contents of the HDLM driver configuration definition file (`/kernel/drv/dlmfdrv.conf`) dynamically without restarting the host. Table 7.2 illustrates the configuration changes applicable and not applicable to the `dlmcfgmgr` utility for managing the HDLM configuration.

Table 7.2 Configuration changes applicable and not applicable to the dlmcfgmgr utility for managing the HDLM configuration

Applicable operation	Inapplicable operation
▪ Adding a new LU ▪ Managing the LU that has been removed from HDLM management	▪ Deleting an existing LU ▪ Changing the settings of an existing LU (changing an LUN mapping or storage port) ▪ Adding a path ▪ Deleting a path ▪ Removing an existing LU from HDLM management

Once an LU is recognized by the HDLM driver, HDLM can use the LU without restarting the host.

If the definition of the HDLM driver configuration definition file (`/kernel/drv/dlmfdrv.conf`) has been changed due to configuration changes by deleting or changing LUs and paths, the changes are not applied to the HDLM driver even though you execute the `dlmcfgmgr` utility. Restart the host with the reconfiguration option specified, and then change the configuration.

Executing the `dlmcfgmgr` utility outputs the KAPL10373-I message. Entering `y` in response to this message performs dynamic reconfiguration of the HDLM devices and then deletes the logical device files for the sd or ssd devices managed by HDLM. Also, the detailed information is output to the `/var/opt/DynamicLinkManager/log/dlmcfgmgr1.log`. Entering `n` in response to this message cancels the dynamic reconfiguration processing.

The `dlmcfgmgr` utility is valid only in a Solaris 8 version1/01 or later or solaris 9 system.

7.3.1 Format

```
/opt/DynamicLinkManager/bin/dlmcfgmgr { -a | -h }
```

Parameters

`-a`

Applies the contents of the HDLM driver configuration definition file (`/kernel/drv/dlmfdrv.conf`) to HDLM driver. The added LU becomes valid after the command with this option is executed.

-h

Displays help about the usage of the command.

Examples

Example 1: The -a option is specified:

Information on the newly added LUs is applied to the HDLM driver based on the contents of the HDLM driver configuration definition file (/kernel/drv/dlmfdrv.conf). Information on LUs that are already managed by HDLM is not updated.

```
# /opt/DynamicLinkManager/bin/dlmcfgmgr -a
```

Example 2: Options are omitted:

In this example, the dlmcfgmgr utility outputs an error message and displays help about the usage of the command.

```
# /opt/DynamicLinkManager/bin/dlmcfgmgr
KAPL10380-E A parameter is invalid.
KAPL10368-I usage: dlmcfgmgr {-a | -h}
#
```

Note

- Execute the dlmsetconf utility for creating the HDLM driver configuration definition file before executing the dlmcfgmgr utility.
- If you use VxVM, do not use the dlmcfgmgr utility. Instead, specify the reconfiguration option, and then restart the host. To register the HDLM device to VxVM, you need to restart with the reconfiguration option specified.
- The correspondence between the paths and path IDs may differ depending on whether dynamic reconfiguration of the HDLM device is performed by executing the dlmcfgmgr utility or by restarting the host.
- You must restart the host after performing an upgrade installation or re-installation. LUs will not be added dynamically if you execute the dlmcfgmgr utility before restarting the host.
- If the HDLM driver configuration definition file (/kernel/drv/dlmfdrv.conf) contains the definition of a newly added LU, the dlmcfgmgr utility makes the HDLM driver recognize the new LU so that the LU is to be managed by HDLM.
- Dynamic reconfiguration is not performed for the HDLM driver for deleted LUs, LUs whose path is added or deleted, and LUs whose path setting (such as the instance number) is changed. If you have deleted an LU, or added, deleted, or changed a path, restart the host specifying the reconfiguration option to change the configuration.
- To remove the LU managed by HDLM from HDLM management, you need to restart the host specifying the reconfiguration option. Do not use the dlmcfgmgr utility.
- If you want to specify that some of the newly added LUs are not to be managed by HDLM, you need to remove these LUs from HDLM management before executing the dlmcfgmgr utility. If the newly added LUs have been removed from HDLM management, reconfiguration of these LUs is not performed for the HDLM driver.

- After HDLM device reconfiguration is performed, executing the `dlmcfgmgr` utility will not dynamically add LUs if one of the following occurs:

  -For the LU, you changed the LUN mapping, the mapping-destination storage port, or the host bus adapter port on the recognizing host.

  -You mapped another logical device to a LUN of the storage port being used by the configured LU.

  -When adding a new LU, you allocated an existing instance number to the path of the added LU.

  -You dynamically added a host bus adapter.

  -The major or minor number of the sd or ssd device for an LU to be newly added matches the major or minor number of an sd or ssd device previously registered in HDLM.
- If the HDLM driver configuration definition file (`/kernel/drv/dlmfdrv.conf`) does not contain HDLM device information, the `dlmcfgmgr` utility outputs the message and terminates.
- If no change is made to the HDLM-managed devices when the `dlmcfgmgr` utility is executed, the utility outputs the message and terminates.
- You cannot concurrently execute more than one `dlmcfgmgr` utility on the same host. If you attempt such execution, the utility outputs the message and terminates.

7.4 dlmlisthdev (utility for assisting HDLM transitions)

Use the `dlmlisthdev` utility when migrating to HDLM 5.9 from HDLM 04-00 or 04-01 installed in a Solaris 8 or 9 environment. The utility creates a correspondence table between the logical device file names of the HDLM devices used in version 04-00 or 04-01 and those of the HDLM devices created in version 5.9.

7.4.1 Format

`/opt/DynamicLinkManager/bin/dlmlisthdev` *input-file-name*

7.4.2 Parameters

input-file-name

Specifies, as an input file, the name of the output file that stores the output result of the `dlnkmgr` command's `view -path` operation executed before HDLM 04-00 or 04-01 was uninstalled. You cannot specify multi-byte characters.

Examples

Example

Referring to the output results of the dlnkmgr command's view -path operation, which are saved in the /var/tmp/old_view_path_list.txt file, results in the creation of a list that shows the correspondences of the logical device file names of HDLM devices.

```
# /opt/DynamicLinkManager/bin/dlmlisthdev /var/tmp/old_view_path_list.txt
```

When upgrading HDLM 04-00 to HDLM 5.9 in a Solaris 8 environment

To make HDLM devices usable in a new environment when HDLM is upgraded from 04-00 to 5.9 in Solaris 8:

1. Save the logical device file names of the HDLM devices used in HDLM 04-00.

   Before uninstalling HDLM 04-00, execute the dlnkmgr command's view -path operation to save the output result in a file that has any desired file name.

   ```
   # /opt/DynamicLinkManager/bin/dlnkmgr view -path > /var/tmp/old_view_path_list.txt
   ```

2. After uninstalling HDLM 04-00, use the reconfiguration option to restart the host.
3. Install HDLM 5.9.
4. To confirm that HDLM has been successfully installed, check the path configuration in HDLM 5.9.
5. Create a list that shows the correspondences between the logical device file names of the HDLM devices in HDLM 04-00 and those in HDLM 5.9.

   As the input file for the command, specify the file that stores the output result of the dlnkmgr command's view -path operation executed before HDLM 04-00 was uninstalled. Executing the following command creates a list for the correspondence of the logical device file names, and displays the result in the standard output. The result of the command execution is as follows:

```
# /opt/DynamicLinkManager/bin/dlmlisthdev /var/tmp/old_view_path_list.txt
OldHDevName NewHDevName DskName                                  iLU
c3t0d10     c5t0d10     HITACHI .DF500F          .0207       0010
c3t0d11     -           HITACHI .DF500F          .0207       0011
-           c5t0d12     HITACHI .DF500F          .0207       0012
...
#
```

Table 7.3 lists and describes the items to be displayed.

Table 7.3 Information displayed by the dlmlisthdev utility for assisting HDLM transitions

Item	Explanation
OldHDevName	The logical device file name (HDevName) of an HDLM device in HDLM 04-00 is displayed. If a hyphen (–) is displayed, it means that the corresponding HDev did not exist in the previous environment.
NewHDevName	The logical device file name (HDevName) of an HDLM device in HDLM

	5.9 is displayed. If a hyphen (–) is displayed, it means that the corresponding HDev does not exist in the new environment.
DskName	Indicates the storage subsystem name
iLU	Indicates the LU number in the storage subsystem

6. Register the HDLM device.

   Use the list showing the correspondence between the logical device file names of the HDLM devices in HDLM 04-00 and those in HDLM 5.9, created in step 5, to verify the logical device file names of the HDLM devices to be registered. After verifying the file names, if needed, register the HDLM devices to the application. For details on how to register, see 3.11.

Note

The file that stores the output result of the `dlnkmgr` command's `view -path` operation must be created in a directory that is not deleted when HDLM is uninstalled. If not, the file will be deleted during uninstallation. The following are the directories that are deleted during uninstallation. Make sure that you do not specify any of these directories as the output destination for the file:

`/opt/DynamicLinkManager`

`/opt/hitachi`

`/var/opt/DynamicLinkManager`

`/var/opt/hitachi`

7.5 dlmsetconf (utility for creating the HDLM driver configuration definition file)

This `dlmsetconf` utility obtains device information retained by Solaris, and defines the storage subsystem's sd or ssd device entries in the HDLM driver configuration definition file (`/kernel/drv/dlmfdrv.conf`) and the HDLM nexus driver configuration definition file (`/kernel/drv/dlmndrv.conf`).

The sd or ssd device entries are sorted in ascending order of the LDEV information (for example, 9970/9980.15001.066A).

If an existing HDLM driver configuration definition file (`/kernel/drv/dlmfdrv.conf`) is available, the existing definitions are inherited. If an existing HDLM nexus driver configuration definition file (`/kernel/drv/dlmndrv.conf`) is available, the existing definitions are inherited. An HDLM driver unconfiguration definition file (`/opt/DynamicLinkManager/config/dlmfdrv.unconf`) is created.

All HDLM-managed devices are defined in the HDLM driver unconfiguration definition file (`/opt/DynamicLinkManager/config/dlmfdrv.unconf`). The devices in the file are managed in units of LUs and begin with a hash mark (#) as comment lines. When you want to specify that a device is not to be managed by HDLM, use a text editor to edit the HDLM driver unconfiguration definition file (`/opt/DynamicLinkManager/config/dlmfdrv.unconf`) and remove the comment-out status in the LDEV line (that is, delete `#` at the beginning of the line).

When you want to specify that a disk not managed by HDLM is to be managed by HDLM, use a text editor to edit the HDLM driver unconfiguration definition file (/opt/DynamicLinkManager/config/dlmfdrv.unconf) and then comment out the LDEV line (that is, add # at the beginning of the line). In either case, do not make any changes to the file except for adding or deleting # at the beginning of the LDEV line.

When the configuration of the storage subsystem is modified (for example, when LUs are added or deleted) or when the hardware configuration on the host side is modified, it is necessary to execute the dlmsetconf utility, and then use one of the following commands to restart the host in order to reconfigure the devices. The HDLM devices are reconfigured.

```
# reboot -- -r
```

At the ok prompt, execute the following:

```
boot -r
```

In the subsequent explanations, the configuration definition file is indicated as the HDLM driver configuration definition file (/kernel/drv/dlmfdrv.conf) and the HDLM nexus driver configuration definition file (/kernel/drv/dlmndrv.conf).

7.5.1 Format

```
/opt/DynamicLinkManager/bin/dlmsetconf [-d output-directory-name] [-r | -u]
```

7.5.2 Parameters

-d *output-directory-name*

Specifies the name of the directory to which the configuration definition file created by the dlmsetconf utility is to be output. This directory must be specified as an absolute path. You cannot specify multi-byte characters. If *output-directory-name* is omitted, the configuration definition file will be created in the /kernel/drv directory. Create the output directory before you execute the utility.

-r

Specify this parameter to invalidate the definition of an existing configuration definition file to create a new configuration definition file and an HDLM driver unconfiguration definition file. When this parameter is omitted, definition information from the existing configuration definition file and HDLM driver unconfiguration definition file is inherited, the information on the changed configuration is added to and/or deleted from the definitions in the HDLM driver configuration definition file (/kernel/drv/dlmfdrv.conf), and the information on the changed configuration is added to the definitions in the HDLM nexus driver configuration definition file (/kernel/drv/dlmndrv.conf).

Do not specify the -r parameter when SDS or SVM is used. A change in instance numbers (that is, the minor numbers for HDLM devices) due to HDLM device reconfiguration will make metadevices and volumes unavailable. If the -r parameter is not used, the instance numbers defined in the existing configuration definition files are retained, and the metadevices and volumes continue to be available even though the configuration of the storage subsystem has been modified.

This parameter cannot be used together with the -u parameter.

Note

When the reconfiguration option is specified to restart the host after executing the dlmsetconf utility with the -r parameter specified, a new logical device file might be created. If a new logical device file is created, delete the old logical device file.

-u

Depending on the contents of the HDLM driver unconfiguration definition file (/opt/DynamicLinkManager/config/dlmfdrv.unconf) that defines the LU you want to remove from HDLM management, in the existing entries of the HDLM driver configuration definition file (/kernel/drv/dlmfdrv.conf), only the attribute value that specifies whether or not the LU is managed by HDLM is updated, and applied in the system when the reconfiguration option is specified to restart the host.

When the -u parameter is specified, the entry is not added or deleted even though the system configuration is changed (such as by adding or deleting LUs). If the attribute value that specifies whether or not the LU is managed by HDLM is not changed when the dlmsetconf utility is executed with the -u parameter specified, the HDLM driver configuration definition file (/kernel/drv/dlmfdrv.conf) is not updated. If the HDLM driver unconfiguration definition file (/opt/DynamicLinkManager/config/dlmfdrv.unconf) contains an invalid entry or does not exist, an error occurs. This parameter cannot be specified together with the -r parameter.

Examples

Example 1

Parameters are omitted:

A configuration definition file is created in the /kernel/drv directory. If the configuration definition file already exists, the existing definition information is inherited, and the parts for the changed configuration are added and/or deleted.

```
# /opt/DynamicLinkManager/bin/dlmsetconf
```

Example 2

The -d parameter is specified:

A configuration definition file is created in the /tmp directory:

```
# mkdir /tmp/dlmdir
# /opt/DynamicLinkManager/bin/dlmsetconf -d /tmp/dlmdir
```

Example 3

The -r parameter is specified:

A new configuration definition file is created in the /kernel/drv directory. Existing information is not inherited.

```
# /opt/DynamicLinkManager/bin/dlmsetconf -r
```

Example 4

The -u parameter is specified:

Defines a non- HDLM-managed LU based on the contents of the definitions in the HDLM driver unconfiguration definition file (/opt/DynamicLinkManager/config/dlmfdrv.unconf):

```
# /opt/DynamicLinkManager/bin/dlmsetconf -u
```

Note

- Execution of the dlmsetconf utility causes sd or ssd devices to be defined in the HDLM driver configuration definition file (/kernel/drv/dlmfdrv.conf) in units of LUs.
- Executing the dlmsetconf utility outputs the KAPL10242-I message. Entering y in response to this message creates or updates the HDLM driver configuration definition file (/kernel/drv/dlmfdrv.conf), the HDLM driver unconfiguration definition file (/opt/DynamicLinkManager/config/dlmfdrv.unconf), and the HDLM nexus driver configuration definition file (/kernel/drv/dlmndrv.conf). Entering n in response to this message does not create or update those files.
- When the reconfiguration option is specified to restart the host after executing the dlmsetconf utility, HDLM deletes the HDLM-managed logical device files for sd or ssd devices.
- The HDLM driver configuration definition file (/kernel/drv/dlmfdrv.conf) and HDLM nexus driver configuration definition file (/kernel/drv/dlmndrv.conf) must not be edited. Editing these configuration definition files may result in a failure to create the correct logical device files for HDLM devices.
- If the dlmfdrv.conf and dlmndrv.conf files exist in either the /kernel/drv directory or in a user-defined directory, the files are renamed dlmfdrv.conf.bak and dlmndrv.conf.bak respectively, and saved as the previous generation of the configuration definition files. If the dlmfdrv.unconf file exists either in the /opt/DynamicLinkManager/config directory or in a user-defined directory, the file is renamed dlmfdrv.unconf.bak and saved as the previous generation of the definition file.
- Even when the -d parameter is specified, existing definition information is read from the /kernel/drv/dlmfdrv.conf and /kernel/drv/dlmndrv.conf files, rather than from the dlmfdrv.conf and dlmndrv.conf files in the specified directory.
- If the dlmsetconf utility finds an added path, in the HDLM driver configuration definition file (/kernel/drv/dlmfdrv.conf) the utility adds an entry (that is, a definition) to the definition of the LU for which the path has been added. If any path has been deleted, the entry (definition) is deleted from the definition of the affected LU. Whether a path is added or deleted, existing instance numbers are inherited unless a new HDLM driver configuration definition file (/kernel/drv/dlmfdrv.conf) file is created.

- If the dlmsetconf utility finds any storage subsystems connected to a newly added HBA, in the HDLM nexus driver configuration definition file (/kernel/drv/dlmndrv.conf) the utility adds the entry (definition) for the new HBA. Even when an HBA, to which one of the Hitachi storage subsystems is connected, is deleted, the entry (definition) remains. Unless a new HDLM nexus driver configuration definition file (/kernel/drv/dlmndrv.conf) is created, the existing instance numbers are inherited, and the entry is not deleted.
- If no configuration definition files (/kernel/drv/dlmfdrv.conf and /kernel/drv/dlmndrv.conf) are found, the configuration definition files (/kernel/drv/dlmfdrv.conf and /kernel/drv/dlmndrv.conf) are created regardless of whether or not the -r parameter is specified.
- If an LU is deleted, you must delete the logical device files of the HDLM devices corresponding to the deleted LU, which are stored in the /dev/dsk or /dev/rdsk directory, before executing the dlmsetconf utility. If the dlmsetconf utility is executed before the logical device files of the HDLM devices corresponding to the deleted LU are deleted, then the format command may incorrectly display the logical device file name of an HDLM device corresponding to the deleted LU, for the newly recognized LU.
- When using SDS or SVM and executing the dlmsetconf utility after making changes in the path configuration, you must not specify the -r parameter. If you do so, the metadevices and volumes might be unavailable.
- Executing the dlmsetconf utility with the -u parameter specified does not update the unconf property in the HDLM driver configuration definition file (/kernel/drv/dlmfdrv.conf) if no change is made to the HDLM-managed devices. After the reconfiguration option is specified to restart the host, the HDLM driver configuration definition file (/kernel/drv/dlmfdrv.conf) that was updated takes effect in the system.
- If you want HDLM to manage a device that is not currently managing, comment out the LDEV line (add a hash mark (#) at the beginning of the line) of the definitions for the desired LU in the HDLM driver unconfiguration definition file (/opt/DynamicLinkManager/config/dlmfdrv.unconf), specify the reconfiguration option, and then restart the host.
- When you specify that a device is, or is not, to be managed by HDLM, edit the HDLM driver unconfiguration definition file (/opt/DynamicLinkManager/config/dlmfdrv.unconf) to execute the dlmsetconf utility with the -u option specified, and then restart the host.
- If the below message appears after you execute the dlmsetconf utility, it means the device ID could not be acquired from the SCSI driver. If this happens, restart the host and then re-execute the dlmsetconf utility.

```
KAPL10237-E The property information (aa...aa) on a device (bb...bb) could not be collected.
```

aa...aa: property name

bb...bb: sd or ssd device name

If the device ID cannot be acquired from the SCSI driver, refer to the bug report for Bug ID 4897065 provided by SUNSOLVE ONLINE (http://sunsolve.sun.com/) for more information. To view the bug report, you need a SUNSOLVE ONLINE account.

7.6 dlmvxexclude (utility for assisting creation of the VxVM configuration file)

This utility outputs the file (vxvm_exclude.txt) having the same format as the VxVM management target exclusion file (/etc/vx/vxvm.exclude) that matches the definition in the HDLM driver configuration definition file (/kernel/drv/dlmfdrv.conf) to the directory specified by the -d parameter. If the -d parameter is not specified, the dlmvxexclude utility for assisting creation of the VxVM configuration file outputs the contents of the file to the standard output.

7.6.1 Format

```
/opt/DynamicLinkManager/bin/dlmvxexclude [-d output-directory-name| -h ]
```

7.6.2 Parameters

-d *output-directory-name*

Specifies the path name of the directory to which the file is to be output. You cannot specify multi-byte characters. If the directory specified here contains a regular file having the same name as vxvm_exclude.txt, the confirmation prompt appears. Entering Y or y creates a backup file with the name vxvm_exclude.txt.bak and creates a new vxvm_exclude.txt file. An error occurs if a file name is specified for the output directory, the directory does not exist, or the output directory name is not specified with the -d parameter.

-h

Displays help about the usage of the command.

Examples

Example 1

In this example, the dlmvxexclude utility outputs the settings of the VxVM management exclusion for the contents of the current HDLM configuration definition file (/kernel/drv/dlmfdrv.conf) to standard output:

```
# /opt/DynamicLinkManager/bin/dlmvxexclude
```

Example 2

In this example, the dlmvxexclude utility outputs the settings of the VxVM management exclusion for the contents of the current HDLM configuration definition file (/kernel/drv/dlmfdrv.conf) to the /tmp/vxvm exclude.txt file.

```
# /opt/DynamicLinkManager/bin/dlmvxexclude -d /tmp
```

Display example

Example

The entries for the following devices are output:

- sd or ssd devices for HDLM-managed devices
- HDLM devices that are not actually created (when using VxVM 3.1.1, or VxVM 3.5 or later)

  Only one HDLM device is created for one HDLM-managed device (one slice). However, VxVM assumes that the number of HDLM devices created for one HDLM-managed device is the same as the number of physical paths. Thus, executing the `dlmvxexclude` utility outputs entries for HDLM devices that have not actually been created.

Entries for these devices are output for each LU, and are sorted in alphabetic order, using the logical device file name as a key.

```
exclude_all 0
paths
c3t0d10s2 /pci@8,600000/pci@1/JNI,FCE@4/sd@0,a
c3t0d11s2 /pci@8,600000/pci@1/JNI,FCE@4/sd@0,b
c3t0d12s2 /pci@8,600000/pci@1/JNI,FCE@4/sd@0,c
c3t0d13s2 /pci@8,600000/pci@1/JNI,FCE@4/sd@0,d
c4t1d10s2 /pci@8,600000/pci@1/JNI,FCE@5/sd@1,a
c4t1d11s2 /pci@8,600000/pci@1/JNI,FCE@5/sd@1,b
c4t1d12s2 /pci@8,600000/pci@1/JNI,FCE@5/sd@1,c
c4t1d13s2 /pci@8,600000/pci@1/JNI,FCE@5/sd@1,d
c6t1d10s2 /pseudo/dlmndrv@2/dlmfdrv@1,a
c6t1d11s2 /pseudo/dlmndrv@2/dlmfdrv@1,b
c6t1d12s2 /pseudo/dlmndrv@2/dlmfdrv@1,c
c6t1d13s2 /pseudo/dlmndrv@2/dlmfdrv@1,d
#
controllers
#
product
#
pathgroups
#
```

Note

- Only a user with root privileges can execute the `dlmvxexclude` utility for assisting creation of the VxVM configuration file.
- After installing HDLM, you need to execute the `dlmvxexclude` utility after executing the `dlmsetconf` utility for creating the HDLM driver configuration definition file.
- Do not use the output result of the `dlmvxexclude` utility as is for the `vxvm.exclude` file. If a device other than the HDLM-managed devices has already been invalidated in VxVM, then the invalidated setting might be canceled.

Chapter 8 Messages

This chapter provides information about viewing messages output by HDLM. It also lists and explains the HDLM messages and shows the actions to be taken in response to each message.

Section 8.15 describes the meanings and actions to be taken of return codes output by HDLM when the HDLM remote access interface sends a request to HDLM.

- 8.1 Before Viewing the List of Messages
- 8.2 HDLM Command (dlnkmgr and Operations) Messages
- 8.3 HDLM API Messages
- 8.4 HDLM Manager Messages
- 8.5 HDLM Driver (Filter Component) Messages or HDLM Nexus Driver Messages
- 8.6 HDLM Alert Driver Messages
- 8.7 HDLM Driver (Core Logic Component) Messages
- 8.8 HDLM Management Target Messages
- 8.9 HDLM Installation Program Messages
- 8.10 Messages from the DLMgetras Utility for Collecting HDLM Error Information
- 8.11 Messages from the dlmsetconf Utility for Creating the HDLM Driver Configuration Definition File or the dlmlisthdev Utility for assisting HDLM Transitions
- 8.12 Messages from the dlmcfgmgr Utility for Managing the HDLM Configuration
- 8.13 Messages from the dlmvxexclude Utility for Assisting the Creation of the VxVM Configuration File
- 8.14 Messages from HDLM remote access interface
- 8.15 Return Codes for the HDLM Remote Access Interface

8.1 Before Viewing the List of Messages

This section explains the following information that is needed to locate messages and understand the message explanations in the sections from 8.2.

- Format and meaning of message IDs
- Terms used in messages and message explanations

The information is described below.

8.1.1 Format and Meaning of Message IDs

Each message has a message ID. The following shows the format and meaning of message IDs.

Table 8.1 Format and Meaning of the Message ID KAPL*mmnnn-l*

Format	Meaning
KAPL	Indicates that the message is an HDLM message.
mm	Number of the HDLM module that issued the message. 01: HDLM command 03: HDLM API 04: HDLM manager 05: HDLM driver (filter component) or HDLM nexus driver 06: HDLM alert driver 07: HDLM driver (core logic component) 08: HDLM management target 09: HDLM installation program 10: The following utilities: ▪ DLMgetras utility (KAPL100*nn*) ▪ dlmsetconf utility, dlmlisthdev utility (KAPL102*nn*) ▪ dlmcfgmgr utility (KAPL103*nn*) ▪ dlmvxexclude utility (KAPL107*nn*) 11: HDLM remote access interface
nnn	Message serial number for the module
l	Message level C: Critical E: Error W: Warning I: Information

8.1.2 Terms Used in Messages and Message Explanations

The following shows the terms that may appear in messages and message explanations.

Table 8.2 Terms Used in Messages and Message Explanations

Terms	Meaning
aa...aa	Variable (If there is more than one variable in a single message, they appear in succession as follows: *bb...bb*, *cc...cc*, and so on.)
Operation name	Type of the operation that is entered after dlnkmgr in the command.

8.1.3 Components That Output Messages to Syslog

Some messages for the following components are output to syslog:

- HDLM manager
- HDLM driver (filter component)

- HDLM alert driver
- HDLM management target

8.2 HDLM Command (dlnkmgr and Operations) Messages

KAPL01001-I	The HDLM command completed normally. Operation name = *aa...aa*, completion time = *bb...bb*	Details The HDLM command completed successfully. *aa...aa*: `clear`, `help`, `monitor`, `offline`, `online`, `set`, `view` *bb...bb*: Year (4 numeric characters)/Month (01-12)/Day (01-31) hour:minute:second Action None.
KAPL01002-I	The HDLM command started. Operation name = *aa...aa*	Details The HDLM command was executed. *aa...aa*: `clear`, `monitor`, `offline`, `online`, `set`, `view` Action None.
KAPL01003-W	No operation name is specified.	Action Specify the operation name, and then retry.
KAPL01004-W	The operation name is invalid. Operation name = *aa...aa*	Details *aa...aa*: Specified operation name Action Execute the `help` operation of the HDLM command (`dlnkmgr`) to check the operation name, and then retry. For details on the `help` operation, see 6.3.
KAPL01005-W	A parameter is invalid. Operation name = *aa...aa*, parameter = *bb...bb*	Details *aa...aa*: `clear`, `set`, `monitor`,`online`, `offline`, `view` *bb...bb*: Specified parameter Action Execute `help` *operation-name* of the HDLM command (`dlnkmgr`) to check the parameter, and then retry. For details on the `help` operation, see 6.3.

KAPL01006-W	A necessary parameter is not specified. Operation name = *aa...aa*	Details The specified operation does not contain the necessary parameter. *aa...aa*: `clear`, `set`, `monitor`, `offline`, `view` Action Execute `help` *operation-name* of the HDLM command (`dlnkmgr`) to check the parameter. Specify the correct parameter, and then retry. For details on the `help` operation, see 6.3.
KAPL01007-W	A duplicate parameter is specified. Operation name = *aa...aa*, parameter = *bb...bb*	Details *aa...aa*: `clear`, `monitor`,`offline`, `online`, `set`, `view` *bb...bb*: Duplicate parameter Action Delete the duplicate parameter, and then retry.
KAPL01008-W	A necessary parameter value is not specified. Operation name = *aa...aa*, parameter = *bb...bb*	Details *aa...aa*: `monitor`,`offline`, `online`, `set`, `view` *bb...bb*: Parameter name Action Specify the parameter value, and then retry.
KAPL01009-W	A parameter value is invalid. Operation name = *aa...aa*, parameter = *bb...bb*, parameter value = *cc...cc*, Valid value = *dd...dd*	Details *aa...aa*: `monitor`,`offline`, `online`, `set`, `view` *bb...bb*: Parameter name *cc...cc*: Specified parameter value *dd...dd*: Specifiable parameter value range Action Specify the correct value for the parameter, and then retry.

KAPL01012-E	Could not connect the HDLM manager. Operation name = *aa...aa*	Details In the `view -sys -sfunc` operation, information must be collected from the HDLM manager but the manager cannot be accessed. *aa...aa*: `view` Action Execute the `view` operation of the HDLM command (`dlnkmgr`) to check whether the HDLM manager has started. Start the HDLM manager if it has not started, and then retry the HDLM command. For details on the `view` operation, see 6.7.
KAPL01013-E	An error occurred in internal processing of the HDLM command. Operation name = *aa...aa* details = *bb...bb*	Details An error whose cause does not seem to be a user operation occurred during command processing. *aa...aa*: `clear`, `monitor`, `offline`, `online`, `set`, `view` *bb...bb*: Function name and contents of error Action Execute the `DLMgetras` utility for collecting HDLM error information, and then contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM. For details on the `DLMgetras` utility, see 7.2.
KAPL01014-W	No authority to execute the HDLM command. Operation name = *aa...aa*	Details You are not authorized to execute the HDLM command as the administrator. *aa...aa*: `clear`, `monitor`, `offline`, `online`, `set`, `view` Action Execute the command as a user with root privileges.

KAPL01015-W	The target HBA was not found. Operation name = *aa...aa*	Details The path having the port number and path number specified in the `-hba` parameter could not be found. *aa...aa*: `offline`, `online` Action Execute the `view` operation of the HDLM command (`dlnkmgr view -path`) and check the value displayed in `PathName`. Specify the two leftmost digits of `PathName` for the relevant HBA port, and then retry. For details on the `view` operation, see 6.7.
KAPL01016-W	The target CHA port was not found. Operation name = *aa...aa*	Details The path ID indicated by `-pathid` and required by the `-cha` parameter is not an object of HDLM management. *aa...aa*: `monitor`, `offline`, `online` Action Execute the `view` operation of the HDLM command (`dlnkmgr view -path`) to check the target CHA port and the AutoPATH_ID that passes through that CHA port. Then, specify the corresponding AutoPATH_ID, and then retry the operation. For details on the `view` operation, see 6.7.
KAPL01018-W	The target device was not found. Operation name = *aa...aa*	Details The specified host device name could not be found. *aa...aa*: `view` Action Execute the `view` operation of the HDLM command (`dlnkmgr view -path`) to check the value displayed in `HDevName`. Specify the value of `HDevName` for the host device to be operated and then retry. For details on the `view` operation, see 6.7.

KAPL01019-W	The target path was not found. Operation name = *aa...aa*	Details *aa...aa*: `offline`, `online`, `monitor`, `view` ▪ `offline`/`online` operation The specified path does not exist. ▪ `view`/`monitor` operation The paths have not been configured because creation of the HDLM environment or configuration changes to the HDLM operating environment have not finished. Action ▪ `offline`/`online` operation Use the `view` operation of the HDLM command (`dlnkmgr`) to check the specification, and then retry. For details on the `view` operation, see 6.7. ▪ `view`/`monitor` operation Refer to Chapter 3 or 4.4, and then configure the path. If the same error occurs again, execute the `DLMgetras` utility for collecting HDLM error information, acquire the error information, and then contact your HDLM vendor or the company for which you have a service contract. For details on the `DLMgetras` utility, see 7.2 *DLMgetras (utility for collecting HDLM error information)*.
KAPL01021-E	Cannot execute the HDLM command due to insufficient memory.	Details Memory required for HDLM command processing could not be allocated. Action Terminate unnecessary applications to increase the amount of free memory, and then retry.
KAPL01023-W	The last Online path for the device cannot be placed Offline(C).	Details The path specified in the `offline` operation cannot be placed Offline(C) because it is the last path for the LU. Action Use the `view` operation of the HDLM command (`dlnkmgr`) to check the status of the path. For details on the `view` operation, see 6.7.

KAPL01024-W	The specified parameters cannot be specified at the same time. Operation name = *aa...aa*, parameters = *bb...bb*	Details *aa...aa*: `monitor`, `offline`, `online`, `set`, `view` *bb...bb*:Parameters cannot be specified at the same time Action Execute `help` *operation-name* of the HDLM command (`dlnkmgr`) to check the parameter that can be specified, and then retry. For details on the `help` operation, see 6.3.
KAPL01036-E	The Offline path cannot be placed online. PathID = *aa...aa*	Details *aa...aa*: Path ID (decimal number) Action Remove the error in the path, and then retry.
KAPL01039-W	During the online operation processing of the HDLM command, a path that cannot be placed in the Online status was detected. PathID = *aa...aa* Would you like to continue the processing of the online operation? [y/n]:	Details A path that cannot be placed Online was detected during multipath online processing. To ignore this path and perform online processing for the next path, enter `y`. To cancel processing, enter `n`. *aa...aa*: Path ID (decimal number) Action If you want to continue processing of the `online` operation of the HDLM command for other paths, enter `y`. If you want to terminate the processing, enter `n`. For details on the `online` operation, see 6.5.
KAPL01040-W	The entered value is invalid. Re-enter [y/n]:	Details A value other than y and n was entered. Enter y or n. Action Enter `y` or `n`.
KAPL01041-E	The entered value is invalid. The operation stops. Operation name = *aa...aa*	Details Command processing will be aborted because an incorrect response was made three times in reply to the request. *aa...aa*: `clear`, `offline`, `online`, `set` Action To execute the operation, re-execute the HDLM command.

KAPL01044-W	A duplicate parameter value is specified. Operation name = *aa...aa*, parameter = *bb...bb*, parameter value = *cc...cc*	Details *aa...aa*: `view` *bb...bb*: Parameter name *cc...cc*: Duplicated parameter value Action Delete the duplicate parameter value name, and then retry.
KAPL01045-W	Too many parameter values are specified. Operation name = *aa...aa*, parameters = *bb...bb*, parameter value = *cc...cc*	Details *aa...aa*: `monitor`, `offline`, `online`, `set`, `view` *bb...bb*: Parameter name *cc...cc*: Parameter value Action Execute `help` *operation-name* of the HDLM command (`dlnkmgr`) to check the parameter value, and then retry. For details on the `help` operation, see 6.3.
KAPL01048-W	Help information cannot be found. Operation name = *aa...aa*.	Details The specified operation is not an operation of the HDLM command. *aa...aa*: Specified operation name Action Use the `help` operation of the HDLM command (`dlnkmgr`) to check the operation name. And then retry. For details on the `help` operation, see 6.3.
KAPL01049-I	Would you like to execute the operation? Operation name = *aa...aa* [y/n]:	Details The `clear`/`set` operation will be started. To continue the operation, enter `y`. To cancel the operation, enter `n`. *aa...aa*: `clear`, `set` Action If you want to execute operation of the HDLM command, enter `y`. If you want to terminate the processing, enter `n`. For details on the `clear` operation, see 6.2. For details on the `set` operation, see 6.6.
KAPL01050-I	The currently selected paths will be changed to the Online status. Is this OK? [y/n]:	Details The `online` operation will be started. To continue the `online` operation, enter `y`. To cancel the operation, enter `n`. Action If you want to execute the online processing, enter `y`. If you want to terminate the processing, enter `n`. For details on the `online` operation, see 6.5.

KAPL01051-I	Because no path has been selected among the currently displayed paths, the paths in the Offline(C), Offline(E), and Online(E) statuses will be changed to the Online status. Is this OK? [y/n]:	Details All the paths will be placed `Online` because the path selection parameter is not specified in the `online` operation. To place all the paths Online, enter `y`. To not place them online, enter `n`. Action If you want to execute the online processing, enter `y`. If you want to terminate the processing, enter `n`. For details on the `view` operation, see 6.7. For details on the `online` operation, see 6.5.
KAPL01052-I	The currently selected paths will be changed to the Offline(C) status. Is this OK? [y/n]:	Details The `offline` operation will be started. To continue the `offline` operation, enter `y`. To cancel the operation, enter `n`. Action If you want to execute the offline processing, enter `y`. If you want to terminate the processing, enter `n`. For details on the `offline` operation, see 6.4.
KAPL01053-I	If you are sure that there would be no problem when the path is placed in the Offline(C) status, enter y. Otherwise, enter n. [y/n]:	Details The `offline` operation will be started. To continue the `offline` operation, enter `y`. To cancel the operation, enter `n`. Action If you want to execute the offline processing, enter `y`. If you want to terminate the processing, enter `n`. For details on the `offline` operation, see 6.4.
KAPL01054-W	During the offline operation processing of the HDLM command, a path that cannot be placed in the Offline(C) status was detected. PathID = *aa...aa* Would you like to continue the processing of the offline operation? [y/n]:	Details A path that cannot be placed Offline(C) was detected during multipath offline processing. To ignore this path and perform offline processing for the next path, enter `y`. To cancel offline processing, enter `n`. *aa...aa*: Path ID (decimal number) Action If you want to continue processing of the `offline` operation of the HDLM command for other paths, enter `y`. If you want to terminate the processing, enter `n`. For details on the `offline` operation, see 6.4.

KAPL01055-I	All the paths which pass the specified *aa...aa* will be changed to the Offline(C) status. Is this OK? [y/n]:	Details Multiple paths will be collectively placed Offline(C) because the `-hba` or `-cha` parameter was specified. To collectively place multiple paths Offline(C), enter `y`. To not collectively place them Offline(C), enter `n`. *aa...aa*: `CHA port`, `HBA` Action If you want to execute the offline processing for the paths that pass the specified target, enter `y`. If you want to terminate the processing, enter `n`.
KAPL01056-I	If you are sure that there would be no problem when all the paths which pass the specified *aa...aa* are placed in the Offline(C) status, enter y. Otherwise, enter n. [y/n]:	Details This message re-asks the user whether to place all the paths Offline(C). To place all the paths Offline(C), enter y. To not place them Offline(C), enter n. *aa...aa*: `CHA port`, `HBA` Action If you want to execute the offline processing for the paths that pass the specified target, enter `y`. If you want to terminate the processing, enter `n`.
KAPL01057-I	All the paths which pass the specified *aa...aa* will be changed to the Online status. Is this OK? [y/n]:	Details Multiple paths will be collectively placed Online because the `-hba` or `-cha` parameter was specified. If you want to continue, enter `y`. To cancel the operation, enter `n`. *aa...aa*: `CHA port`, `HBA` Action If you want to execute the online processing for the paths that pass the specified target, enter `y`. If you want to terminate the processing, enter `n`.
KAPL01058-W	The specified parameter value is not needed. Operation name = *aa...aa*, parameter = *bb...bb*, parameter value = *cc...cc*	Details A parameter value was specified in a parameter that does not need any parameter value. *aa...aa*: `clear`, `monitor`, `offline`, `online`, `set`, `view` *bb...bb*: Parameter name *cc...cc*: Parameter value Action Execute `help` *operation-name* of the HDLM command (`dlnkmgr`) to check the parameter and parameter value, and then retry. For details on the `help` operation, see 6.3.

KAPL01059-W	Cannot specify the parameter *aa...aa* at the same time if you specify parameter *bb...bb* and parameter value *cc...cc*. Operation name = *dd...dd*	Details A parameter value conflicts with the specification of another parameter. *bb...bb*: Parameter name *cc...cc*: Parameter value *aa...aa*: Parameter name *dd...dd*: `view`, `set` Action Execute `help` *operation-name* of the HDLM command (`dlnkmgr`) to check the parameter and parameter value, and then retry. For details on the `help` operation, see 6.3.
KAPL01060-I	The user terminated the operation. Operation name = *aa...aa*	Details The command processing will be aborted because n was entered in reply to the acknowledgment. *aa...aa*: `online`, `offline`, `set`, `clear` Action None.
KAPL01061-I	*aa...aa* path(s) were successfully placed *bb...bb*; *cc...cc* path(s) were not. Operation name = *dd...dd*	Details This message indicates the number of the paths processed in the `online`/`offline` operation. *aa...aa*: Number of paths where `online`/`offline` operation is successful (decimal number) *bb...bb*: Online or Offline(C) *cc...cc*: Number of paths where `online`/`offline` is unsuccessful (decimal number) *dd...dd*: `online`, `offline` Action None. For details on the `online` operation, see 6.5. For details on the `offline` operation, see 6.4.
KAPL01063-I	The target path(s) are already *aa...aa*.	Details As a result of executing the `online`/`offline` operation, the specified path is already placed Online/Offline(C). *aa...aa*: Online or Offline(C) Action Use the `view` operation of the HDLM command (`dlnkmgr`) to check the status of the path. For details on the `view` operation, see 6.7. For details on the `online` operation, see 6.5. For details on the `offline` operation, see 6.4.

KAPL01068-I	Enter a license key:	Details The license key will now be renewed. Enter a license key. Action None.
KAPL01069-W	The entered license key is invalid.	Details The entered license key is invalid. Action Enter a valid license key.
KAPL01070-E	The entered license key is invalid. Renewal of the license key will now stop.	Details The license key renewal will be aborted because an invalid license key was entered three times. Action Obtain a valid license key, and then retry.
KAPL01071-I	The permanent license was installed.	Details The license was renewed into a permanent license. Action None.
KAPL01072-I	The emergency license was installed. The license expires on *aa...aa*.	Details A license was renewed into the emergency license. *aa...aa*: Year (4 numeric characters)/Month (01-12)/Day (01-31) Action Install a permanent license by the expiration day.
KAPL01073-E	The temporary license expired.	Details The temporary license expired. Register a permanent license. Action Register a permanent license.
KAPL01074-E	The emergency license expired.	Details The emergency license expired. Register a permanent license. Action Register a permanent license.
KAPL01075-E	A fatal error occurred in HDLM. The system environment is invalid.	Details The license information file is missing. Action Re-install HDLM.

KAPL01076-I	The permanent license has been installed.	Details You need not install a license because a permanent license has already been installed. Action None.
KAPL01078-W	The operation terminated because the path configuration changed during execution of the HDLM command. Operation name = aa...aa	Details *aa...aa*: `monitor` Action After the dynamic reconfiguration of the HDLM device has finished, re-execute the HDLM command.
KAPL01079-W	The intermittent error monitoring function cannot be set up because automatic failback is disabled.	Details The intermittent error monitoring function cannot be set up because automatic failback is disabled. Action Enable automatic failback, and then re-execute.
KAPL01080-W	The error monitoring interval and the number of times that the error is to occur conflict with the automatic failback checking interval.	Details An intermittent error cannot be detected by using the values specified for the following: the checking interval for automatic failback, the error monitoring interval, and the number of times the error is to occur. Action Set the intermittent error monitoring interval to a value that is equal to or more than (*automatic-failback-checking-interval* x *number-of-times-error-is-to-occur-for-intermittent-error-monitoring*).
KAPL01081-E	The license key file is invalid. File name = *aa...aa*	Details The format of the license key file is invalid. *aa...aa*: `/var/tmp/hdlm_license` Action Save the correct license key file in the designated, and then re-execute. `/var/tmp/hdlm_license`

KAPL01082-E	There is no installable license key in the license key file. File name = *aa...aa*	Details There is no HDLM-installable license key in the license key file. *aa...aa*: `/var/tmp/hdlm_license` Action Make sure that the license key file is correct, and then re-execute. `/var/tmp/hdlm_license`
KAPL01083-I	There is no license key file. File name = *aa...aa*	Details There is no license key file in the designated directory: *aa...aa*: `/var/tmp/hdlm_license` Action When the message that prompts you to enter the license key is displayed, enter the license key. Alternatively, cancel the HDLM command, save the correct license key file in the designated directory, and then re-execute the HDLM command. *aa...aa*: `/var/tmp/hdlm_license`
KAPL01084-W	An attempt to delete the license key file has failed. File name = *aa...aa*	Details *aa...aa*: `/var/tmp/hdlm_license` Action If a license key file exists, delete it. `/var/tmp/hdlm_license`
KAPL01088-W	The specified parameter values cannot be specified at the same time. Operation name = *aa...aa*, parameter = *bb...bb*, parameter values = *cc...cc*	Details *aa...aa*:`view` *bb...bb*: Parameter name *cc...cc*: Parameter values cannot be specified at the same time Action Execute `help` *operation-name* of the HDLM command (`dlnkmgr`) to check the parameter that can be specified, and then retry. For details on the `help` operation, see 6.3.

KAPL01089-E	One of the following was executed at the same time as an HDLM command set -lic operation: another set -lic operation, or an update of the license for an update installation.	Action Check the license by using the HDLM command's `view -sys -lic` operation. Then, if necessary, re-execute the HDLM command's `set -lic` operation. If the same error message is output, contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM. Do not perform the following operations: ▪ Simultaneous executions of the HDLM command's `set -lic` operation ▪ Execution of the HDLM command's `set -lic` operation simultaneously with an update of the license for an upgrade installation or re-installation
KAPL01090-E	The Offline path cannot be placed Online because a path error occurred. path ID = *aa...aa*, details = *bb...bb*	Details An attempt to recover the path failed because an error occurred in the path. *aa...aa*: Path ID *bb...bb*: Code showing reason for error Action Resolve the error in the relevant path, and then re-execute.
KAPL01091-E	The Offline path cannot be placed Online because memory is insufficient. path ID = *aa...aa*, details = *bb...bb*	Details An attempt to recover the path failed because of insufficient memory. Online processing will be stopped if online processing is running for multiple paths. *aa...aa*: Path ID *bb...bb*: Code showing reason for error Action Re-execute. If the error occurs again, confirm that the system memory is sufficient.

KAPL01092-E	The Offline path cannot be placed Online because a Reservation Conflict occurred. path ID = *aa...aa*, Details = *bb...bb*	Details An attempt to recover the path failed because the LU for the path to be put online has been reserved by another path. *aa...aa*: Path ID *bb...bb*: Code showing reason for error Action Use the `format` command to confirm that the LU for the path to be put online is not reserved by another host. If the LU is incorrectly reserved, remove the reservation, and then re-execute the `online` operation for the HDLM command. If the LU is not incorrectly reserved or if the same error occurs even after removing the reservation, execute the `DLMgetras` utility for collecting HDLM error information to collect the error information, and then contact your HDLM vendor or maintenance company if you have a maintenance contract.
KAPL01093-E	The offline path cannot be placed online because an attempt to read the disk label failed. path ID = *aa...aa*, details = *bb...bb*	Details An attempt to recover the path failed because the correct disk label has not been set in the LU for the path to be put online. *aa...aa*: Path ID *bb...bb*: Code showing reason for error Action Use the `format` command to set the VTOC disk label in the LU for the relevant path, and then re-execute.
KAPL01094-E	The Offline path cannot be placed Online because an error occurred. path ID = *aa...aa*, details = *bb...bb*	Details An attempt to recover the path failed because an error occurred while accessing the SCSI driver. *aa...aa*: Path ID *bb...bb*: Code showing reason for error Action Re-execute. If the same error occurs even after removing the reservation, execute the `DLMgetras` utility for collecting HDLM error information to collect the error information, and then contact your HDLM vendor or maintenance company if you have a maintenance contract.

KAPL01095-E	An attempt to acquire the HDLM version information has failed. details = *aa...aa*	Details *aa...aa*: Code showing reason for error Action Re-execute. If the same error occurs again, execute the `DLMgetras` utility for collecting HDLM error information, acquire the error information, and then contact your HDLM vendor or the company for which you have a service contract.
KAPL01096-E	An attempt to acquire the Service Pack version information has failed. details = aa...aa	Details *aa...aa*: Code showing reason for error Action Re-execute. If the same error occurs again, execute the `DLMgetras` utility for collecting HDLM error information, acquire the error information, and then contact your HDLM vendor or the company for which you have a service contract.
KAPL01097-W	All the current trace files will be deleted. Is this OK? [y/n]	Details If you set a value less than the current value of the trace file size or number of trace files, all the current trace files will be deleted. To continue the operation, enter `y`. To cancel the operation, enter `n`. Action If you want to execute operation of the HDLM command, enter `y`. If you want to terminate the processing, enter `n`.
KAPL01100-I	*aa...aa*	Details This message indicates the executed command line. Action None.
KAPL01101-W	The target HBA port was not found. Operation name = *aa...aa*	Details The HBA having the HBA_ID specified in the `-hbaid` parameter could not be found. *aa...aa:* `monitor`, `offline`, `online` Action Use the view operation of the HDLM command (`dlnkmgr view -hba`) to check the target HBA port and the HBA_ID of the target HBA port. After that, specify the appropriate HBA_ID, and then retry.

KAPL01102-I	All the paths which pass the specified *aa...aa* port will be changed to the Offline(C) status. Is this OK? [y/n]:	Details Multiple paths will be collectively placed Offline(C) because the `-hbaid` or `-chaid` parameter was specified. To collectively place multiple paths Offline(C), enter `y`. To not collectively place them Offline(C), enter `n`. *aa...aa:* CHA, HBA Action If you want to execute the offline processing for the paths which pass the specified target, enter `y`. If you want to terminate the processing, enter `n`.
KAPL01103-I	If you are sure that there would be no problem when all the paths which pass the specified *aa...aa* port are placed in the Offline(C) status, enter y. Otherwise, enter n. [y/n]:	Details This message re-asks the user whether to place all the paths Offline(C). To place all the paths Offline(C), enter `y`. To not place them Offline(C), enter `n`. *aa...aa:* CHA, HBA Action If you want to execute the offline processing for the paths which pass the specified target, enter `y`. If you want to terminate the processing, enter `n`.
KAPL01104-I	All the paths which pass the specified *aa...aa* port will be changed to the Online status. Is this OK? [y/n]:	Details Multiple paths will be collectively placed Online because the `-hbaid` or `-chaid` parameter was specified. To collectively place multiple paths Online, enter `y`. To not collectively place them Online, enter `n`. *aa...aa:* CHA, HBA Action If you want to execute the Online processing for the paths which pass the specified target, enter `y`. If you want to terminate the processing, enter `n`.
KAPL01107-I	The load balancing type specified for individual LUs will become invalid when this operation is executed. Do you want to execute the operation anyway? Operation name = set [y/n]:	Action If you want to set the load balancing type for the system, enter `y`. If you want to terminate the processing, enter `n`.

8.3 HDLM API Messages

KAPL03001-I	HDLM API information - *aa...aa*	Details This information is required for determining the cause of the problem (if any). *aa...aa*: Trace information Action None.
KAPL03003-E	HDLM API Error information - *aa...aa*	Details This information is required for determining the cause of the problem (if any). *aa...aa*: API trace error information Action Execute the `DLMgetras` utility for collecting HDLM error information, and then contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM. For details on the `DLMgetras` utility, see 7.2.
KAPL03004-C	A critical error occurred in the HDLM API. (*aa...aa*)	Details This information is required for determining the cause of the problem (if any). *aa...aa*: API trace error information Action Execute the `DLMgetras` utility for collecting HDLM error information, and then contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM. For details on the `DLMgetras` utility, see 7.2.
KAPL03006-E	An access to the HDLM driver causes an error. (*aa...aa*)	Details This information is required for determining the cause of the problem (if any). *aa...aa*: API trace error information Action Execute the `DLMgetras` utility for collecting HDLM error information, and then contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM. For details on the `DLMgetras` utility, see 7.2.

KAPL03007-E	An error occurred during communication with the HDLM manager. (*aa...aa*)	Details This information is required for determining the cause of the problem (if any). *aa...aa*: API trace error information Action Execute the `DLMgetras` utility for collecting HDLM error information, and then contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM. For details on the `DLMgetras` utility, see 7.2.
KAPL03008-E	An error occurred during log input to the HDLM alert driver. (*aa...aa*)	Details This information is required for determining the cause of the problem (if any). *aa...aa*: API trace error information Action Execute the `DLMgetras` utility for collecting HDLM error information, and then contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM For details on the `DLMgetras` utility, see 7.2.
KAPL03999-E	An unexpected error occurred.	Details A conflict occurred in the versions of the modules that HDLM uses internally. Action Execute the `DLMgetras` utility for collecting HDLM error information, and then contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM. For details on the `DLMgetras` utility, see 7.2.

8.4 HDLM Manager Messages

KAPL04001-I	HDLM manager started.	Action None.
KAPL04002-E	Could not start the HDLM manager.	Details HDLM manager failed to start because the environment is incorrect for the manager to run properly. Action Execute the `DLMgetras` utility for collecting HDLM error information, and then contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM. For details on the `DLMgetras` utility, see 7.2 .
KAPL04003-E	The startup parameter is invalid.	Details The parameter held internally by the HDLM manager is incorrect. Action Execute the `DLMgetras` utility for collecting HDLM error information, and then contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM. For details on the `DLMgetras` utility, see see 7.2.
KAPL04004-I	HDLM manager will now terminate.	Action None.
KAPL04008-E	Cannot open the option definition file (*aa...aa*).	Details HDLM manager cannot start normally (unable to open the option definition file). *aa...aa*: Option definition file name Action Check whether another program is using the file (or has opened the file with Notepad), or whether the file has been deleted inadvertently.

KAPL04009-E	The option definition is invalid.	Details HDLM manager cannot start normally (some of the definition in the option definition file is invalid). Action If the KAPL04033-W message is output after this message, execute `dlnkmgr view -sys -sfunc` and check the option settings. For options with setting values that have returned to default values, use the `dlnkmgr set` operation to reset the values. If the KAPL04033-W message is not output, restart HDLM Manager. If the same error occurs, re-install HDLM.For details on the `view` operation, see 6.7. For details on the `set` operation, see 6.6.
KAPL04010-E	Could not open the error log file.	Details HDLM manager cannot start normally (unable to open the error log file `/var/opt/DynamicLinkManager/log/dlmmgr[1-16].log`). Action Check whether another program is using the file (or has opened the file with Notepad), or whether the error log file has been deleted inadvertently.
KAPL04011-E	Could not output the error log file.	Details The log information could not be output to the error log file `/var/opt/DynamicLinkManager/log/dlmmgr[1-16].log`. Action Check that the disk has sufficient free space.
KAPL04012-E	Could not create a communication pipe. RC = *aa...aa*	Details HDLM manager cannot start normally (unable to create a pipe file to be used in communication with HDLM commands). *aa...aa*: OS error code (decimal number) Action Execute the `DLMgetras` utility for collecting HDLM error information, and then contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM. For details on the `DLMgetras` utility, see 7.2.

KAPL04013-E	Input is impossible via the communication pipe. RC = *aa...aa*	Details Data could not be read from the pipe file during the communication with the HDLM command. *aa...aa*: OS error code (decimal number) Action Execute the `DLMgetras` utility for collecting HDLM error information, and then contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM. For details on the `DLMgetras` utility, see 7.2.
KAPL04014-E	Output is impossible via the communication pipe. RC = *aa...aa*	Details Data could not be written to the pipe file during the communication with the HDLM command. *aa...aa*: OS error code (decimal number) Action Execute the `DLMgetras` utility for collecting HDLM error information, and then contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM. For details on the `DLMgetras` utility, see 7.2.
KAPL04019-E	Could not collect the error information. RC = *aa...aa*	Details An attempt to read the log information from the alert driver failed. *aa...aa*: API return code (decimal number) Action Execute the `DLMgetras` utility for collecting HDLM error information, and then contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM. For details on the `DLMgetras` utility, see 7.2.
KAPL04021-I	HDLM manager information - *aa...aa*	Details This information is required for determining the cause of the problem (if any). *aa...aa*: HDLM manager trace information Action None.

KAPL04022-W	HDLM manager warning information - *aa...aa*	Details This information is required for determining the cause of the problem (if any). *aa...aa*: HDLM manager trace warning information Action Execute the `DLMgetras` utility for collecting HDLM error information, and then contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM. For details on the `DLMgetras` utility, see 7.2.
KAPL04023-E	HDLM manager error information - *aa...aa*	Details This information is required for determining the cause of the problem (if any). *aa...aa*: HDLM manager trace error information Action Execute the `DLMgetras` utility for collecting HDLM error information, and then contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM. For details on the `DLMgetras` utility, see 7.2.
KAPL04024-C	A critical error occurred in the HDLM manager. (*aa...aa*)	Details This information is required for determining the cause of the problem (if any). *aa...aa*: HDLM manager trace error information Action Execute the `DLMgetras` utility for collecting HDLM error information, and then contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM. For details on the `DLMgetras` utility, see 7.2.
KAPL04025-C	A memory shortage occurred in the HDLM manager.	Details Sufficient memory to process the HDLM manager could not be obtained. Action To increase the amount of free memory, terminate unnecessary applications or restart the host.

KAPL04026-I	The temporary license is valid. The license expires in *aa...aa* days on (*bb...bb*).	Details *aa...aa*: Number of remaining valid days *bb...bb:* Expiration day, Year (4 numeric characters)/Month (01-12)/Day (01-31) Action Install the permanent license by the expiration day.
KAPL04027-I	The emergency license is valid. The license expires in *aa...aa* days on (*bb...bb*).	Details *aa...aa*: Number of remaining valid days *bb...bb:* Expiration day, Year (4 numeric characters)/Month (01-12)/Day (01-31) Action Install the emergency license by the expiration day.
KAPL04028-E	The temporary license expired.	Action Install the permanent license.
KAPL04029-E	The emergency license expired.	Action Install the permanent license.
KAPL04030-E	The temporary license has already expired.	Action Install the permanent license.
KAPL04031-E	The emergency license has already expired.	Action Install the permanent license.
KAPL04032-C	A fatal error occurred in HDLM. The system environment is invalid	Details A part of the HDLM configuration file is missing. Action Re-install HDLM.
KAPL04033-W	The option definition file was re-created.	Details An option definition file was re-created using the default values. The specified values are set when some of the options have been read. Action As for the options other than the defaults, use the `dlnkmgr set` operation to set the options again.For details on the `set` operation, see 6.6.

KAPL04034-E	An attempt to create the option definition file has failed.	Details An attempt to re-create an option definition file (`/etc/opt/DynamicLinkManager/dlmmgr.xml`) using the default values has failed. Action Remove unnecessary files to secure free space on the file system, or check the write permissions for the directory and file.
KAPL04035-I	The path health check will now start. Total number of paths = *aa...aa*	Details *aa...aa*: Total number of paths Action None.
KAPL04036-I	The path health check for the path *aa...aa* was executed. Number of error paths = *bb...bb*	Details *aa...aa*: Number of paths targeted for the path health check. *bb...bb*: Number of error paths by the path health check. Action None.
KAPL04037-I	The path health check completed normally. Path ID = *aa...aa*	Details There were no error paths as a result of executing the path health check. *aa...aa*: PathID for executing the path health check. Action None.
KAPL04042-I	HDLM SNMP TRAP information - *aa...aa*	Details *aa...aa*: Start or Stop Action None.
KAPL04045-I	HDLM SNMP TRAP was sent. Trap ID = *aa...aa*, IP Address = *bb...bb*, Port Number= *cc...cc*, Community = *dd...dd*, Trap Data = *ee...ee*	Details *aa...aa*: Trap ID *bb...bb*: Destination IP address of the trap *cc...cc*: Destination port number of the trap *dd...dd*: Community Name given to the trap *ee...ee*: Send data Action None.

8.5 HDLM Driver (Filter Component) Messages or HDLM Nexus Driver Messages

KAPL05003-I	The HDLM driver (filter component) was successfully attached to Disk (*aa...aa*), Partition (*bb...bb*).	Details The path for the disk and partition was registered successfully in the core logic. *aa...aa*: minor number of the sd or ssd device (decimal number) *bb...bb*: slice number Action None.
KAPL05008-E	Could not allocate memory. (*aa...aa*:*bb...bb*)	Details The OS memory reservation function started but returned an error. *aa...aa*: File ID, line number (hexadecimal number) *bb...bb*: Target memory size (hexadecimal number) Action Check whether the HDLM driver has started normally. If it has not started or contains an error, execute the `DLMgetras` utility for collecting HDLM error information, and then contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM. For details on the `DLMgetras` utility, see 7.2.
KAPL05011-E	Could not attach the HDLM driver (filter component) to Disk (*aa...aa*), Partition (*bb...bb*). (*cc...cc*:*dd...dd*)	Details An attempt to register the path for the disk and partition into the core logic has failed. *aa...aa*: Minor number of the sd or ssd device (decimal number) *bb...bb*: Slice number (decimal number) *cc...cc*: Error code (hexadecimal number) *dd...dd*: Address of table for managing the HDLM filter driver (hexadecimal number) Action Check whether the HDLM driver has started normally. If it has not started or contains an error, contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM and report the error and detail code.

KAPL05014-I	The device object (*aa...aa*) was registered as the path (*bb...bb*).	Details The HDLM filter driver was successfully attached and registered as a path. *aa...aa*: Address of table for managing the HDLM filter driver (hexadecimal number) *bb...bb*: Path ID (hexadecimal number) Action None.
KAPL05018-W	The FO processing in the path (*aa...aa*) failed. (*bb...bb*:*cc...cc*)	Details *aa...aa*: Core logic path identifier for the failed FO (hexadecimal number) *bb...bb*: Error code (hexadecimal number) *cc...cc*: 0 (fixed) Action The I/O being processed is discarded. Check the status of the device path and take an appropriate action.
KAPL05019-I	The FO processing in the path (*aa...aa*) finished. The I/O request was processed in the path (*bb...bb*).	Details *aa...aa*: Path ID (hexadecimal number) *bb...bb*: Path ID (hexadecimal number) Action None.
KAPL05020-I	Processing of IOCTL(*aa...aa*) will now start.	Details An IOCTL request was received.. *aa...aa*: IOCTL code (hexadecimal number) Action None.
KAPL05021-I	Processing of IOCTL(*aa...aa*) completed normally.	Details The processing for the requested IOCTL operation was successful. *aa...aa* (hexadecimal number): IOCTL code Action None.

KAPL05023-E	Could not process the IOCTL(*aa...aa*). (*bb...bb*:*cc...cc*)	Details An attempt to process the requested IOCTL call has failed. *aa...aa*: IOCTL code (hexadecimal number) *bb...bb*: Error code (hexadecimal number) *cc...cc*: 0 (fixed) Action Check the message of the HDLM command (`dlnkmgr`) or the HDLM manager, and then take the appropriate action. If you do not know the appropriate action, execute the `DLMgetras` utility for collecting HDLM error information, and then contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM and report the error and detail code. For details on the `DLMgetras` utility, see 7.2.
KAPL05039-W	Received the IOCTL error from Disk Driver. errno = *aa...aa* (*bb...bb*:*cc...cc*)	Details The HDLM filter driver received an IOCTL error from a lower sd or ssd driver. *aa...aa*: Error code (decimal number) *bb...bb*: IOCTL code (hexadecimal number) *cc...cc*: Minor number of the sd or ssd driver (hexadecimal number) Action Check whether the path is correctly connected to the disk.
KAPL05040-E	The disk geometry or partition information could not be set. IOCTL = *aa...aa*, ErrorCode = *bb...bb*, PathID = *cc...cc*	Details The HDLM filter driver sent a request to the lower sd or ssd driver for IOCTL for the disk geometry or the partition information settings, but an error was returned. *aa...aa*: IOCTL code (hexadecimal number) *bb...bb*: Error code (decimal number) *cc...cc*: Path ID (decimal number) Action Remove the error from the paths, and then place the paths in the Online status. If the error occurred for all of the paths, retry the operation.

KAPL05041-E	The contents of the configuration definition file (dlmfdrv.conf) are invalid. (Instance = *aa...aa*) Execute the dlmsetconf command, and then reboot. (*bb...bb*:*cc...cc*:*dd...dd*)	Details *A definition in the configuration definition file for the HDLM filter driver is invalid.* *aa...aa*:*HDLM filter driver instance number* *bb...bb*: *Code indicating the error type* *cc...cc*: *Error information 1* *dd...dd*: *Error information 2* Action Make sure the path is connected to the disk correctly. Execute the `dlmsetconf` utility for creating the HDLM driver configuration definition file, and then reboot.
KAPL05411-W	The instance number *aa...aa* for "dlmfdrv" already exists. The HDLM device could not be attached.	Details The instance number of the HDLM filter driver instance that you tried to attach is already being used. *aa...aa*: Instance number (decimal number) Action Contact the HDLM vendor, or the maintenance company if you have a maintenance contract for HDLM.
KAPL05412-W	An attempt to allocate a bus address for the HDLM device instance *aa...aa* failed. The HDLM device could not be attached.	Details An attempt to assign a physical path name (device address) to the HDLM filter driver instance that you tried to attach has failed. *aa...aa*: Instance number (decimal number) Action Execute the `dlmsetconf` utility for creating the HDLM driver configuration definition file, and then reboot.

KAPL05437-E	The contents of the configuration definition file *aa...aa*.conf are invalid. Execute the dlmsetconf command and reboot.	Details An entry in a configuration definition file (`/kernel/drv/dlmfdrv.conf` or `/kernel/drv/dlmndrv.conf`) of the filter driver or nexus driver is invalid. *aa...aa*: Driver name Action Check whether the path is correctly connected to the disk. Execute the `dlmsetconf` utility for creating the HDLM driver configuration definition file, and then reboot the host.
KAPL05438-E	The contents of the configuration definition file *aa...aa*.conf are invalid. Instance = *bb...bb*. Execute the dlmsetconf command and reboot.	Details An entry in a configuration definition file (`/kernel/drv/dlmfdrv.conf` or `/kernel/drv/dlmndrv.conf`) of the filter driver or nexus driver is invalid. *aa...aa*: Driver name *bb...bb*: Instance number (decimal number) Action Check whether the path is correctly connected to the disk. Execute the `dlmsetconf` utility for creating the HDLM driver configuration definition file, and then reboot the host.
KAPL05901-E	Failed to open *aa...aa* directory. errno = *bb...bb*	Details An attempt by `dlmsetdevname` to open a directory has failed. *aa...aa*: Directory name *bb...bb*: Error code Action Remove the error by referencing errno.
KAPL05902-E	Failed to allocate memory. errno = *aa...aa*	Details An attempt by `dlmsetdevname` to assign memory has failed. *aa...aa*: Error code (decimal number) Action Check the memory capacity which can be used.

KAPL05903-E	Failed to search *aa...aa* directory.	Details An attempt by `dlmsetdevname` to search for a directory has failed. *aa...aa*: Directory name Action Check the permission for the directory.
KAPL05904-E	Failed to set device name.	Details An attempt by the `dlmsetconf` utility to report the logical device file name of the HDLM device to the filter driver has failed. Action Check whether the HDLM driver has started normally.
KAPL05905-W	No authority to execute the dlmsetdevname command.	Details You do not have a permission required to execute `dlmsetdevname`. Action Do not execute the command unless it is in a boot script.
KAPL05921-E	There is no symbolic link file for the HDLM nexus driver (dlmndrv@*aa...aa*) in the *bb...bb* directory.	Details LinkGen has searched the *bb...bb* directory, but there is no symbolic link file for the HDLM nexus driver (dlmndrv@*aa...aa*) in the *bb...bb* directory. *bb...bb*: Directory name *aa...aa*: Instance number of HDLM nexus driver (decimal number) Action Execute the `dlmsetconf` utility for creating the HDLM driver configuration definition file, and then reboot the host.

8.6 HDLM Alert Driver Messages

KAPL06003-I	Initialization of the HDLM alert driver (*aa...aa*) was successful.	Details *aa...aa*: Address of the table for managing the alert driver (hexadecimal number) Action None.
KAPL06004-E	Could not allocate memory. (*aa...aa*:*bb...bb*)	Details An attempt to reserve memory to save alert information has failed. *aa...aa*: Program line (hexadecimal number) *bb...bb*: Target memory size (hexadecimal number) Action Check whether the HDLM driver has started normally. If it has not started or contains an error, execute the `DLMgetras` utility for collecting HDLM error information, and then contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM and report the error and detail code. For details on the `DLMgetras` utility, see 7.2.
KAPL06007-I	IOCTL processing (*aa...aa*) will now start.	Details An IOCTL request was received. *aa...aa*: IOCTL code (hexadecimal number) Action None.
KAPL06008-I	IOCTL processing (*aa...aa*) completed normally.	Details An IOCTL request was processed normally. *aa...aa*: IOCTL code (hexadecimal number) Action None.
KAPL06009-I	Invalid IOCTL (*aa...aa*) was received. The processing is canceled.	Details A request to the alert driver was made with an invalid IOCTL code. *aa...aa*: IOCTL code (hexadecimal number) Action None.

KAPL06010-E	Could not process the IOCTL(*aa...aa*). (*bb...bb*:*cc...cc*)	Details An IOCTL request was not processed normally. *aa...aa*: IOCTL code (hexadecimal number) *bb...bb*: IOCTL end code (hexadecimal number) *cc...cc*: 0 (fixed) Action Check the message of the HDLM command (`dlnkmgr`) or HDLM manager, and then take the appropriate action. If you do not know the appropriate action, execute the `DLMgetras` utility for collecting HDLM error information, and then contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM and report the error and detail code. For details on the `DLMgetras` utility, see 7.2.
KAPL06013-E	Could not write log information into the log buffer. (*aa...aa*:*bb...bb*)	Details Log information from the filter driver was destroyed without being written to a log buffer because the attempt to reserve memory for the log information failed. *aa...aa*: Message code (hexadecimal number) *bb...bb*: Buffer size (hexadecimal number) Action Check whether any other error occurred. The information that could not be written is discarded. Review the actual memory size when another error does not occur. When the actual memory size is insufficient, increase the actual memory size. When the actual memory size is sufficient, execute the `DLMgetras` utility for collecting HDLM error information, and then contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM. For details on the `DLMgetras` utility, see 7.2.

KAPL06014-E	Could not write emergency information into the emergency information buffer. (*aa...aa*:*bb...bb*)	Details Urgent information from the filter driver was destroyed without being written to the urgent information buffer because the attempt to reserve memory for the information failed. *aa...aa*: Message code (hexadecimal number) *bb...bb*: Buffer size (hexadecimal number) Action Check whether any other error occurred. The information that could not be written is discarded. Review the actual memory size if another error does not occur. If the actual memory size is insufficient, increase the actual memory size. If the actual memory size is sufficient, execute the `DLMgetras` utility for collecting HDLM error information, and then contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM. For details on the `DLMgetras` utility, see 7.2.

8.7 HDLM Driver (Core Logic Component)Messages

KAPL07819-I	Data for maintenance: *aa...aa bb...bb cc...cc dd...dd*.	Details This message is generated by the core logic for maintenance. *aa...aa*: Detailed information 1 (decimal number) *bb...bb*: Internal function number of the core logic (decimal number) *cc...cc*: Detailed information 2 (decimal number) *dd...dd*: Detailed information 3 (decimal number) Action None.

8.8 HDLM Management Target Messages

KAPL08019-E	The path (*aa...aa*) detected an error (*bb...bb*). (*cc...cc*)	Details An error occurred in the path because of a condition such as a disconnection. *aa...aa*: Path identifier (hexadecimal number) *bb...bb*: Error code (hexadecimal number) ▪ When the path error was detected by a path health checking or the `online` operation: Displays 0x000F0000(Fixed). ▪ When the path error was detected through an I/O error: Displays the OS error code. *cc...cc*: Fixed at 0x00000000 Action Check the path in which the error was detected.

KAPL08022-E	A path error occurred. ErrorCode = *aa...aa*, PathID = *bb...bb*, PathName = *cc...cc.dd...dd.ee...ee.ff...ff*, DNum = *gg...gg*, HDevName = *hh...hh*	Details A physical or logical error occurred in the path. *aa...aa*: Error code (hexadecimal number) ▪ When the path error was detected by a path health checking or the `online` operation: Displays 0x000F0000(Fixed). ▪ When the path error was detected through an I/O error: Displays the OS error code. *bb...bb*: Path ID (same as PathID of `view -path`) (decimal number) *cc...cc*: Host Port number(same as PathName of `view -path`) (hexadecimal number) *dd...dd*: Bus number (hexadecimal number) *ee...ee*: Target ID (hexadecimal number) *ff...ff*: HLU number (same as PathName of `view -path`) (hexadecimal number)*gg...gg*: Dev number (same as DNum of `view -path`) (decimal number) *hh...hh*: Host device name Action There could be an error in the path. See 5.3 and restore the path displayed in the message to running status.
KAPL08023-I	A path was recovered. PathID = *aa...aa*, PathName = *bb...bb.cc...cc.dd...dd.ee...ee*, DNum = *ff...ff*, HDevName = *gg...gg*	Details *aa...aa*: Path ID (same as PathID of `view -path`) (decimal number) *bb...bb*: Host Port number(same as PathName of `view -path`) (hexadecimal number) *cc...cc*: Bus number(same as PathName of `view -path`) (hexadecimal number) *dd...dd*: Target ID(same as PathName of `view -path`) (hexadecimal number) *ee...ee*: HLU number (same as PathName of `view -path`) (hexadecimal number) *ff...ff*: Device number (same as DNum of `view -path`) *gg...gg*: Host Dev name (same as HDevName of `view -path`) Action None.

KAPL08026-E	An error occurred on all the paths of the LU. PathID = *aa...aa*	Details An error occurred in the last path of one LU because of a condition such as a disconnection. *aa...aa*: Path ID (same as PathID of `view -path`) (decimal number) Action Errors are detected in all the paths connected to the LUs. See 5.3 to make the path shown in the error message or the paths connected to the target LU.
KAPL08027-E	A path was excluded from the items subject to automatic failback. PathID = *aa...aa*	Details A path was excluded from the items subject to automatic failback because the system judged that an intermittent error was occurring in that path. *aa...aa*: Path ID (same as PathID of `view -path`) (decimal number) Action An intermittent error was occurring. The path may have a problem. For details on the action to be taken, See 5.3 and switch the path shown in the message into Online.

8.9 HDLM Installation Program Messages

KAPL09002-E	The disk does not have sufficient free space.	Details The file system of the installation does not have sufficient free space to install HDLM. Action Delete unnecessary files to increase the amount of free space, and then retry.
KAPL09008-W	The license code is invalid.	Action Check the license code, and then re-enter it.
KAPL09009-E	The license code is invalid. The HDLM installation program will now terminate.	Details The HDLM installation program is terminating because an attempt to enter the license code failed multiple times. Action Check the license code, and then re-execute the installation program.

KAPL09017-E	Prerequisite patch is not applied. The sufficient OS patches for the Dynamic Link Manager have not been applied yet. Reinstall the software after those patches, which are described in the Software Description, are applied to the Operating System.	Action Re-execute the installation program after applying the prerequisite patch described in 3.1.1.1.
KAPL09019-E	An attempt to cancel the registration of the bundle PP name of Hitachi Network Objectplaza Trace Library 2 failed.	Action Manually cancel the registration of the bundle PP name and uninstall Hitachi Network Objectplaza Trace Library 2. If the attempt to cancel the registration of the bundle PP name and to uninstall Hitachi Network Objectplaza Trace Library 2 fails again, contact your HDLM vendor or the maintenance company if there is a maintenance contact of HDLM.
KAPL09020-E	An attempt to uninstall Hitachi Network Objectplaza Trace Library 2 failed.	Details An attempt to uninstall HNTRLib2 has failed. Action Manually uninstall Hitachi Network Objectplaza Trace Library 2. If the attempt to uninstall Hitachi Network Objectplaza Trace Library 2 fails again, contact your HDLM vendor or the maintenance company if there is a maintenance contact of HDLM.
KAPL09021-E	An attempt to register the bundle PP name of Hitachi Network Objectplaza Trace Library 2 failed.	Details An attempt to register the PP name of HNTRLib 2 has failed. Action Contact your HDLM vendor or the maintenance company if there is a maintenance contact of HDLM.
KAPL09023-E	A file or directory related to HDLM could not be found.	Details A target file to copy to the directory of HiCommand products other than HDLM could not be found among the files related to HDLM Action Re-execute the HDLM installation program.

KAPL09024-E	An attempt to copy a file or directory related to HDLM has failed.	Details An attempt to copy a file related to HDLM to the directory of HiCommand products other than HDLM Agent has failed. Action If you detect this message while executing the HDLM installation program, re-execute the HDLM installation program. If you detect this message while executing the HiCommand products other than HDLM installation program, re-execute the HiCommand products other than HDLM installation program.

KAPL09025-W	An attempt to delete a file or directory has failed.	Details
		An attempt to delete a file related to HDLM from the directory of HiCommand products other than HDLM has failed.
		Action
		If the following files or directories are on the host, remove them manually.
		`/opt/HDVM/agent/docroot/webstart/HDLM.jnlp`
		`/opt/HDVM/agent/classes/com/Hitachi/soft/HiCommand/DVM/agent/module/HDLMManager.class`
		`/opt/HDVM/agent/docroot/webstart/hdlm`
		`/opt/HDVM/agent/docroot/hdlmhelp`
		`/opt/HDVM/agent/classes/jp`
		`/opt/HDVM/agent/classes/com/Hitachi/soft/HiCommand/DVM/agent/module/hdlm`

KAPL09026-I	Hitachi Network Objectplaza Trace Library 2 wasn't uninstalled because it was being used for other products.	Details Hitachi Network Objectplaza Trace Library 2 was not uninstalled because it was being used for other products. Action None.
KAPL09027-E	This version of the OS is not supported.	Action See 3.1.1.1 and install HDLM in a supported OS system.
KAPL09029-E	This version of HDLM cannot be updated by installation. Uninstall the already installed version of HDLM.	Action Uninstall the already installed version of HDLM.

KAPL09030-E	The installed version of HDLM is not supported for the present OS version. Uninstall the already installed version of HDLM.	Action Uninstall the already installed version of HDLM.
KAPL09031-E	A version of HDLM with a different language environment is already installed. Uninstall the already installed version of HDLM.	Action Uninstall the already installed version of HDLM.
KAPL09032-I	Update installation completed successfully. Reboot the machine to enable the new HDLM modules.	Action Reboot the machine.
KAPL09033-E	HDLM supports only SPARC platform systems.	Action See 3.1.1.1 and install HDLM in a supported OS system.
KAPL09076-I	The permanent license was installed.	Details The permanent license was installed. Action None.
KAPL09077-I	The temporary license was installed. The license expires on *aa...aa*.	Details The temporary license was installed. *aa...aa*: Year (4 numeric characters)/Month (01-12)/Day (01-31) Action Install the permanent license by the expiration day.
KAPL09078-I	The emergency license was installed. The license expires on *aa...aa*.	Details The emergency license was installed. *aa...aa*: Year (4 numeric characters)/Month (01-12)/Day (01-31) Action Install the permanent license by the expiration day.
KAPL09079-I	The permanent license has been installed.	Action None.
KAPL09080-I	The temporary license has been installed. The license expires on *aa...aa*.	Details The temporary license has been installed. *aa...aa*: Year (4 numeric characters)/Month (01-12)/Day (01-31) Action Install the permanent license by the expiration day.

KAPL09081-I	The emergency license has been installed. The license expires on *aa...aa*.	Details The emergency license has been installed. *aa...aa*: Year (4 numeric characters)/Month (01-12)/Day (01-31) Action Install the permanent license by the expiration day.
KAPL09082-W	The temporary license expired.	Action Enter the permanent license key.
KAPL09083-W	The emergency license expired.	Action Enter the permanent license key.
KAPL09086-W	The entered license key is invalid.	Action Enter a valid license key.
KAPL09087-E	The entered license key is invalid. Renewal of the license key will now stop.	Details The renewal of the license key will be aborted because an invalid license key was entered three times. Action Obtain a valid license key, and then retry.
KAPL09088-E	The entered license key is invalid. The HDLM installation will now terminate.	Action Retry installation.
KAPL09089-W	License information cannot be acquired.	Details You need to install a license because the license information cannot be acquired. Action Case of the license key file being placed in the designated directory (`/var/tmp/hdlm_license`): In this case, a license is installed from the license key file. No special operations are needed. For cases other than above: When the message that prompts you to enter the license key is displayed, enter the license key. Alternatively, cancel the installation, save the correct license key file in the designated directory, and then re-execute the HDLM installation program. `/var/tmp/hdlm_license`

KAPL09090-W	This operation will now be continued without updating the license.	Details This operation will be continued without updating the license. Action Install the permanent license later.
KAPL09091-E	A fatal error occurred in HDLM. The system environment is invalid.	Details A part of the HDLM configuration file is missing. Action Re-install HDLM.
KAPL09112-E	The license key file is invalid. File name = *aa...aa*	Details The format of the license key file is invalid. *aa...aa*: `/var/tmp/hdlm_license` Action Save the correct license key file in the designated directory, and then re-execute. `/var/tmp/hdlm_license`
KAPL09113-E	There is no installable license key in the license key file. File name = *aa...aa*	Details There is no HDLM-installable license key in the license key file. *aa...aa*: `/var/tmp/hdlm_license` Action In the license key file, store a license key that is available to install HDLM, and then re-install HDLM.
KAPL09114-I	There is no license key file. File name = *aa...aa*	Details There is no license key file in the designated directory. *aa...aa*: `/var/tmp/hdlm_license` Action When the message that prompts you to enter the license key is displayed, enter the license key. Alternatively, cancel the installation, save the correct license key file in the designated directory, and then re-execute the HDLM installation program. `/var/tmp/hdlm_license`

KAPL09115-W	An attempt to delete the license key file has failed. File name = *aa...aa*	Details An attempt to delete the license key file has failed. *aa...aa*: `/var/tmp/hdlm_license` Action If a license key file exists, delete it. `/var/tmp/hdlm_license` Store a correct license key file in the designated directory, and then re-execute.
KAPL09133-E	The following patch(es) required for HDLM has not been applied: *aa...aa* Apply the OS patch(es) before installing the HDLM package. For details on patches required for HDLM, see the HDLM document.	Details *aa...aa*: the number of the patch not yet applied Action Check the prerequisite patches described in 3.1.1.1. Apply the patch whose number is indicated, and then install HDLM.
KAPL09135-E	One of the following was executed at the same time as an HDLM command set -lic operation: another set -lic operation, or an update of the license for an update installation.	Action Check the license by using the HDLM command's `view -sys -lic` operation. Then, if necessary, update the license by using the `set -lic` operation during or after installation. If the same error message is output, contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM. Do not perform the following operation: Execution of the HDLM command's `set -lic` operation simultaneously with an update of the license for an upgrade or re-installation.
KAPL09137-W	An internal error occurred when the system information was set.	Details During the installation of HDLM, an error considered not to be the result of a user operation occurred. Action Please perform the utility for collecting HDLM error information (`DLMgetras`), and acquire obstacle information, and then contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM and report the error and detail code.

KAPL09140-E	HDLM cannot be installed in this zone. Install it in the global zone.	Details HDLM cannot be installed in a non-global zone. Install HDLM in a global zone. Action Install HDLM in a global zone.
KAPL09142-E	HDLM *aa...aa* cannot be performed. Wait a while, and then perform *aa...aa* again. Error Code = *bb...bb*	Details HDLM cannot be installed or uninstalled. *aa...aa* : "installation" or "uninstallation" *bb...bb* : Internal code (decimal number) Action Wait a while, and then reperform the installation or uninstallation. If this message is still displayed, execute the `hbsasrv stop -f` command, and then reperform the installation or uninstallation.
KAPL09143-E	HDLM *aa...aa* cannot be performed. Wait a while, and then perform *aa...aa* again. Error Code = *bb...bb*	Details HDLM cannot be installed or uninstalled. *aa...aa* : "installation" or "uninstallation" *bb...bb* : Internal code (decimal number) Action

KAPL09148-W	Could not modify file. File name = *aa...aa*, Details = *bb...bb*	Details An HDLM driver configuration definition file or HDLM driver unconfiguration definition file, *aa...aa*, could not be updated. *aa...aa*: One of the following file names: ▪ `/kernel/drv/dlmfdrv.conf` ▪ `/opt/DynamicLinkManager/config/dlmfdrv.unconf` ▪ `/kernel/drv/dlmfdrv.conf.bak` ▪ `/opt/DynamicLinkManager/config/dlmfdrv.unconf.bak` *bb...bb*: Internal code (decimal number) Action 1. Execute the `DLMgetras` utility for collecting HDLM error information to obtain the error information. Contact your HDLM vendor or your maintenance company if you have a maintenance contract for HDLM. For details on the `DLMgetras` utility, see 7.2. If *aa...aa* is `/kernel/drv/dlmfdrv.conf.bak`, or `/opt/DynamicLinkManager/config/dlmfdrv.unconf.bak`, use the file as it is because the file will not affect the system. If *aa...aa* is `/kernel/drv/dlmfdrv.conf`, or `/opt/DynamicLinkManager/config/dlmfdrv.unconf`, perform the following steps. Execute the `dlmsetconf` utility for creating the HDLM driver configuration definition file to update the file. For details on the `dlmsetconf` utility, see 7.5. Execute either of the following commands to restart the host: `# reboot -- -r` `boot -r` (at the `ok` prompt)

KAPL09149-E	An attempt to update the file has failed. File name = *aa...aa*, Details = *bb...bb*	Details An HDLM driver configuration definition file or HDLM driver unconfiguration definition file, *aa...aa*, could not be updated. *aa...aa*: One of the following file names: ▪ `/kernel/drv/dlmfdrv.conf` ▪ `/opt/DynamicLinkManager/config/dlmfdrv.unconf` ▪ `/kernel/drv/dlmfdrv.conf.bak` ▪ `/opt/DynamicLinkManager/config/dlmfdrv.unconf.bak` *bb...bb*: Internal code (decimal number) Action Execute the `DLMgetras` utility for collecting HDLM error information to obtain the error information. Contact your HDLM vendor or your maintenance company if you have a maintenance contract for HDLM. For details on the `DLMgetras` utility, see 7.2.
KAPL09176-I	The installation completed successfully. Execute the dlmsetconf utility and, if you are using VxVM, execute the dlmvxexclude utility. After that, specifying the -r (reconfigure) option, restart the host machine.	Action Create HDLM driver configuration definition file(`/kernel/drv/dlmfdrv.conf`) by executing `dlmsetconf` utility. When you want to specify that a device is not to be managed by HDLM, edit the HDLM driver unconfiguration definition file (`/opt/DynamicLinkManager/config/dlmfdrv.unconf`) and re-execute dlmsetconf utility with `-u` option. When you use VxVM, remove HDLM devices from VxVM. The `dlmvxexclude` utility assists this. Restart the host with the reconfiguration option specified.
KAPL09501-E	HDLM is not installed on this system.	Details The Service Pack cannot be applied because HDLM is not installed on this system. Action Check whether HDLM has been correctly installed.

KAPL09504-E	The language environments of HDLM and the Service Pack are different.	Details The Japanese SP was applied to the English edition of HDLM, or vice versa. Action Acquire the SP that has the same language as the installed HDLM, and then try again.
KAPL09505-E	*aa...aa* cannot be applied to the installed *bb...bb*.	Details This HDLM or Service Pack (*aa...aa*) cannot be used for an upgrade or re-installation. installation on the already installed HDLM or installed Service Pack (*bb...bb*). *aa...aa*: the-HDLM-to-be-installed HDLM-version or the-ServicePack-to-be-installed HDLM-version *bb...bb*: the-installed-HDLM HDLM-version or the-installed-ServicePack HDLM-version Action An upgrade or re-installation cannot be performed on the already installed HDLM or Service Pack.When installing a HDLM:Uninstall the installed HDLM or Service Pack, and then perform installation.When installing a Service Pack: Obtain, and then install, a Service Pack or corrected version that can be applied to the installed HDLM.
KAPL09509-E	Service Pack *aa...aa* cannot be installed. The same version has already been installed.	Details The same version as that of service pack to install is installed. Installation of Service Pack is stopped. *aa...aa*: version of the Service Pack to install Action You do not have to install the Service Pack. Keep using HDLM already installed.

KAPL09510-E	Service Pack *aa...aa* cannot be installed. A newer *bb...bb* version has already been installed.	Details A newer version of the Service pack (*bb...bb)* is already installed. Installation of Service Pack (*aa...aa*) is stopped. *aa...aa*: the version of the Service Pack to install *bb...bb*: the version of the installed HDLM or the version of the installed Service Pack Action You do not have to install the Service Pack. Keep using HDLM already installed.

8.10 Messages from the DLMgetras Utility for Collecting HDLM Error Information

KAPL10001-W	No parameter has been specified.	Details No parameter (directory to which collected information is output) has been specified. Action Check the parameters of the `DLMgetras` utility for collecting HDLM error information, and then retry. For details on the `DLMgetras` utility, see 7.2.
KAPL10002-W	Too many parameters have been specified.	Details Four or more parameters have been specified. Action Check the parameters of the `DLMgetras` utility for collecting HDLM error information, and then retry. For details on the `DLMgetras` utility, see 7.2.
KAPL10003-W	The first parameter has not been set to a directory. Value = *aa...aa*	Details The first parameter must be set to a directory to which collected information is output. *aa...aa*: First parameter Action Check the parameters of the `DLMgetras` utility for collecting HDLM error information, and then retry. For details on the `DLMgetras` utility, see 7.2.

KAPL10004-W	The parameter contains an incorrect value. Value = *aa...aa*	Details The first parameter must be a directory. The second parameter must be `-f`. *aa...aa*: Invalid parameter Action Check the parameters of the `DLMgetras` utility for collecting HDLM error information, and then retry. For details on the `DLMgetras` utility, see 7.2.
KAPL10005-W	The number of parameters is insufficient.	Details The `-f` parameter exists but the file for defining the information to be collected does not exist. Action Check the parameters of the `DLMgetras` utility for collecting HDLM error information, and then retry. For details on the `DLMgetras` utility, see 7.2.
KAPL10006-W	The file for defining the information to be collected does not exist, or cannot be read. Value = *aa...aa*	Details The file for defining the information to be collected does not exist, or the specified file exists but the permission to read the file is missing. *aa...aa*: Name of the file for defining the information to be collected Action Check whether the specified file for defining the information to be collected exists, and check whether you have access permission for the specified file.
KAPL10007-W	A directory has been specified in the third parameter. Value = *aa...aa*	Details The `-f` parameter is specified to a directory. *aa...aa*: Third parameter Action Check the parameters of the `DLMgetras` utility for collecting HDLM error information, and then retry. For details on the `DLMgetras` utility, see 7.2.

KAPL10008-W	You lack write permission for the specified directory. Value = *aa...aa*	Details You do not have write permission for the specified directory, or the creation of a subdirectory of the specified directory failed. *aa...aa*: First parameter Action Check the following. 1. Check whether you have access permission for the specified directory. Check whether the specified directory name is correct. Check that the disk has sufficient free space.
KAPL10009-W	The specified directory already exists. Do you want to overwrite it? [y/n]:	Details The specified directory already exists. Enter `y` to overwrite it or `n` to cancel. Action The specified directory already exists. Enter `y` to overwrite the existing file. Enter `n` or press any other key to terminate the `DLMgetras` utility for collecting HDLM error information without executing it. For details on the `DLMgetras` utility, see 7.2.
KAPL10010-W	A root directory has been specified. Line = *aa...aa*	Details The root "/" has been specified as a directory to be collected in the file for defining the information to be collected. *aa...aa*: Line number of the file for defining information to be collected (decimal number) Action Delete the coding of the root directory from the specified file. The displayed directory will be ignored and the `DLMgetras` utility for collecting HDLM error information will continue. For details on the `DLMgetras` utility, see 7.2.

KAPL10011-W	More than one file or directory has been specified on one line. Line = *aa...aa,* Value = *bb...bb*	Details Two or more file names or directory names exist in the file for defining the information to be collected. *aa...aa*: Line number of the file for defining information to be collected (decimal number) *bb...bb*: Indicated contents in a line Action After the `DLMgetras` utility for collecting HDLM error information terminates, check the contents of the file for defining the information to be collected. This file is shown in the message. If the contents of the file are incorrect, correct them and then try to collect error information again. The `DLMgetras` utility will ignore the specified file or directory and continue processing. For details on the `DLMgetras` utility, see 7.2.
KAPL10012-W	The specified file or directory does not exist. Line = *aa...aa,* Value = *bb...bb*	Details The specified file or directory does not exist in the file for defining the information to be collected. *aa...aa*: Line number of the file for defining information to be collected (decimal number) *bb...bb*: Indicated contents in a line Action After the `DLMgetras` utility for collecting HDLM error information terminates, check the contents of the file for defining the information to be collected. This file is shown in the message. If the contents of the file are incorrect, correct them and then try to collect error information again. The `DLMgetras` utility will ignore the specified file or directory and continue processing. For details on the `DLMgetras` utility, see 7.2.

KAPL10013-W	You lack read permission for the specified file. Line = *aa...aa,* Value = *bb...bb*	Details You lack read permission for the specified file in the file for defining information to be collected. *aa...aa*: Line number of the file for defining information to be collected (decimal number) *bb...bb*: Indicated contents in a line Action After the `DLMgetras` utility for collecting HDLM error information terminates, check the contents of the file for defining the information to be collected. This file is shown in the message. If the contents of the file are incorrect, correct them and then try to collect error information again. The `DLMgetras` will ignore the specified file and continue processing. For details on the `DLMgetras` utility, see 7.2.
KAPL10014-W	You lack read permission for the specified directory. Line = *aa...aa,* Value = *bb...bb*	Details You lack read permission for the specified directory in the file for defining information to be collected. *aa...aa*: Line number of the file for defining information to be collected (decimal number) *bb...bb*: Indicated contents in a line Action After the `DLMgetras` utility for collecting HDLM error information terminates, check the contents of the file for defining the information to be collected. This file is shown in the message. If the contents of the file are incorrect, correct them and then try to collect error information again. The `DLMgetras` will ignore the specified file and continue processing. For details on the `DLMgetras` utility, see 7.2.
KAPL10015-W	The file format is invalid. Value = *aa...aa*	Details The file format in the file for defining information to be collected is not a text file. *aa...aa*: Third parameter Action After the `DLMgetras` utility for collecting HDLM error information terminates, check whether the file for defining the information to be collected is a text file. The file is shown in the message. For details on the `DLMgetras` utility, see 7.2.

KAPL10016-W	The root directory has been specified in the first parameter.	Details A root "/" cannot be specified in a directory to which collected information is output. Action Check the parameters of the `DLMgetras` utility for collecting HDLM error information, and then re-execute. For details on the `DLMgetras` utility, see 7.2.
KAPL10017-W	You lack privileges for executing the utility for collecting HDLM error information.	Details The `DLMgetras` utility for collecting HDLM error information must be executed by a user with root privileges. Action Re-execute as a user with root .privileges. For details on the `DLMgetras` utility, see 7.2.
KAPL10020-I	The file has been obtained successfully. File = *aa...aa*, Collection time = *bb...bb*(GMT:*bb...bb*)	Details The file to be collected has been obtained. *aa...aa*: Collected file name *bb...bb*: Year (4 numeric characters)/Month (01-12)/Day (01-31) hour:minute:second Action None.
KAPL10021-I	Processing terminated before completion because a signal was received.	Details The process has been terminated by an operation such as Ctrl+C. Action The utility for collecting HDLM error information terminated before completion. If the directory is unnecessary, delete directory. For details on the `DLMgetras` utility, see 7.2.
KAPL10022-I	The utility for collecting HDLM error information completed normally.	Details Error information has been collected. Action None. For details on the `DLMgetras` utility, see 7.2.
KAPL10030-I	A user terminated the utility for collecting HDLM error information.	Details Processing of the `DLMgetras` utility has been terminated because `n` was sent as a confirmation reply. Action None. For details on the `DLMgetras` utility, see 7.2.

KAPL10031-W	The entered value is invalid. Continue operation ? [y/n]:	Details A value other than y or n has been entered for a [y/n] request. Enter `y` or `n`. Action Enter `y` or `n`.
KAPL10032-W	The entered value is invalid. The utility for collecting HDLM error information stops.	Details Processing of the `DLMgetras` utility for collecting HDLM error information will terminate because an invalid response was sent three times to a request. Action Re-execute the `DLMgetras` utility. For details on the `DLMgetras` utility, see 7.2.
KAPL10033-W	The file does not exist. Filename = *aa...aa*	Details The file to collect does not exist. *aa...aa:* File to collect Action None.
KAPL10034-E	The file could not be copied. Filename = *aa...aa*, Details = *bb...bb*	Details Execution of the cq command failed. *aa...aa*: File name you tried to copy *bb...bb*: cq output message Action An error occurred during copying of the file to be collected. The user environment might have been unstable. Check the system configuration.
KAPL10035-E	An attempt to archive the error information failed. Details = *aa...aa*	Details Execution of the tar command failed. *aa...aa*: tar output message Action See the details in the message, and then remove the cause of the error. For information about the error, collect the archive in the output directory specified at the time of execution, and then contact your HDLM vendor or your maintenance company if you have a maintenance contract for HDLM.

KAPL10036-E	An attempt to compress the error information failed. Details = *aa...aa*	Details Execution of the compress command failed. *aa...aa*: compress output message Action See the details in the message, and then remove the cause of the error. For information about the error, collect the archive in the output directory specified at the time of execution, and then contact your HDLM vendor or the maintenance company if you have a maintenance contract for HDLM.

8.11 Messages from the dlmsetconf Utility for Creating the HDLM Driver Configuration Definition File or the dlmlisthdev Utility for assisting HDLM Transitions

KAPL10201-W	usage: dlmsetconf [-d output-directory-name] [-r|-u]	Details Displays how to use `dlmsetconf`. Action Specify appropriate parameters, and then retry.
KAPL10202-W	A backup file (*aa...aa*) could not be created. Errno = *bb...bb* (*cc...cc*)	Details An attempt to create a backup file has failed. *aa...aa*: Backup file name *bb...bb*: Errno *cc...cc*: Errno string Action Check the capacity of the root file system or the file system including the directory specified by the `-d` option. Check the memory capacity that can be used.
KAPL10204-I	dlmsetconf completed successfully.	Details The `dlmsetconf` utility for creating the HDLM driver configuration definition file completed successfully. Action None.
KAPL10220-E	An error occurred when collecting the HDLM support disk information.	Details An attempt to acquire information about HDLM-managed devices has failed. Action Check the system configuration.

KAPL10221-E	An error occurred when creating the configuration definition file (*aa...aa*).	Details An attempt to create a configuration definition file has failed. *aa...aa*: Name of the configuration definition file Action Check the system configuration.
KAPL10223-E	An error occurred when releasing the memory.	Details An attempt to release memory has failed. Action Contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM if this error occurs a number of times.
KAPL10224-E	An error occurred when creating the configuration definition information.	Details An attempt to create a configuration definition file has failed. Action Check the system configuration.
KAPL10225-W	The existing configuration definition file could not be opened. Errno = *aa...aa* (*bb...bb*)	Details An attempt to open a configuration definition file has failed. *aa...aa*: Errno (decimal number) *bb...bb*: Errno string Action Remove the error by referencing Errno, and then retry.
KAPL10227-E	The text *aa...aa* could not be found in the existing configuration definition file (*bb...bb*).	Details The property name of the search string does not exist in the configuration definition file. *aa...aa*: Characters to be searched for *bb...bb*: Name of the configuration definition file Action Check the existing *bb...bb* file.
KAPL10228-W	The HDLM support disk information does not exist.	Details HDLM-managed devices cannot be found. Action Check the system configuration.

KAPL10229-E	Internal processing error.	Details An internal processing error occurred. Action Internal processing error. Contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM.
KAPL10230-E	The *aa...aa* directory could not be opened. Errno = *bb...bb* (*cc...cc*)	Details An attempt to open a directory has failed. *aa...aa*: Directory name *bb...bb*: Errno (decimal number) *cc...cc*: Errno string Action Remove the error by referencing Errno, and then retry.
KAPL10231-W	The *aa...aa* file could not be deleted. Errno = *bb...bb* (*cc...cc*)	Details An attempt to delete a file has failed. *aa...aa*: File name *bb...bb*: Errno (decimal number) *cc...cc*: Errno string Action Remove the error by referencing Errno, and then retry.
KAPL10234-W	Get file(*aa...aa*) status failed. Function = stat, Errno = *bb...bb* (*cc...cc*)	Details An attempt to acquire information about a file has failed. *aa...aa*: File name *bb...bb*: Errno (decimal number) *cc...cc*: Errno string Action Remove the error by referencing Errno, and then retry.
KAPL10235-W	Get symbolic link file(*aa...aa*) status failed. Function = lstat, Errno = *bb...bb* (*cc...cc*)	Details An attempt to acquire information about the symbolic link file has failed. *aa...aa*: File name *bb...bb*: Errno (decimal number) *cc...cc*: Errno string Action Remove the error by referencing Errno, and then retry.

KAPL10236-E	Could not allocate memory. Errno = *aa...aa* (*bb...bb*)	Details An attempt to reserve memory has failed. *aa...aa*: Errno (decimal number) *bb...bb*: Errno string Action Check memory usage, and then restore unnecessary memory space.
KAPL10237-E	The property information (*aa...aa*) on a device (*bb...bb*) could not be collected.	Details An attempt to acquire property information for the sd or ssd device has failed. *aa...aa*: property name *bb...bb*: sd or ssd device name (physical path name) Action Check whether the sd or ssd driver is attached. If there is no problem, contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM.
KAPL10238-W	The device (*aa...aa*) could not be opened. Errno = *bb...bb* (*cc...cc*)	Details An attempt to open an sd or ssd device failed. *aa...aa*: sd or ssd device name *bb...bb*: Errno (decimal number) *cc...cc*: Errno string Action Remove the error by referencing Errno, and then retry.
KAPL10239-W	The ioctl() operation failed. FD = *aa...aa*, Errno = *bb...bb* (*cc...cc*)	Details An attempt to issue an IOCTL call has failed. *aa...aa*: File descriptor *bb...bb*: Errno (decimal number) *cc...cc*: Errno string Action Remove the error by referencing Errno, and then retry.
KAPL10240-E	The directory specified by the -d option does not exist.	Details The directory specified in the -d option was not found. Action Check whether the specified directory exists.

KAPL10241-E	An error occurred when collecting the catalog message. Errno = *aa...aa* (*bb...bb*)	Details A message could not be acquired from the catalog file. *aa...aa*: Errno (decimal number) *bb...bb*: Errno string Action Remove the error by referencing Errno, and then retry.
KAPL10242-I	To configure HDLM, device files of sd/ssd devices managed by HDLM are removed after reboot. Do you want to continue? [y/n]	Details This message is confirming execution of the `dlmsetconf` utility for creating the HDLM driver configuration definition file. To configure HDLM, the logical device files of the HDLM-managed sd or ssd device are deleted at reboot. Action After you execute the `dlmsetconf` utility, the logical device files of the HDLM-managed sd or ssd device are deleted at reboot. To configure HDLM, enter `y` if you want to execute the `dlmsetconf` utility, enter `n` to stop.
KAPL10243-W	A lock error occurred during open processing for the physical path (*aa...aa*).	Details An attempt to open a path failed because an exclusive open is being executed for another process. *aa...aa*: Physical path Action Comply with the following KAPL10244-I messages.
KAPL10244-I	The definition of the physical path was reconfigured since the same physical path existed in the configuration definition file (dlmfdrv.conf).	Details Information about the same physical path (the same as the path where the open processing failed) was already defined in the HDLM driver configuration file (`/kernel/drv/dlmfdrv.conf`). The path information was redefined on the basis of that information. Action Remove the physical path that was shown in the above KAPL10243-W messages from the configuration definition file if that physical path is not needed.

KAPL10251-W	An attempt to retrieve information for the directory *aa...aa* failed. Function = readdir, errno = *bb...bb* (*cc...cc*)	Details An attempt to collect directory information has failed. *aa...aa*: Directory name *bb...bb*: Errno (decimal number) *cc...cc*: Errno string Action See errno to eliminate the error, and then retry.
KAPL10252-W	An attempt to read the contents of symbolic link *aa...aa* failed. Function = readlink, errno = *bb...bb* (*cc...cc*)	Details An attempt to acquire link destination information about the symbolic link file has failed. *aa...aa*: File name *bb...bb*: Errno (decimal number) *cc...cc*: Errno string Action See errno to eliminate the error, and then retry.
KAPL10253-W	An attempt to create symbolic link *aa...aa* to a file failed. Function = symlink, errno = *bb...bb* (*cc...cc*)	Details An attempt to create a symbolic link file has failed. *aa...aa*: File name *bb...bb*: Errno (decimal number) *cc...cc*: Errno string Action See errno to eliminate the error, and then retry.
KAPL10254-W	The entered value is invalid. Re-enter [y/n]:	Details An incorrect value was entered when the message confirming execution of the `dlmsetconf` utility for creating the HDLM driver configuration definition file was output. Action If you want to continue the `dlmsetconf` utility processing, enter `y`. If you want to stop, enter `n`.
KAPL10255-E	The entered value is invalid. The operation stops.	Details Processing will stop because an incorrect value was entered three or more times for the message confirming the execution of the `dlmsetconf` utility for creating the HDLM driver configuration definition file. Action When executing the utility, enter the correct value.

KAPL10256-I	The user terminated the operation.	Details The user has interrupted processing. Action None.
KAPL10257-W	The output file (*aa...aa*) could not be opened. Errno = *bb...bb* (*cc...cc*)	Details An attempt to open the output file has failed. *aa...aa*: Output file name *bb...bb*: Errno (decimal number) *cc...cc*: Errno string Action Take countermeasures according to the Errno.
KAPL10297-E	The specified file does not exist. Value = *aa...aa*	Details The file that was specified at execution of the `dlmlisthdev` utility for assisting HDLM transitions does not exist. *aa...aa*: Input file name Action Check the input file, and then retry. For details on the `dlmlisthdev` utility, see7.4.
KAPL10298-E	Could not get host device information. The "dlnkmgr view -path" command failed.	Details An attempt to acquire HDevList failed because the `dlnkmgr view -path` command failed. Action Remove from the `dlnkmgr view -path` command error, and execute the command again.
KAPL10299-E	The input file is invalid. The input file is not the output result of the "dlnkmgr view -path" command.	Details The content of a file specified in a parameter of the `dlmlisthdev` command differs from the output of the `dlnkmgr view -path` command. Action Check the input file, and then retry. For details on the `dlmlisthdev` utility, see 7.4.

8.12 Messages from the dlmcfgmgr Utility for Managing the HDLM Configuration

KAPL10316-E	Could not allocate memory. Size = *aa...aa*	Details Securing of memory failed. *aa...aa*: Target memory size Action Contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM.
KAPL10320-E	The file format is invalid. File name = *aa...aa*	Details The file format is invalid. *aa...aa*: File name Action Re-execute the HDLM-configuration definition utility. If the same message is output, contact your HDLM vendor or maintenance company if you have a maintenance contract for HDLM. For details on the `dlmcfgmgr` utility, see 7.3.
KAPL10328-E	The entered value is invalid. The HDLM-configuration definition utility processing will now stop.	Details Processing will stop because an incorrect value was entered three or more times for the message confirming execution of the `dlmcfgmgr` utility. Action Re-execute the `dlmcfgmgr` utility. For details on the `dlmcfgmgr` utility, see 7.3.
KAPL10329-E	A file could not be opened. File name = *aa...aa*, Errno = *bb...bb*	Details The file could not be opened. *aa...aa*: File name *bb...bb*: Errno (decimal number) Action Confirm the existence of the file and access permission. If there is no problem, contact your HDLM vendor or maintenance company if you have a maintenance contract for HDLM.

KAPL10334-E	Ioctl *aa...aa* failed with errno = *bb...bb*.	Details An attempt to issue an IOCTL call has failed. *aa...aa*: Ioctl Code *bb...bb*: Errno (decimal number) Action Contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM.
KAPL10368-I	usage: dlmcfgmgr {-a|-h}	Details Displays how to use `dlmcfgmgr`. Action Specify appropriate parameters, and then retry. For details on the `dlmcfgmgr` utility, see 7.3.
KAPL10369-I	dlmcfgmgr *aa...aa* started: *bb...bb*	Details Shows the execution option and the execution start time of the `dlmcfgmgr` utility. *aa...aa*: Parameter *bb...bb*: Date (character string) Action None. For details on the `dlmcfgmgr` utility, see 7.3.
KAPL10370-I	dlmcfgmgr completed successfully.	Details The `dlmcfgmgr` utility for managing the HDLM configuration completed successfully. Action None. For details on the `dlmcfgmgr` utility, see 7.3.
KAPL10371-W	The HDLM support disk information does not exist in the configuration file(dlmfdrv.conf).	Details The `/kernel/drv/dlmfdrv.conf` file does not contain the data of the HDLM-managed devices. Action Check the system configuration, and then retry.
KAPL10372-E	An error occurred when collecting the catalog message. ErrorCode = *aa...aa*(*bb...bb*)	Details A message could not be acquired from the message catalog. *aa...aa*: Errno *bb...bb*: Errno string Action Remove the error by referencing ErrorCode(errno), and then retry.

KAPL10373-I	This operation will change the configuration of HDLM devices, and delete device files of sd/ssd devices managed by HDLM.Do you want to continue? [y/n]:	Details This message confirms the execution of the `dlmcfgmgr` utility for managing the HDLM configuration. Action If you want to continue the `dlmcfgmgr` utility processing, enter `y`. If you want to stop, enter `n`. For details on the `dlmcfgmgr` utility, see 7.3.
KAPL10374-I	The user terminated the operation.	Details The execution of the `dlmcfgmgr` utility for managing the HDLM configuration was stopped. Action None. For details on the `dlmcfgmgr` utility, see 7.3.
KAPL10375-W	The entered value is invalid. Please re-enter it [y/n]:	Details An incorrect value was entered for the message confirming the execution of the `dlmcfgmgr` utility for managing the HDLM configuration. Action Enter y or n. For details on the `dlmcfgmgr` utility, see 7.3.
KAPL10376-E	The device (*aa...aa*) could not be opened. ErrorCode = *bb...bb* (*cc...cc*)	Details An attempt to open the device managing the HDLM filter driver has failed. *aa...aa*: Name of the device managing the HDLM filter driver *bb...bb*: Errno (decimal number) *cc...cc*: Errno string Action Take action in accordance with ErrorCode (errno), and then retry.
KAPL10377-I	The HDLM device configurations have not been changed.	Details There was no change in the HDLM device configuration. Action None.
KAPL10378-I	The LU (LDEV = *aa...aa*) was successfully added.	Details *aa...aa*: LDEV name Action None.

KAPL10379-E	An operation (*aa...aa*) of the lock file (*bb...bb*) has failed. ErrorCode = *cc...cc* (*dd...dd*)	Details An attempt to acquire a lock for simultaneous execution has failed. *aa...aa*: Command parameters (character strings like `F_SETLK` or `F_GETLK`) *bb...bb*: Filename *cc...cc*: Errno (decimal number) *dd...dd*: Errno string Action Contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM.
KAPL10380-E	A parameter is invalid.	Details An execution option of the `dlmcfgmgr` utility for managing the HDLM configuration is incorrect. Action Specify appropriate parameters, and then retry. For details on the `dlmcfgmgr` utility, see 7.3.
KAPL10382-W	There is no instance of the HDLM nexus driver that manages the path of the LU. This LU (LDEV = *aa...aa*) is not added. FilterPath = *bb...bb*	Details This LU was not added because the HDLM nexus driver instance does not exist in the LU. *aa...aa*: LDEV name *bb...bb*: Filter Path Action If you want to use the LU, reboot the host.
KAPL10383-E	The text *aa...aa* could not be found in the existing configuration definition file (*bb...bb*).	Details The property name of the search string does not exist in the `/kernel/drv/dlmfdrv.conf` file. *aa...aa*: Search string *bb...bb*: Configuration definition file name Action Contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM.

KAPL10384-E	The value of the *aa...aa* property in the configuration definition file (*bb...bb*) is invalid.	Details The value of the property defined in the `/kernel/drv/dlmfdrv.conf` file is invalid. *aa...aa*: Property name *bb...bb*: Configuration definition file name Action Contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM.
KAPL10385-I	The LU (LDEV = *aa...aa*) is not configured.	Details Unconf is specified for this LU in the HDLM driver configuration file (`/kernel/drv/dlmfdrv.conf`). *aa...aa*: LDEV name Action None.
KAPL10386-I	The LU (LDEV = *aa...aa*) is already managed by HDLM.	Details This LU is already managed in dlmfdrv. *aa...aa*: LDEV name Action If the path configuration of the LU shown by the LDEV name has changed, reboot the host. If there is no change to the configuration, do nothing.
KAPL10387-W	The path of the another LU (LDEV = *aa...aa*) managed by HDLM has the same value as the property (*bb...bb*). The configuration of the LU (LDEV = *cc...cc*) will not be changed.	Details The property value of the path described in the configuration definition file has the same value as a different LU path managed by dlmfdrv. *aa...aa*: LDEV name *bb...bb*: Property name *cc...cc*: LDEV name Action If you want to use the LU, reboot the host.

KAPL10388-E	The path (*aa...aa*) could not be added. (Function: *bb...bb*, Error Code = *cc...cc* (*dd...dd*))	Details An attempt to add a path has failed. *aa...aa*: Physical device file path for the sd or ssd device corresponding to the HDLM device that could not be added *bb...bb*: Function name *cc...cc*: Errno (decimal number) *dd...dd*: Errno string Action Contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM.
KAPL10390-W	The backup file (*aa...aa*) could not be created. ErrorCode = *bb...bb* (*cc...cc*)	Details An attempt to create a backup file has failed. *aa...aa*: Backup file name *bb...bb*: Errno (decimal number) *cc...cc*: Errno string Action Check the capacity of the root file system or the file system including the file. Check the memory capacity which can be used.
KAPL10391-E	The output file (*aa...aa*) could not be opened. ErrorCode = *bb...bb* (*cc...cc*)	Details An attempt to open the output file has failed. *aa...aa*: Output file name *bb...bb*: Errno (decimal number) *cc...cc*: Errno string Action Remove the error by referencing ErrorCode (errno).
KAPL10392-E	The catalog file (*aa...aa*) could not be opened. ErrorCode = *bb...bb* (*cc...cc*)	Details An attempt to open the catalog file has failed. *aa...aa*: Catalog file name *bb...bb*: Errno (decimal number) *cc...cc*: Errno string Action Contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM.

KAPL10395-E	Multiple paths of the LU (LDEV = *aa...aa*) to be added have the same value as the property (*bb...bb)*. The configuration of the LU will not be changed.	Details The same property value exists for an identical LU path defined in the configuration definition file. *aa...aa*: LDEV name *bb...bb*: Property name = property value Action Contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM.
KAPL10396-W	No authority to execute the dlmcfgmgr utility.	Details You do not have permission to execute the `dlmcfgmgr` utility for managing the HDLM configuration. Action Retry as a user with root privileges. For details on the `dlmcfgmgr` utility, see 7.3.
KAPL10397-E	Multiple instances of the dlmcfgmgr utility cannot be executed concurrently.	Details The `dlmcfgmgr` utility for managing the HDLM configuration cannot be simultaneously executed for multiple processes. Action Wait until the `dlmcfgmgr` utility ends, and then re-execute the utility. For details on the `dlmcfgmgr` utility, see 7.3.
KAPL10398-E	An attempt to acquire the process ID has failed. ErrorCode = *aa...aa* (*bb...bb*)	Details An attempt to acquire the process ID that is used to acquire locks for multiple simultaneous executions has failed. *aa...aa*: Errno (decimal number) *bb...bb*: Errno string Action Contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM.

8.13 Messages from the dlmvxexclude Utility for Assisting the Creation of the VxVM Configuration File

KAPL10701-I	usage: dlmvxexclude [-d dir | -h]	Details Displays how to use `dlmvxexclude`. Action Specify appropriate parameters, and then retry. For details on the `dlmvxexclude` utility, see 7.6.
KAPL10702-E	A backup file (*aa...aa*) could not be created. Errno = *bb...bb* (*cc...cc*)	Details An attempt to create a backup file has failed. *aa...aa*: Backup file name *bb...bb*: Errno (decimal number) *cc...cc*: Errno string Action Check the capacity of the root file system or the capacity of the file system containing the specified files. Check the available memory space.
KAPL10703-E	This version of the OS is not supported.	Details The `dlmvxexclude` utility for assisting creation of the VxVM configuration file was executed on a non-supporting OS. Action It is not necessary to execute this command. For details on the `dlmvxexclude` utility, see 7.6.
KAPL10704-I	dlmvxexclude completed successfully.	Details The `dlmvxexclude` utility for assisting creation of the VxVM configuration file completed successfully. Action None. For details on the `dlmvxexclude` utility, see 7.6.
KAPL10705-E	Internal processing error.	Details An internal error occurred while the `dlmvxexclude` utility for assisting creation of the VxVM configuration file was executing. Action Check the system configuration and then retry. For details on the `dlmvxexclude` utility, see 7.6.

KAPL10706-E	Invalid argument.	Details An invalid argument was specified. Action Specify appropriate parameters, and then retry.
KAPL10707-E	This version of the VxVM is not supported.	Details The `dlmvxexclude` utility for assisting creation of the VxVM configuration file was executed in a non-supporting VxVM environment. Action It is not necessary to execute this command. For details on the `dlmvxexclude` utility, see 7.6.
KAPL10708-E	No authority to execute dlmvxexclude.	Details You do not have permission to execute the `dlmvxexclude` utility for assisting creation of the VxVM configuration file. Action Execute by a user with root privileges. For details on the `dlmvxexclude` utility, see 7.6.
KAPL10709-E	An error occurred when creating the file *aa...aa*.	Details An attempt to create a file has failed. *aa...aa*: File name Action Check the system configuration.
KAPL10710-E	The existing configuration definition file *aa...aa* could not be opened. Errno = *bb...bb* (*cc...cc*)	Details An attempt to open a configuration definition file has failed. *aa...aa*: Name of the configuration definition file *bb...bb*: Errno (decimal number) *cc...cc*: Errno string Action Take countermeasures according to the Errno and then execute the command again.

KAPL10711-E	The text *aa...aa* could not be found in the existing configuration definition file (*bb...bb*).	Details The property name of the search string does not exist in the `/kernel/drv/dlmfdrv.conf` file. *aa...aa*: Search string *bb...bb*: Name of the configuration definition file Action Check the existing configuration definition file (`/kernel/drv/`*aa...aa*).
KAPL10712-E	The HDLM support disk information does not exist.	Details A valid entry does not exist in the `/kernel/drv/dlmfdrv.conf` file. Action Check the system configuration and then execute the `dlmsetconf` utility for creating the HDLM driver configuration definition file.
KAPL10713-E	The *aa...aa* directory could not be opened. Errno = *bb...bb* (*cc...cc*)	Details An attempt to open a directory has failed. *aa...aa*: Directory name *bb...bb*: Errno (decimal number) *cc...cc*: Errno string Action Check the system configuration and then retry.
KAPL10714-E	Could not allocate memory. Errno = *aa...aa* (*bb...bb*).	Details An attempt to reserve memory has failed. *aa...aa*: Errno (decimal number) *bb...bb*: Errno string Action Check if more than necessary memory has been allocated. If so, release the unnecessary resources.

KAPL10715-E	The information on a device (*aa...aa*) could not be collected.	Details An attempt to acquire information about the device managed by HDLM has failed. *aa...aa*: Physical path name Action Execute the `dlmsetconf` utility for creating the HDLM driver configuration definition file, and then re-execute the `dlmvxexclude` utility for assisting creation of the VxVM configuration file. For details on the `dlmvxexclude` utility, see 7.6.
KAPL10716-E	*aa...aa* specified by the -d option is not a directory.	Details The specified file name is not the name of a directory. *aa...aa*: Directory name Action Specify the correct directory.
KAPL10717-E	The directory specified by the -d option does not exist.	Details A non-existing directory was specified. Action Specify the correct directory.
KAPL10718-E	An error occurred when collecting the catalog message. Errno = *aa...aa* (*bb...bb*)	Details A message could not be acquired from the message catalog. *aa...aa*: Errno *bb...bb*: Errno string Action Take countermeasure according to the Errno and then execute the command again.
KAPL10719-I	The file *aa...aa* already exists. Do you want to overwrite it? [y/n]:	Details This message confirms the execution of the `dlmvxexclude` utility for assisting creation of the VxVM configuration file. *aa...aa*: File name Action Enter `y` to continue. Enter `n` to cancel. For details on the `dlmvxexclude` utility, see 7.6.

KAPL10720-W	The entered value is invalid. Re-enter [y/n]:	Details An incorrect value was entered for the message confirming execution of the `dlmvxexclude` utility for assisting creation of the VxVM configuration file. Action Enter `y` or `n`. For details on the `dlmvxexclude` utility, see 7.6.
KAPL10721-E	The entered value is invalid. The processing will now stop.	Details Processing was stopped because an incorrect value was entered three or more times for the message confirming execution of the `dlmvxexclude` utility for assisting creation of the VxVM configuration file. Action Retry `dlmvxexclude` command. For details on the `dlmvxexclude` utility, see 7.6.
KAPL10722-I	The user terminated the operation.	Details The execution of the `dlmvxexclude` utility for assisting creation of the VxVM configuration file was stopped. Action None. For details on the `dlmvxexclude` utility, see 7.6.
KAPL10723-I	The execution of dlmvxexclude is not required for the configuration.	Details The devices from VxVM do not need to be disabled so the `dlmvxexclude` utility does not need not be executed. Action It is not necessary to execute this utility. For details on the `dlmvxexclude` utility, see 7.6.

8.14 Messages from HDLM remote access interface

KAPL11901-I	*aa...aa* has started.	Details The operation has started on the host. *aa...aa*: Operation (character string) ▪ Get Path Information ▪ Get Option Information ▪ Set Option Information ▪ Clear Data ▪ Get HDLM Manager Status ▪ Get HDLM Driver Status ▪ Get HDLM Alert Driver Status Action None.
KAPL11902-I	*aa...aa* has started. PathID = *bb...bb*	Details The operation has started on the host. *aa...aa*: Operation (character string) ▪ `Online` ▪ `Offline` *bb...bb*: The Path ID of the target path for the operation.(decimal number) Action None.

KAPL11903-I	*aa...aa* has completed normally.	Details Operation has completed normally on the host. *aa...aa*: Operation (character string) ▪ Get Path Information ▪ Get Option Information ▪ Set Option Information ▪ Clear Data ▪ Get HDLM Driver Status ▪ Get HDLM Manager Status ▪ Get HDLM Alert Driver Status ▪ `Online` ▪ `Offline` ▪ Get SNMP Trap Information ▪ Set SNMP Trap Information ▪ Set LU Load Balance Action None.
KAPL11904-E	*aa...aa* has completed abnormally. Error status = *bb...bb*	Details Operation has completed abnormally on the host. *aa...aa*: Operation (character string) ▪ Get Path Information ▪ Get Option Information ▪ Set Option Information ▪ Clear Data ▪ Get HDLM Driver Status ▪ Get HDLM Manager Status ▪ Get HDLM Alert Driver Status ▪ `Online` ▪ `Offline` ▪ Get SNMP Trap Information ▪ Set SNMP Trap Information ▪ Set LU Load Balance *bb...bb*: Error-status-returned-from-API (character string) Action Execute the `DLMgetras` utility for collecting HDLM error information, and then contact your HDLM vendor or the maintenance company if you have a maintenance contract for HDLM. For details on the `DLMgetras` utility, see 7.2.

KAPL11905-E	An unexpected error occurred.	Details An exception occurred during processing in the host. Action Execute the `DLMgetras` utility for collecting HDLM error information, and then contact your HDLM vendor or the maintenance company if you have a maintenance contract for HDLM. For details on the `DLMgetras` utility, see 7.2.
KAPL11906-I	GUI information - *aa...aa*	Details This information is required for determining the cause of the problem (if any). *aa...aa*: Trace information (character string) Action None.
KAPL11907-I	XML reception - *aa...aa*	Details This information is required for determining the cause of the problem (if any). *aa...aa*: XML information (character string) Action None.
KAPL11908-I	XML transmission - *aa...aa*	Details This information is required for determining the cause of the problem (if any). *aa...aa*: XML information (character string) Action None.

8.15 Return Codes for the HDLM Remote Access Interface

The HDLM remote access interface is an interface used by other HiCommand products to access hosts. HDLM outputs a return code when the operation that the HDLM remote access interface requested to HDLM terminates abnormally, or when there is a warning while the operation terminates normally.

1002	Details There is no path on which the operation can be performed. Action Refresh the host information, check the path status, and then perform the operation again.
1003	Details No path was detected. Action Check whether a path between the host and the storage subsystem is connected. If a path is connected, check whether HDLM is configured correctly.
1004	Details Memory required for HDLM internal processing could not be allocated. Action Terminate unnecessary applications to increase free memory, or restart the host.
1006	Details The Offline path cannot be placed Online. Action Remove the error in the path, and then retry.
1007	Details The target path of the offline operation is the last path of the device and cannot be placed in the offline status. Action Click Refresh to update the host information, check the path status, and then retry the offline operation.
1015	Details The Offline path cannot be placed Online. Action Remove the error in the path, and then retry.
1016	Details The target path(s) are already Online. Action Update the host information, and then check the path status.
1017	Details The target path(s) are already Offline(C). Action Update the host information, and then check the path status.

1019	Details An error occurred during HDLM internal processing. Action Execute the DLMgetras utility for collecting HDLM error information, and then contact your HDLM vendor or the maintenance company if you have a maintenance contract for HDLM. For details on the DLMgetras utility, see 7.2.
1020	Details An unexpected error occurred during HDLM internal processing. Action Execute the DLMgetras utility for collecting HDLM error information, and then contact your HDLM vendor or the maintenance company if you have a maintenance contract for HDLM. For details on the DLMgetras utility, see 7.2.
1025	Details A parameter value is invalid. Action Refresh the host information, and then perform the operation again. If the same error occurs even after removing the reservation, execute the DLMgetras utility for collecting HDLM error information, and then contact your HDLM vendor or maintenance company if you have a maintenance contract for HDLM. For details on the DLMgetras utility, see 7.2.
1026	Details The getting Path Information has been stopped because the configuration of paths was changed during the processing of the getting Path Information. Action Refresh the host information, check the path status, and then perform the operation again.
1027	Details The error monitoring interval and the number of times that the error is to occur conflict with the automatic failback checking interval. Action Set the monitoring interval to a value that is equal to or more than (auto failback checking interval x number of times).
1028	Details An attempt to recover the path failed because an error occurred in the path. Action Resolve the error in the relevant path, and then re-execute.
1029	Details An attempt to recover the path failed because of insufficient memory. Online processing will be stopped if online processing is running for multiple paths. Action Re-execute. If the error occurs again, confirm that the host memory is sufficient.

1030	Details An attempt to recover the path failed because the LU for the path to be put online has been reserved by another path. Action Use the `format` command to confirm that the LU for the path to be put online is not reserved by another host. If the LU is incorrectly reserved, remove the reservation, and then re-execute the `online` operation for the HDLM command. If the LU is not incorrectly reserved or if the same error occurs even after removing the reservation, execute the `DLMgetras` utility for collecting HDLM error information to collect the error information, and then contact your HDLM vendor or maintenance company if you have a maintenance contract. Refresh the host information, and then perform the operation again. If the same error occurs even after removing the reservation, execute the DLMgetras utility for collecting HDLM error information, and then contact your HDLM vendor or maintenance company if you have a maintenance contract for HDLM. For details on the `DLMgetras` utility, see 7.2.
1031	Details An attempt to recover the path failed because an error occurred on the host. Action Re-execute. If the same error occurs even after removing the reservation, execute the `DLMgetras` utility for collecting HDLM error information to collect the error information, and then contact your HDLM Refresh the host information, and then perform the operation again. If the same error occurs even after removing the reservation, execute the DLMgetras utility for collecting HDLM error information, and then contact your HDLM vendor or maintenance company if you have a maintenance contract for HDLM. For details on the `DLMgetras` utility, see 7.2.
1032	Details An attempt to recover the path failed because the correct disk label has not been set in the LU for the path to be put online. Action Use the `format` command to set the VTOC disk label in the LU for the relevant path, and then re-execute.
1033	Details An attempt to acquire the HDLM version information failed. Action Re-execute. If the same error occurs even after removing the reservation, execute the `DLMgetras` utility for collecting HDLM error information to collect the error information, and then contact your HDLM vendor or maintenance company if there is a maintenance contract for HDLM. For details on the `DLMgetras` utility, see 7.2.
1034	Details An attempt to acquire information about the HDLM version or Service Pack version has failed. Action Re-execute. If the same error occurs even after removing the reservation, execute the `DLMgetras` utility for collecting HDLM error information to collect the error information, and then contact your HDLM vendor or maintenance company if there is a maintenance contract for HDLM. For details on the `DLMgetras` utility, see 7.2.

1035	Details An attempt to acquire information about the HDLM version or Service Pack version has failed. Action Re-execute. If the same error occurs even after removing the reservation, execute the DLMgetras utility for collecting HDLM error information to collect the error information, and then contact your HDLM vendor or maintenance company if there is a maintenance contract for HDLM. For details on the DLMgetras utility, see 7.2.
1036	Details An attempt to acquire information about the HDLM version or Service Pack version has failed. Action Re-execute. If the same error occurs even after removing the reservation, execute the DLMgetras utility for collecting HDLM error information to collect the error information, and then contact your HDLM vendor or maintenance company if there is a maintenance contract for HDLM. For details on the DLMgetras utility, see 7.2.
1037	Details A parameter value is invalid. Action Refresh the host information, and then perform the operation again. If the same error occurs even after removing the reservation, execute the DLMgetras utility for collecting HDLM error information, and then contact your HDLM vendor or maintenance company if you have a maintenance contract for HDLM. For details on the DLMgetras utility, see 7.2.
1041	Details An attempt to communicate with the HDLM Manager has failed. Action Check whether the HDLM Manager is running on the host. If it is not running, start the HDLM Manager.
1042	Details Information about the path configuration on the specified LU does not match the path configuration information held by HDLM. Action Refresh the host information, check the path status, and then perform the operation again.
1045	Details A parameter value is invalid. Action Refresh the host information, and then perform the operation again. If the same error occurs even after removing the reservation, execute the DLMgetras utility for collecting HDLM error information, and then contact your HDLM vendor or maintenance company if you have a maintenance contract for HDLM. For details on the DLMgetras utility, see 7.2.

1046	Details A parameter value is invalid. Action Refresh the host information, and then perform the operation again. If the same error occurs even after removing the reservation, execute the `DLMgetras` utility for collecting HDLM error information, and then contact your HDLM vendor or maintenance company if you have a maintenance contract for HDLM. For details on the `DLMgetras` utility, see 7.2.

Appendix A Differences in Functionality Between HDLM Versions

This section gives precautions on differences in functionality between HDLM versions.

A.1 Differences in Functionality Between HDLM 5.9 and Earlier Versions

- Hitachi storage subsystems can be used as a boot disk.
- HDLM supports Universal Storage Platform V.
- The following information is now output to the `getras.tar.Z` file, which contains HDLM and system information:
  - Dump device configuration
  - Host hardware-configuration information
  - Output results of `vxdisk path`
  - `metastat -s` *diskset-name* information
- Solaris-related, HDLM-related, and application-related tasks can now be performed even if the server on which HDLM is installed is started in single-user mode.
- The function for operating HDLM from HDLM GUI or HDLM Web GUI is no longer supported.
- Solaris 2.6 and Solaris 7 are no longer supported.

A.2 Precautions on Differences in Functionality Between HDLM 5.6.1 or earlier and HDLM 5.6.2 or later

For HDLM version 5.6.1 or earlier and version 5.6.2 or later, when one of the following commands is executed for LDEV created on SVS, the output LDEV display names will be different.

Table A.1 shows the commands that output different LDEV display names and the actual output display names.

Table A.1 Differences in LDEV Display Names Between Versions

HDLM version	Operations of the dlnkmgr command	Output LDEV display name
5.6.1 or earlier	▪ `dlnkmgr view -path` ▪ `dlnkmgr view -drv` ▪ `dlnkmgr view -lu`	XP
5.6.2 or later		SVS

Glossary

This glossary explains the terms used in this manual.

A

automatic failback

A function for checking the status of paths at regular intervals, and automatically placing the status of a path recovered from an error into Online status.

If a path with Offline(E) or Online(E) status recovers from an error, automatic failback places the path Online.

Automatic failback checks the status of paths that were placed in Offline(E) or Online(E) status because of an error, but does not check the status of paths that were placed in Offline(C) status by executing the `offline` operation. For details on the `offline` operation, see 6.4 *offline (Places a Path or Paths Offline)*.

AutoPATH_ID

An ID that HDLM assigns to a path at system startup. Every path has a unique AutoPATH_ID.

(See also: *path*)

B

boot disk environment" INDEXITEM="boot disk environment">boot disk environment

An environment in which the startup disk is in a storage subsystem instead of in the host.

C

CHA (Channel Adapter)

An adapter for controlling a channel of a storage subsystem.

CLPR(Cache Logical Partition)

A function, supported by the TagmaStore USP and Universal Storage Platform V, for logically splitting a cache. This function splits the cache by array groups in the storage subsystem, so that other array groups do not affect cache performance.

cluster

A system that connects multiple hosts having the same OS or platform (that is, an environment in which the same application can run) and treats them as one system.

D

Dev

A logical division of an LU that HDLM controls and operates. A Dev is equivalent to a slice or a partition in Solaris.

In Solaris, each LU can be divided into more than one Dev.

Each Dev has a Dev number.

Depending on the access method, Devs are classified into the following two types: block devices and character devices.

(See also: *Dev number*)

Dev number

In UNIX, a device number usually indicates a combination of the major device number and minor device number. In HDLM, however, a device number has the following meaning: a Dev number that is displayed in the `DNum` field.

A number (from 0 to 7 in ascending order) is assigned to each Dev in an LU. In Solaris it is called the slice number.

No corresponding Dev number is allocated, because HDLM manages Devs per LU.

(See also: *Dev*)

dynamic reconfiguration

A reconfiguration that enables HDLM to recognize an HDLM device whose configuration has been changed, without restarting the host.

E

emulation type

An LU type accessible from a host. Because an HDLM host is an open-system host such as a PC or a UNIX machine, the HDLM host can access only the LUs that have open-system emulation types.

For details on emulation types supported by a storage subsystem, see the maintenance manual for that storage subsystem.

F

failback

A function for placing the status of a path recovered from an error into Online status, and switching the access path to the recovered path.

failover

A function for switching to another normal path if there is an error in a path, thereby enabling the system to continue to operate.

FCode

A program written by the interpreter language that is not dependent on the machine. The HBA that can connect to a boot disk might incorporate FCode for use.

H

HBA (Host Bus Adapter)

A device that functions as an interface connecting a host with external devices.

In this manual, the term *HBA* refers to an interface card installed in a host, in configurations where the host and storage units are connected via a SCSI or a Fibre Channel connection.

HDLM alert driver

A program that receives information about an error detected by the HDLM driver, and then reports the error information to the HDLM manager.

HDLM driver

A program that controls the HDLM functionality, manages paths, and detects errors.

HDLM manager

A program that manages error information. The HDLM manager receives the error information from the HDLM alert driver and then collects error logs.

HDLM nexus driver

A program that manages HDLM driver instances. The HDLM nexus driver manages the HDLM driver instances for each HBA port to which an sd or ssd driver, which is managed by HDLM, is connected.

host

A generic term for both servers and clients.

host device

A logical area within a host LU.

(See also: *host LU*, *host device name*)

host device name

A name assigned to a host device. A name of the logical device file for an HDLM device is assigned to a host device name.

A host device name is assigned using the c*U*t*X*d*Y* format.

The variables used in in c*U*t*X*d*Y* indicate the following:

U: Controller number reserved by HDLM

X: Target ID or World Wide Name of the sd or ssd device for the HDLM device

Y: LUN of the sd or ssd device for the HDLM device

(See also: *host device*)

host LU

An LU that the host recognizes.

Each host LU has a *host LU number*.

(See also: *LU*, *host LU number*, *host device*)

host LU number

A number assigned to a host LU. The host LU number is a part of a path name.

(See also: *host LU*, *path name*)

I

intermittent error

An error that occurs irregularly due to reasons such as a loose cable connection.

L

LDEV (Logical Device)

The value that identifies an LDEV consists of a combination of the storage subsystem's product name, serial number, and internal LU.

load balancing

A function for distributing the load across multiple paths that access an area within an LU. To distribute loads, load balancing uses multiple paths to perform I/O operations.

The following two types of algorithms are available: round-robin and extended round-robin.

local boot disk environment

An environment in which a boot disk is located on the host.

logical device file

Device files under `/dev/dsk` and `/dev/rdsk`.

Note that the term *logical device* used here is not an LDEV.

LU (Logical Unit)

A logical unit that is a logical volume defined in the storage subsystem, and with which the host can perform input or output operations.

N

node

A server of cluster members.

non-owner controller

A controller other than an owner controller.

Non-owner controllers exist only when the Freedom Storage 5700E, Freedom Storage 5800, Thunder 9200, Thunder 9500V Series, or TagmaStore AMS/WMS series is in use.

(See also: *owner controller*, *non-owner path*)

non-owner path

A path that passes through a controller other than an owner controller.

Non-owner paths exist only when the Freedom Storage 5700E, Freedom Storage 5800, Thunder 9200, Thunder 9500V Series, or TagmaStore AMS/WMS series is in use.

(See also: *owner path*, *non-owner controller*)

O

owner controller

A controller that has been set as an owner controller for an LU on the storage subsystem when the Freedom Storage 5700E, Freedom Storage 5800, Thunder 9200, Thunder 9500V Series, or TagmaStore AMS/WMS series is in use. A controller other than an owner controller is called a *non-owner controller*.

When the Freedom 7700E, Lightning 9900 Series, Lightning 9900V Series, TagmaStore USP, or Universal Storage Platform V is in use, all controllers are owner controllers.

(See also: *owner path*, *non-owner controller*)

owner path

A path that passes through a controller that has been set as an owner controller for an LU on the storage subsystem when the Freedom Storage 5700E, Freedom Storage 5800, Thunder 9200, Thunder 9500V Series, or TagmaStore AMS/WMS series is in use. A path that passes through a controller other than an owner controller is called a non-owner path.

When the Freedom 7700E, Lightning 9900 Series, Lightning 9900V Series, TagmaStore USP, or Universal Storage Platform V is in use, all paths are owner paths.

(See also: *owner controller*, *non-owner path*)

P

path

An access path from a host to a storage subsystem. Access to an area within an LU in a storage subsystem is made via a cable connecting the host bus adapter on the host and the channel adapter on the storage subsystem. Each path has an `AutoPATH_ID`.

(See also: *AutoPATH_ID*)

path health checking

A function for checking the status of paths at regular intervals.

When an error occurs at a path that was in the Online status, path health checking changes the status of the path to the Offline(E) status.

Path health checking checks only those paths that have Online status.

path name

The name consisting of the following four elements, separated by periods:

- Host port number (hexadecimal number)
- Bus number (hexadecimal number)
- Target ID (hexadecimal number)
- Host LU number (hexadecimal number)

A path name is used to identify a physical path.

(See also: *host LU number*)

persistent reservation

Similar to *reservation*, persistent reservation function enables a host to declare that it has exclusive use of an LU, and prevents other hosts from accessing that LU. Note, however, that reservation allows a host to have exclusive use of only one path to the LU, but persistent reservation allows a host to have exclusive use of more than one path.

If HDLM applies persistent reservation, a host can have exclusive use of multiple paths to an LU, so load balancing among these exclusively used paths is possible.

(See also: *reservation*)

physical device file

A device file under `/devices`.

physical path

An access path from a host to an LU. A path name is used to identify a physical path.

R

reservation

The reservation function enables a host to declare that it has exclusive use of an LU, and prevents other hosts from accessing that LU. Access permission for an LU that has been reserved is given only to the host that issued the reservation, so the LU cannot be accessed from multiple paths simultaneously. As a result, load balancing is not possible.

(See also: *persistent reservation*)

S

SAN (Storage Area Network)

A high-speed network connecting hosts and storage subsystems. This network is independent of a LAN and is dedicated to data transfer. A SAN provides faster accesses to storage subsystems, and prevents the transfer of high-volumes of data from deteriorating LAN performance.

SCSI device

A SCSI disk device

SLPR (Storage Logical Partition)

The SLPR function enables logical splitting of a storage subsystem supported by the TagmaStore USP and Universal Storage Platform V. This function splits the resources (ports, CLPR, and volumes) in the storage subsystem, so that the user can manage each resource independently.