HITACHI
DATA SYSTEMS

# HiCommand Dynamic Link Manager User's Guide for Linux®

MK-92DLM113-08

## Trademarks

SUSE is a registered trademark of SUSE LINUX.

UNIX is a registered trademark of The Open Group in the United States and other countries.

VERITAS, VERITAS Cluster Server, VERITAS File System, and VERITAS Net Backup are trademarks of VERITAS Symantec Corporation.

All other brand or product names are or may be trademarks, registered trademarks, or service marks of and are used to identify products or services of their respective owners.

## 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-92DLM113-00 | November 2003 | Initial Release |
| MK-92DLM113-01 | March 2004 | Revision 1, supersedes and replaces MK-92DLM113-00 |
| MK-92DLM113-02 | June 2004 | Revision 2, supersedes and replaces MK-92DLM113-01 |
| MK-92DLM113-03 | August 2004 | Revision 3, supersedes and replaces MK-92DLM113-02 |
| MK-92DLM113-04 | April 2005 | Revision 4, supersedes and replaces MK-92DLM113-03 |
| MK-92DLM113-05 | September 2005 | Revision 5, supersedes and replaces MK-92DLM113-04 |
| MK-92DLM113-06 | May 2006 | Revision 6, supersedes and replaces MK-92DLM113-05 |
| MK-92DLM113-07 | March 2007 | Revision 7, supersedes and replaces MK-92DLM113-06 |
| MK-92DLM113-08 | May 2007 | Revision 8, supersedes and replaces MK-92DLM113-07 |

# Preface

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

The readers of this document must have a basic knowledge of the following areas:

- Linux® and its management functionality,
- Storage subsystem management 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),
- Cluster software functionality, and
- Volume management software functionality.

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.

***Notice:*** 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 Linux® version 5.9.1.

## Convention for Storage Capacity Values

This document uses the following convention for storage capacity values:

1 KB = 1,000 bytes
1 MB = $1,000^2$ bytes
1 GB = $1,000^3$ bytes
1 TB = $1,000^4$ bytes

## Referenced Documents

- *Hitachi Universal Storage Platform V User and Reference Guide, MK-96RD635*
- *Hitachi Lightning 9900™ V Series User and Reference Guide*, *MK-92RD100*
- *Hitachi Lightning 9900™ User and Reference Guide*, *MK-90RD008*
- *Hitachi Thunder 9500™ V Series User and Reference Guide*, *MK-92DF601*
- *Hitachi Thunder 9200™ User and Reference Guide*, *MK-90DF504*

## 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.


- **E-mail:** doc.comments@hds.com
- **Fax:** 858-695-1186
- **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 4 HDLM Operation

# List of Figures

# List of Tables

# 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?

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 SANs, 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*.

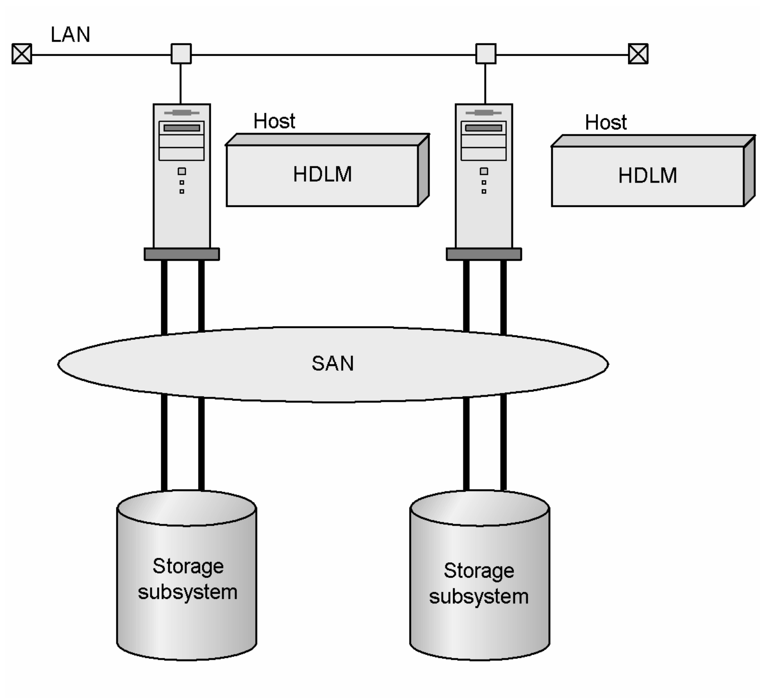


**Figure 1.1 Between Hosts and Storage Subsystems**

HDLM supports the following storage subsystems:

- Thunder 9200
- Lightning 9900 Series
- Thunder 9500V Series
- Lightning 9900V Series
- TagmaStore AMS/WMS series
- TagmaStore USP
- Universal Storage Platform V

## 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 section 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 section 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 an automatic failback, HDLM automatically restores the route to the active state after the user has corrected the problem in the physical route.

For details on failback, see section 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 section 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

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:

- SCSI devices in Hitachi storage subsystems.
- Boot disks

Devices that HDLM cannot manage:

- SCSI devices other than those for Hitachi storage subsystems
- Hitachi storage subsystem command devices, such as Hitachi RAID Manager command devices
- Devices other than disks (such as tape devices)

## 2.2 System Configuration

HDLM manages routes between a host and a storage subsystem by using the SCSI driver. The host and storage subsystems are connected using SAN with fiber 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 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 a physical path and manages the paths on a physical-path basis. When you use HDLM, there is no need to consider the difference between a physical path and a path. Thus, hereafter both physical paths and paths might be called *paths*, without a distinction being made between the two. The ID that HDLM assigns for each physical path is called an *AutoPATH_ID*. Also, a path might be called a *management target*.

Figure 2.1 shows the HDLM system configuration.

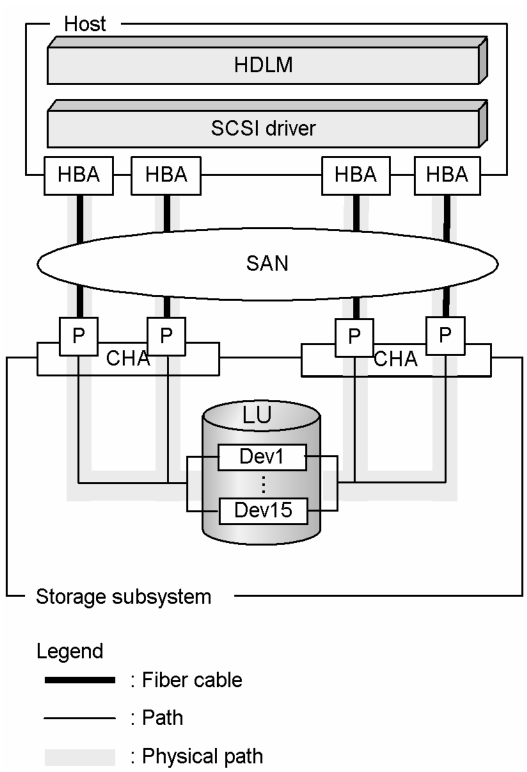


**Figure 2.1 HDLM System Configuration**

Table 2.1 lists the HDLM system components.

**Table 2.1 HDLM System Components**

| Components | Description |
|---|---|
| CHA | A channel adapter. |
| Dev | An area (partition) of a divided LU. |
| HBA | A host bus adapter. This serves as a cable port on the host. |
| 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. |
| P | A port on a channel adapter. This serves as a cable port on a storage subsystem. |
| SAN | A dedicated network that is used for data transfer between the host machine and storage subsystems. |
| 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 a SCSI device is connected to each path.

Thus, a single LU in the storage subsystem is recognized as though there are as many LUs as there are paths.

After installing HDLM:

An HDLM device corresponding one-to-one with an LU in the storage subsystem is created at a level higher than the SCSI device.#

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

#

In addition to the one that indicates the entire LU, a logical device file for the HDLM device is created for each partition.

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 (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 SCSI device.

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

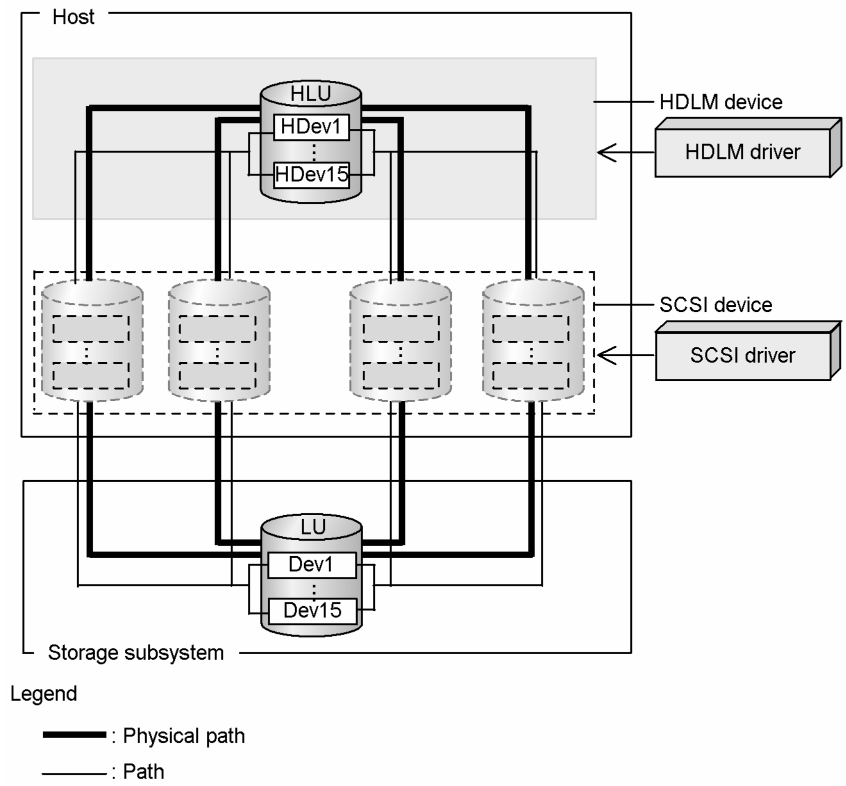


**Figure 2.2 LU Configuration Recognized by the Host After HDLM Installation**

Table 2.2 lists the components recognized by the host.

**Table 2.2 LU Components**

| Components | Description |
|---|---|
| HDev | A Dev (partition) in an LU that the host recognizes via the HDLM driver. It is called a *host device*. One host device is recognized for one Dev in the storage subsystem. |
| HLU | An LU that the host recognizes via the HDLM driver. It is called a *host LU*. Regardless of how many paths exist, only one host LU is recognized for each LU 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.

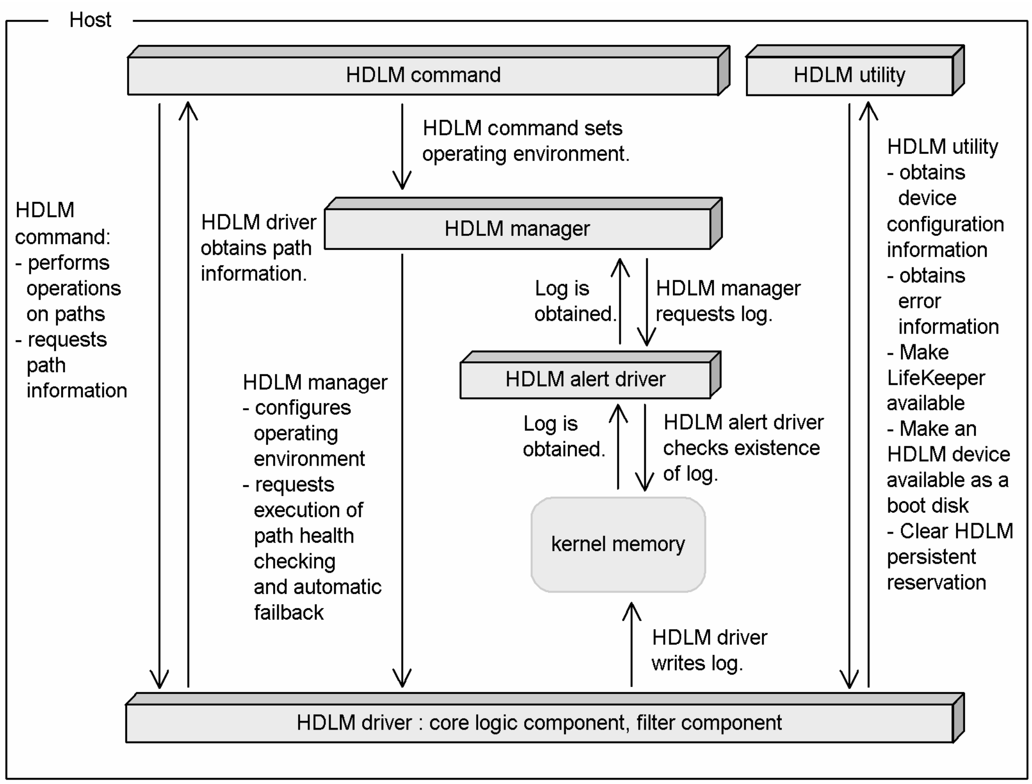


**Figure 2.3 HDLM Program Configuration**

Table 2.3 lists and describes the functionality of these programs.

**Table 2.3 Functionality of HDLM Programs**

| Program name | Functionality |
|---|---|
| HDLM command | Provides the `dlnkmgr` command to enable you to:<br>▪ Manage paths<br>▪ Display error information<br>▪ Set up the HDLM operating environment |
| HDLM utility | Provides HDLM utilities to enable you to:<br>▪ Collect error information<br>▪ Define HDLM device configuration information<br>▪ Make LifeKeeper available<br>▪ Make an HDLM device available as a boot disk |

| Program name | Functionality |
| --- | --- |
| | ▪ Clear HDLM persistent reservation |
| HDLM manager | Provides the HDLM manager to enable you to:<br>▪ Configure the HDLM operating environment<br>▪ Request the execution of path health checks and automatic failback<br>▪ Collect error log data |
| HDLM alert driver | Reports the log information collected by the HDLM driver to the HDLM manager. The driver name is `sddlmadrv`. |
| HDLM driver | Controls HDLM functionality, manages paths, and detects errors. The HDLM driver consists of the following:<br>▪ Core logic component<br>Controls the basic functionality of HDLM.<br>▪ Filter component<br>Sends and receives I/O. The driver name is `sddlmfdrv`. |

## 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 device.

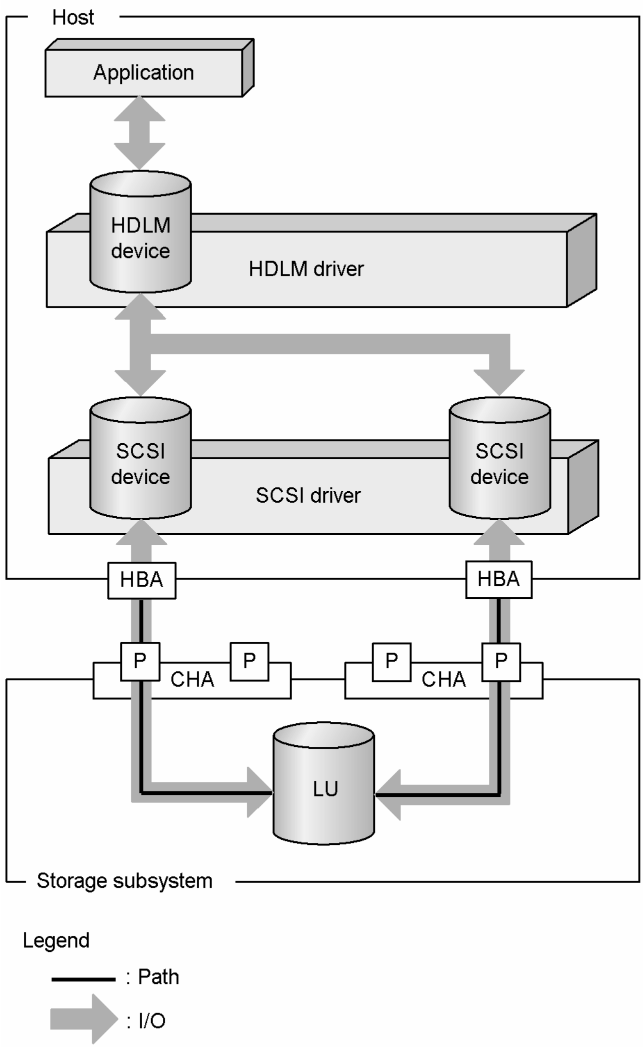


**Figure 2.4 Position of the HDLM Driver and HDLM Devices**

## 2.6 Logical Device Files for HDLM Devices

The logical device file name of an HDLM device is different from the logical device file name of a SCSI device. When you configure the logical device file of an HDLM device for applications such as volume management software, these applications can access the LUs that HDLM manages.

The following shows an example of the logical device file name that the application uses to access the LU (for accesses before and after HDLM installation).

Table 2.4 illustrates the logical device file name of the device that the application uses, for before and after HDLM installation.

**Table 2.4 Example of Using the Logical Device File Name of the Device Used When the Application Accesses the LU**

| Host status | Device file name that the application uses |
|---|---|
| Before installing HDLM | The application uses the logical device file name for the SCSI device.<br>Example:<br>`sda`<br>`sdb` |
| After installing HDLM | The application uses the logical device file name for the HDLM device.<br>Example:<br>`sddlmaa` |

The logical device file name of an HDLM device has the following format:

```
/dev/sddlm[a-p][a-p][1-15]
```

Where:

*a-p*: A character from a to p

*1-15*: A number from 1 to 15

A major number is required for each of the first [*a-p*] values. [*1-15*] indicates a partition number in the applicable LU. For example, if the logical device file name of an HDLM device is `sddlmaa1`, it indicates partition 1 on `sddlmaa`. To specify the entire LU, simply use `sddlmaa`. Note that HDLM creates block device files. The system dynamically selects the major number of the block device that this file uses.

## 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 functionality is called *load balancing,* and it prevents a heavily loaded path from affecting the performance of the entire system.

***Note:*** 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 operation being issued for the same LU from multiple applications.

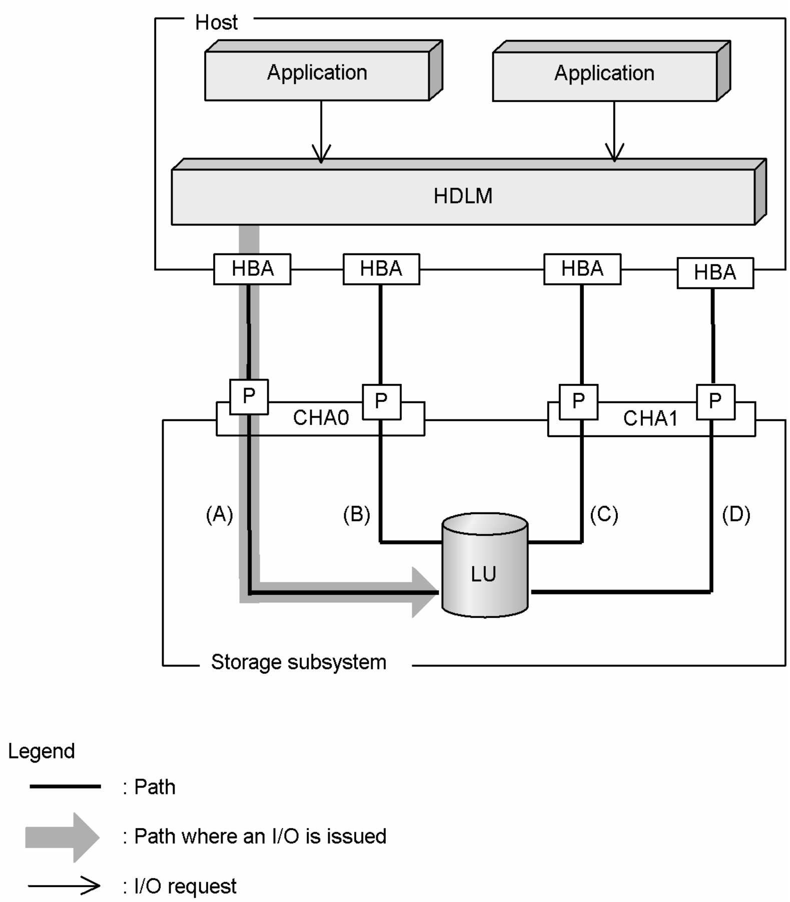


**Figure 2.5 I/O Flow When the Load Balancing Function is Not Being Used**

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

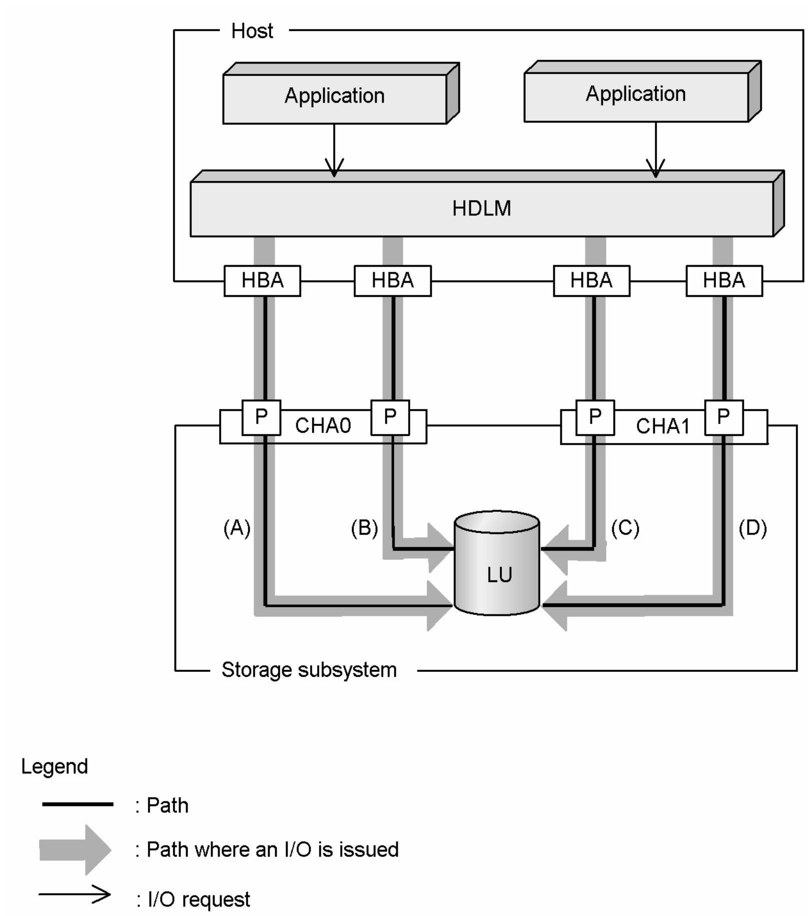


**Figure 2.6 I/O Flow When the Load Balancing Function Is Being Used**

When the load balancing function is being used, I/O operations are distributed via 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 Paths To Which Load Balancing Is Applied

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

#### 2.7.1.1 When Using the Thunder 9200, Thunder 9500V, 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*). Paths used for load balancing are selected from owner paths first, then 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 LU0 is CHA0. When the LU is accessed, the load is balanced between the paths (A) and (B) (that is, between owner paths). When the path (A) cannot be used due to a problem such as a failure, the LU can only be accessed via the 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).

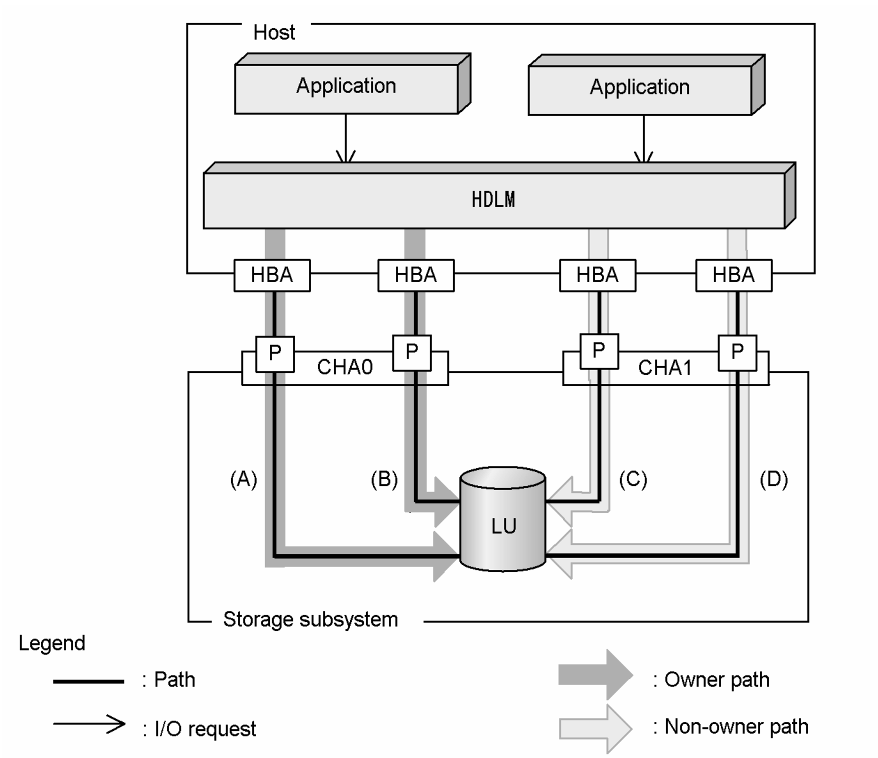


**Figure 2.7 Overview of Load Balancing**

#### 2.7.1.2 When Using the Lightning 9900, Lightning 9900V Series, TagmaStore USP, or Universal Storage Platform V

All the online paths are owner paths. Thus, for the example in Figure 2.6, when the LU is accessed, the load is balanced among the 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 algorithms for load balancing:

- Round robin

  Round robin distributes all I/Os among multiple paths.

- Extended round robin

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

**Table 2.5 Algorithms for Load Balancing**

| Algorithm | For sequential access | For random access |
|---|---|---|
| Round robin | ▪ After an I/O operation is issued to a path once or a certain number of times, the path is switched to the next path.<br>▪ The storage subsystem cache might not be fully usable.<br>▪ When multiple applications that request sequential access are run concurrently, we recommend that you use the round robin algorithm in order to distribute I/Os across multiple paths.# | After an I/O operation 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 operation 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 operation is issued to a path a certain number of times, the path is switched to the next path when sequential access is switched to random access.<br>▪ The storage subsystem cache can be used.<br>▪ 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 section 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 functionality 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 functionality 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 section 6.4.

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

***Note:***

- When using the Thunder 9200, Thunder 9500V Series, or TagmaStore AMS/WMS series make sure that you set the data share mode to **ON** (**Used**).

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

### 2.8.1 Automatic Path Switching

The following describes the automatic failover and automatic failback functionalities that automatically switch a path.

#### 2.8.1.1 Automatic Failover

If you detect an error for a path in use, you can keep operating the system by changing the state to offline, and using other online paths. This functionality 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 section 2.11.2.

When the 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 path (B) and the second candidate is the path (C) or (D).

When the 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 paths (B), (C), and (D).

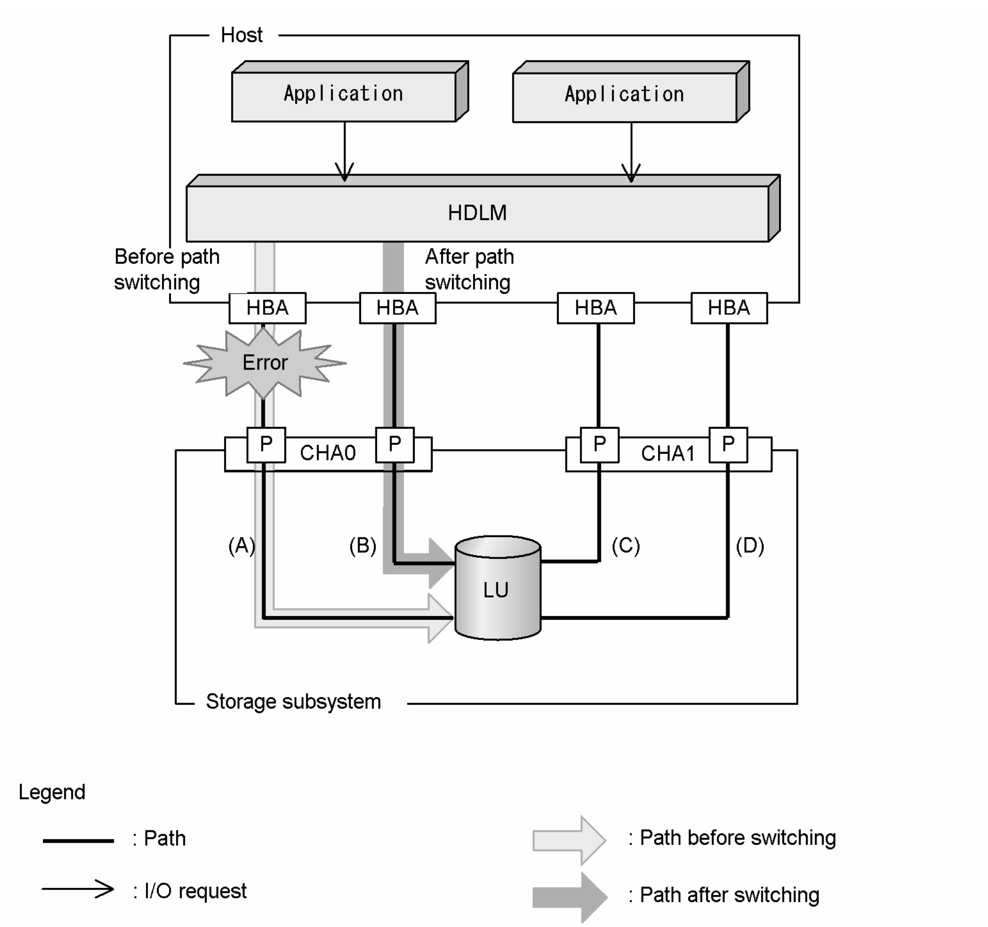


**Figure 2.8 Path Switching**

### 2.8.1.2 Automatic Failback

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

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

When using the Thunder 9200, or Thunder 9500V Series, or TagmaStore AMS/WMS series, 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 to use is switched to the owner path.

When the 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 to use is not switched.

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 by the `dlnkmgr` command's `set` operation. For details on the `set` operation, see section 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 path by manually placing a path online or offline. Manually switching a path 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 section 6.5. For details on the `offline` operation, see section 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 using the Thunder 9200, or Thunder 9500V Series, or TagmaStore AMS/WMS series, HDLM selects the switching destination path from owner paths and then from non-owner paths. When the 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. All 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 section 6.5. After the path status is changed to online (by executing the `online` operation), HDLM selects the path to use in the same way as for automatic path switching. When using the Thunder 9200, or Thunder 9500V Series, or TagmaStore AMS/WMS series, HDLM selects the path to use from online owner paths, and then from online non-owner paths. When the 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 by using the `online` operation.

### 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 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 on the path and, among the paths that access the same LU, none of those are in the `Online` status.

  If none of the paths accessing a single LU are in the `Online` status, one of the paths is changed to the `Online(E)` status. This ensures that the LU can be accessed, by making sure that all paths are not offline.

  The (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 section 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 performed because an error occurred in the path.

  The (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.

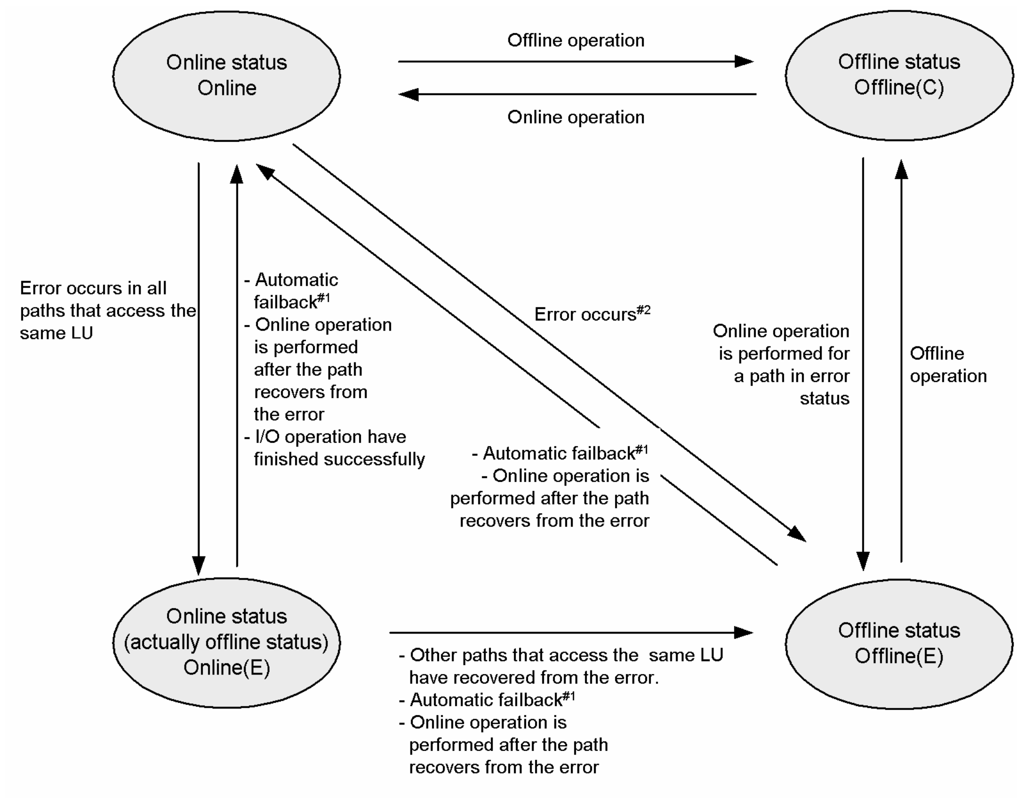


**Figure 2.9 Path Status Transitions**

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 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 operations issued to an LU is successful.

#2

When an Online or 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 section 6.4.

If an error occurs in the last available online path for each LU, the status of the path is changed to `Online(E)`.

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 online 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 made offline by using an HDLM command, Offline(C) might change to Offline(E). If an offline operation was performed, wait for a fixed period of time (about 1 minute), check the path status by using an HDLM 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 a special error status, called *not subject to auto failback*, until the user places the path online. Failback is not performed for the path.

### 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 section 6.7.

### 2.9.2 Setting Up Intermittent Error Monitoring

When you use the intermittent error functionality, you can enable or disable the functionality. If you enable the functionality, 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 section 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 following actions for intermittent error monitoring:

- 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 change to online (`Online`). For details on the conditions, see Figure 2.12.

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 event occurred in one path is described on the time arrow.

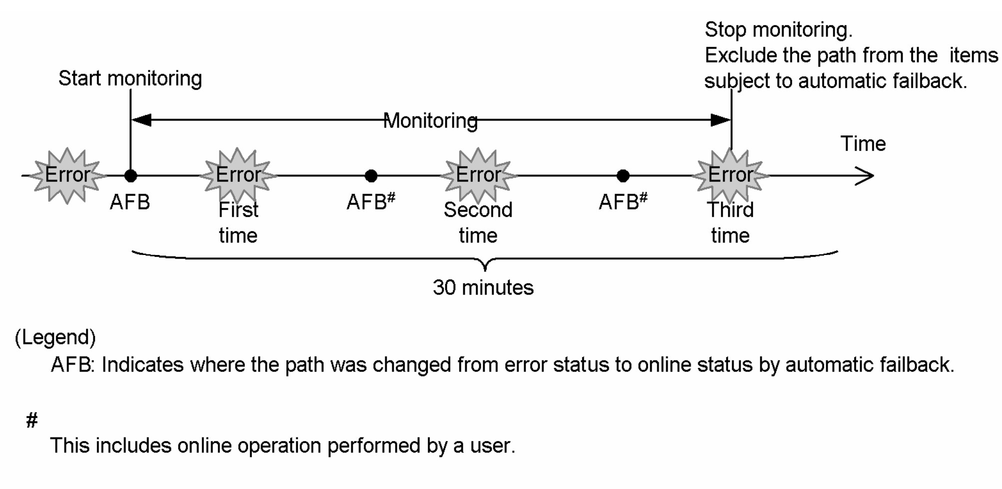


**Figure 2.10 Action When an Intermittent Error Occurs in the Path**

### 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, an intermittent error does not occur. 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 event occurring in one path is described on the time arrow.

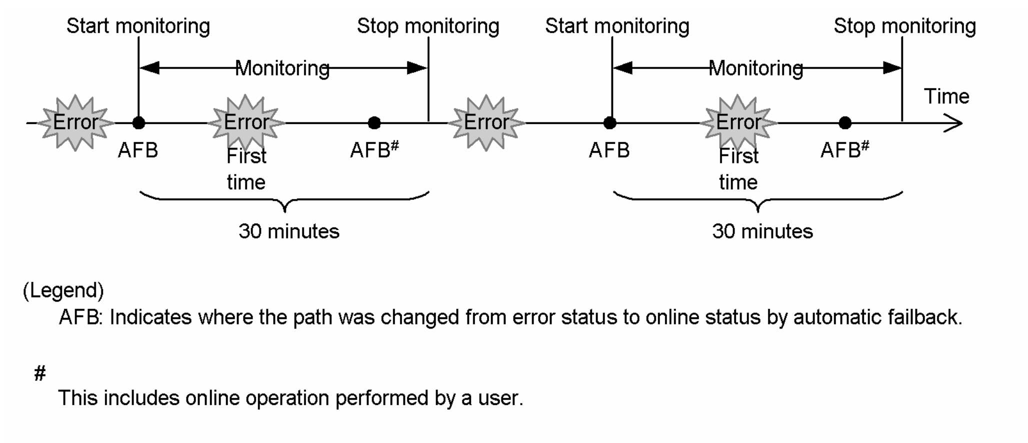


**Figure 2.11 Action When an Intermittent Error Does Not Occur in the Path**

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 operations 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 written on the time arrow.

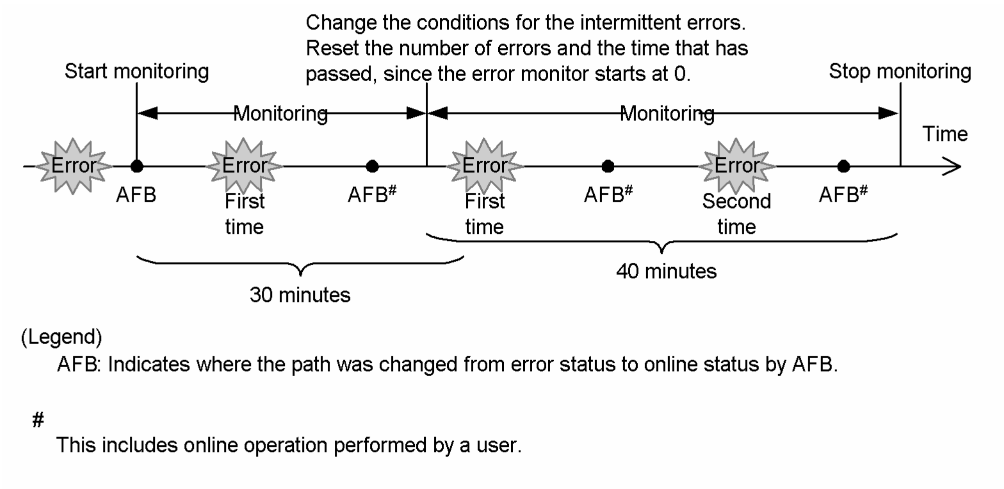


**Figure 2.12 Action When the Conditions for the Intermittent Error Are Changed During Error Monitoring**

## 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 2.6 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 **Intermittent Error Path** item, then intermittent error monitoring is being performed.

**Table 2.6 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 |

| User operation | | Number of errors and time passed since error monitoring started | Information about paths not subject to automatic failback |
|---|---|---|---|
| | Setting `on` during intermittent error monitoring by executing the `set` operation (the conditions for the intermittent error monitoring are not changed) | | |
| | Changing the setting for intermittent error monitoring to 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 change the intermittent error monitoring functionality to off, 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 is reset to `0`, and then monitoring restarts in accordance with the changed monitoring conditions.

#3

The number of errors is reset to `0`, and then monitoring restarts.

## 2.10 Detecting Errors by Using Path Health Checking

HDLM can check the status of paths at regular intervals, and detect errors. This functionality is called *path health checking*.

Without path health checking, an error is not detected unless I/O is performed because the system only checks the path status when I/O is performed. With path health checking, however, the system checks the status of online paths at regular intervals regardless of whether I/O is performed. If an error is detected in a path, path health checking functionality 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 performed 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 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 section 6.6.

## 2.11 Error Management

HDLM collects information for troubleshooting into log files. HDLM can also filter error information according to the error level when collecting the information.

Figure 2.13 shows the data flow when collecting error information.

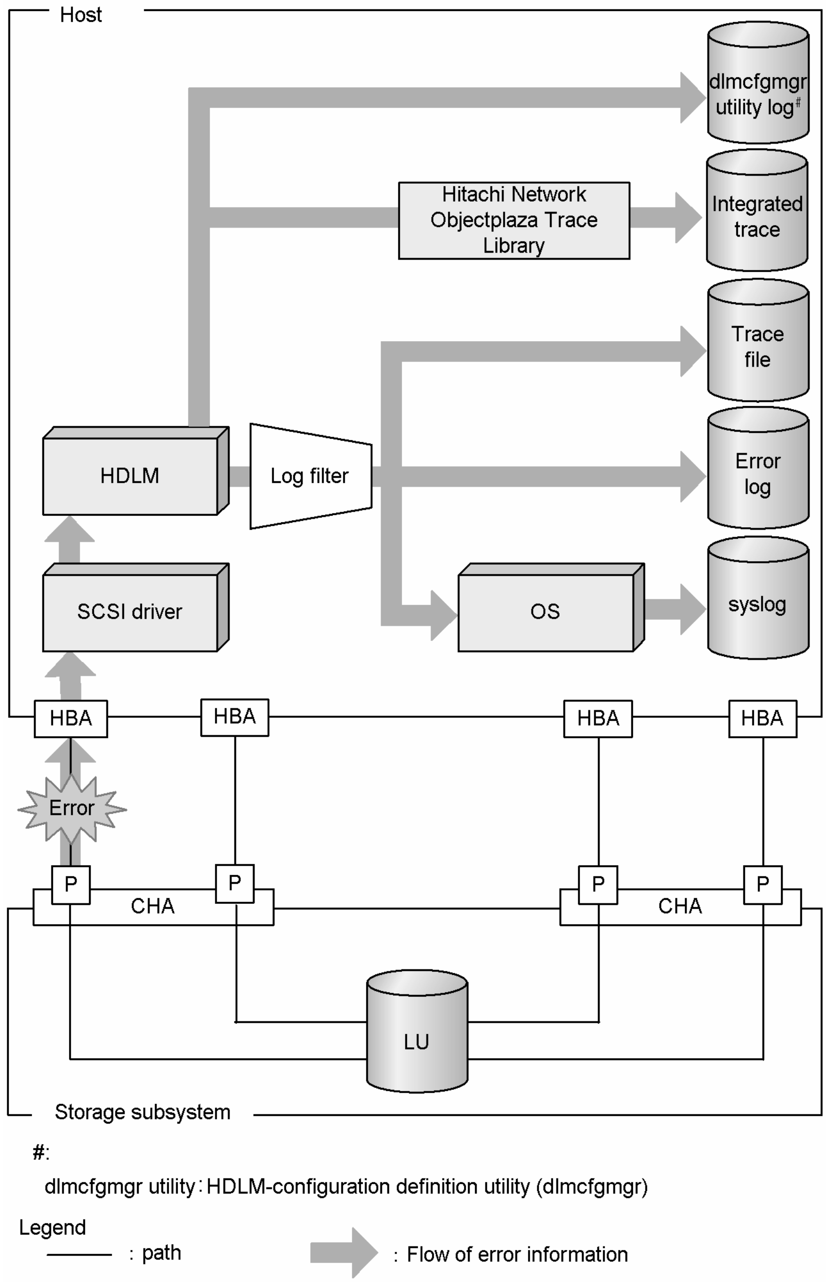


**Figure 2.13 Data Flow When Collecting Error Information**

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

### 2.11.1 Types of Logs Collected

HDLM collects information on the detected error 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 2.7 describes the types of error information.

**Table 2.7 Types of Error Information**

| Log name | Description | Output destination |
|---|---|---|
| Integrated trace file | Operation logs of the HDLM command are collected. | The default file path is `/var/opt/hitachi/HNTRLib2/spool/hntr2[1-16].log`.<br>To specify the output destination directory and the file prefix for integrated trace files, 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 following is the trace file name:<br>`/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:<br>`/var/opt/DynamicLinkManager/log/dlmmgr[1-16].log`<br>HDLM remote access interface logs:<br>`/var/opt/DynamicLinkManager/log/dlmwebagent[1-`*N*`].log`<br>The value *N* depends on the setting in the file `dlmwebagent.properties`. |
| *HDLM-configuration definition utility (dlmcfgmgr) log* | Logs are collected when the `dlmcfgmgr` utility is executed. | The following is the log file name.<br>▪ `/var/opt/DynamicLinkManager/log/dlmcfgmgr[1-2].log`<br>▪ `/var/opt/DynamicLinkManager/log/dlminquiry.log` |
| Syslog | The HDLM messages on or above the level set by the user with syslogd settings file are collected. # | The default file path is `/var/log/messages`.<br>The syslog file path is specified in the syslogd settings file. For details, see the Linux documentation. |

| Log name | Description | Output destination |
|---|---|---|
| | 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. | |

#

When you want to configure the system so that HDLM messages are output to syslog, specify `user` for the facility in the syslog settings 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 section 2.11.2.

### 2.11.2 Filtering of Error Information

Errors that HDLM detects are classified into error levels. Table 2.8 lists the error levels, in order of adverse effect on the system.

**Table 2.8 Error Levels**

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

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 section 8.1.1.

In syslog, the HDLM messages on or above the level set by the user configured in syslogd settings file are collected. It is recommended that you set the information to be output at the Information level or higher.

***Note:*** all the facilities when HDLM outputs messages in syslog are 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 section 3.12.2.

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

HDLM has 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, 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 section 7.2.

### 2.11.4 Utility for Collecting HDLM Installation Error Information (installgetras)

HDLM has a utility for collecting HDLM installation error information (`installgetras`).

By executing this utility, you can collect the logs required for analyzing errors that occurred during installation. You can use the collected information when you contact your HDLM vendor or maintenance company.

For details on the `installgetras` utility, see section 7.8.

## 2.12 Cluster Support

HDLM is also available for cluster configurations.

For details about the cluster software supported by HDLM, see section 3.1.3.

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

The details of host switching depend on the application.

***Notes:***

- The DiskReservation agent of the VERITAS Cluster Server is not supported.
- 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
```

# Chapter 3 Creating a HDLM Environment

This chapter explains the procedure for setting up an HDLM environment and the procedure for canceling the environment settings.

Make sure that HDLM installation and function setup has been performed. Set up volume groups and cluster software according to the environment you are using.

## 3.1 System Requirements

Check the following before installing HDLM:

### 3.1.1 Applicable Models

#### 3.1.1.1 Applicable Hosts

You can install HDLM on a host that is running an OS described in Table 3.1, and meets the requirements listed in Table 3.1.

**Table 3.1 Requirements for Applicable Hosts**

| Items | Requirements |
|---|---|
| CPU | Intel Pentium III or Itanium2 833MHz or more#1 |
| | AMD Opteron |
| Memory | 512 MB or more |
| Disk size | 170 MB#2 or more |

#1

HDLM is compatible with Hyper-Threading technology.

#2

The disk space required for installation.

You can install HDLM on the host that is running the OS listed in Table 3.2.

**Table 3.2 Applicable OSs for Hosts**

| OS#1 | Kernel |
|---|---|
| Red Hat Enterprise Linux AS 3 (IA32)#2<br>Red Hat Enterprise Linux ES 3 (IA32)#2 | 2.4.21-9.0.1.EL, 2.4.21-9.0.1.ELsmp, or 2.4.21-9.0.1.ELhugemem |
| | 2.4.21-15.EL, 2.4.21-15.ELsmp, or 2.4.21-15.ELhugemem |
| | 2.4.21-20.EL, 2.4.21-20.ELsmp, or 2.4.21-20.ELhugemem |
| | 2.4.21-27.EL, 2.4.21-27.ELsmp, or 2.4.21-27.ELhugemem#3 |
| | 2.4.21-27.0.1.EL, 2.4.21-27.0.1.ELsmp, or 2.4.21-27.0.1.ELhugemem |
| | 2.4.21-27.0.2.EL, 2.4.21-27.0.2.ELsmp, or 2.4.21-27.0.2.ELhugemem |

| OS#1 | Kernel |
|---|---|
| | 2.4.21-32.0.1.EL, 2.4.21-32.0.1.ELsmp, or 2.4.21-32.0.1.ELhugemem |
| | 2.4.21-37.EL, 2.4.21-37.ELsmp, or 2.4.21-37.ELhugemem |
| | 2.4.21-47.EL, 2.4.21-47.ELsmp, or 2.4.21-47.ELhugemem |
| Red Hat Enterprise Linux AS 3 (IPF)#4 | 2.4.21-9.0.1.EL |
| | 2.4.21-15.EL |
| | 2.4.21-20.EL |
| | 2.4.21-27.EL |
| | 2.4.21-27.0.1.EL |
| | 2.4.21-27.0.2.EL |
| | 2.4.21-32.0.1.EL |
| | 2.4.21-37.EL |
| | 2.4.21-47.EL |
| Red Hat Enterprise Linux ES 3 (IPF)#4 | 2.4.21-20.EL |
| | 2.4.21-27.EL |
| | 2.4.21-27.0.1.EL |
| | 2.4.21-27.0.2.EL |
| | 2.4.21-32.0.1.EL |
| | 2.4.21-37.EL |
| | 2.4.21-47.EL |
| Red Hat Enterprise Linux AS 3 (EM64T)#5 | 2.4.21-20.EL |
| Red Hat Enterprise Linux ES 3 (EM64T)#5 | 2.4.21-27.EL |
| | 2.4.21-32.0.1.EL |
| | 2.4.21-37.EL |
| | 2.4.21-47.EL |
| Red Hat Enterprise Linux AS 3 (AMD64)#6 | 2.4.21-32.0.1.EL or 2.4.21-32.0.1.ELsmp |
| Red Hat Enterprise Linux ES 3 (AMD64)#6 | 2.4.21-37.EL or 2.4.21-37.ELsmp |
| | 2.4.21-47.EL or 2.4.21-47.ELsmp |
| Red Hat Enterprise Linux AS 4 (IA32)#2<br>Red Hat Enterprise Linux ES 4 (IA32)#2 | 2.6.9-11.EL, 2.6.9-11.ELsmp, or 2.6.9-11.ELhugemem |
| | 2.6.9-22.EL, 2.6.9-22.ELsmp, or 2.6.9-22.ELhugemem |
| | 2.6.9-34.EL, 2.6.9-34.ELsmp, or 2.6.9-34.ELhugemem |

| OS[#1] | Kernel |
| --- | --- |
| | 2.6.9-42.EL, 2.6.9-42.ELsmp[#7], or 2.6.9-42.ELhugemem |
| | 2.6.9-42.0.3.EL, 2.6.9-42.0.3.ELsmp, or 2.6.9-42.0.3.ELhugemem |
| Red Hat Enterprise Linux AS 4 (IPF)[#4] | 2.6.9-11.EL |
| Red Hat Enterprise Linux ES 4 (IPF)[#4] | 2.6.9-22.EL |
| | 2.6.9-34.EL |
| | 2.6.9-42.EL |
| | 2.6.9-42.0.3.EL |
| Red Hat Enterprise Linux AS 4 (EM64T/AMD64)[#8] | 2.6.9-11.EL or 2.6.9-11.ELsmp |
| Red Hat Enterprise Linux ES 4 (EM64T/AMD64)[#8] | 2.6.9-22.EL or 2.6.9-22.ELsmp |
| | 2.6.9-34.EL, 2.6.9-34.ELsmp, or 2.6.9-34.ELhugemem |
| | 2.6.9-42.EL, 2.6.9-42.ELsmp, or 2.6.9-42.ELhugemem |
| | 2.6.9-42.0.3.EL, 2.6.9-42.0.3.ELsmp, or 2.6.9-42.0.3.ELlargesmp |
| SUSE LINUX Enterprise Server 9 (IA32)[#2] | 2.6.5-7.139-default, 2.6.5-7.139-smp, or 2.6.5-7.139-bigsmp[#9] |
| | 2.6.5-7.191-default, 2.6.5-7.191-smp, or 2.6.5-7.191-bigsmp[#10] |
| | 2.6.5-7.201-smp[#10] |
| | 2.6.5-7.244-default, 2.6.5-7.244-smp, or 2.6.5-7.244-bigsmp[#11] |
| SUSE LINUX Enterprise Server 9 (IPF)[#4] | 2.6.5-7.139-default or 2.6.5-7.139-64k-pagesize[#9] |
| | 2.6.5-7.191-default or 2.6.5-7.191-64k-pagesize[#10] |
| | 2.6.5-7.244-default or 2.6.5-7.244-64k-pagesize[#11] |
| SUSE LINUX Enterprise Server 9 (EM64T/AMD64)[#8] | 2.6.5-7.191-default or 2.6.5-7.191-smp[#10] |
| | 2.6.5-7.244-default or 2.6.5-7.244-smp[#11] |
| | 2.6.5-7.252-smp[#11] |
| SUSE LINUX Enterprise Server10 (IA32)[#2 #`12] | 2.6.16.21-0.8-default, 2.6.16.21-0.8-smp, or 2.6.16.21-0.8-bigsmp |
| | 2.6.16.27-0.9-default, 2.6.16.27-0.9-smp, or 2.6.16.27-0.9-bigsmp |
| SUSE LINUX Enterprise Server10 (IPF)[#4 #`12] | 2.6.16.21-0.8-default |
| | 2.6.16.27-0.9-default |
| SUSE LINUX Enterprise Server10 (EM64T/AMD64)[#8 #`12] | 2.6.16.21-0.8-default or 2.6.16.21-0.8-smp |

| OS#1 | Kernel |
|---|---|
| | 2.6.16.27-0.9-default or 2.6.16.27-0.9-smp |

#1

Kernels that are provided by OS distributors in binary format are only supported.

#2

IA32 indicates an environment where an IA32 kernel is installed on a system that uses an Intel CPU.

#3

HDLM also supports environments in which the following kernel packages for Red Hat Enterprise Linux AS/ES 3 for IA32 are installed on a system using an AMD Opteron (Single Core) processor. The HDLM functions will be the same as those applied on systems that use Intel processors.

- kernel-2.4.21-27.EL.i686.rpm

- kernel-smp-2.4.21-27.EL.i686.rpm

- kernel-hugemem-2.4.21-27.EL.i686.rpm

#4

IPF indicates an environment where an IPF kernel is installed on a system that uses an Intel CPU.

#5

EM64T indicates an environment where an EM64T kernel is installed on a system that uses an Intel CPU.

#6

AMD64 indicates an environment where an AMD64 kernel is installed on a system that uses an AMD Opteron CPU (Single Core and Dual Core).

#7

HDLM also supports environments in which the following kernel packages for Red Hat Enterprise Linux AS/ES 4 for IA32 are installed on a system using an AMD Opteron (Single Core and Dual Core) processor.

- kernel-smp-2.6.9-42.EL.i686.rpm

#8

EM64T and AMD64 kernels are the same. EM64T/AMD64 indicates an environment where an EM64T/AMD64 kernel is installed on a system that uses an Intel CPU or AMD Opteron CPU (Single Core and Dual Core).

#9

All of the packages of SP1 for SUSE LINUX Enterprise Server 9 must be installed.

#10

All of the packages of SP2 for SUSE LINUX Enterprise Server 9 must be installed.

#11

All of the packages of SP3 for SUSE LINUX Enterprise Server 9 must be installed.

#12

A gdb package of version 6.5-21.2 or later must be installed.

#### 3.1.1.2 Applicable Host Bus Adapters (HBAs)

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

#### 3.1.1.3 Applicable Storage Subsystems

The following storage subsystems are applicable to HDLM:

- Thunder 9200 (Note that SCSI connections are not supported.)
- Thunder 9500V series
- Lightning 9900 series
- Lightning 9900V series
- TagmaStore USP
- Hitachi Universal Storage Platform V
- TagmaStore AMS
- TagmaStore WMS

Storage subsystems to be applied must have a dual controller configuration. If you use them in a HUB-connected environment, specify a unique loop ID for all the connected hosts and storage subsystems. For details on the microprogram version and storage subsystem setting information required for using HDLM, see HDLM *Release Notes*.

***Note:*** BladeSymphony supports the Thunder 9500V series, the Lightning 9900V series, the TagmaStore USP, and the Universal Storage Platform V. In a BladeSymphony environment, the Lightning 9900V series cannot be used as a boot disk.

### 3.1.2 HDLM Package Configuration

Table 3.3 lists the components included in the HDLM package.

**Table 3.3 Components Packaged With This Product**

| Components packaged with this product | | Quantity |
|---|---|---|
| Product medium | CD-ROM | 1 |
| Documentation | Release Notes | 1 |

| Components packaged with this product | | Quantity |
|---|---|---|
| | HiCommand Dynamic Link Manager User's Guide for Linux® (this manual) | 1 |
| | Statement of Delivery and Usage Consent Form | 1 |

### 3.1.3 Related Programs

#### 3.1.3.1 When combining cluster configurations

Table 3.4 lists the related programs used when creating a cluster.

**Table 3.4 Related Programs Used When Creating a Cluster**

| OS | Related programs |
|---|---|
| Red Hat Enterprise Linux AS 3 (IA32)<br>Red Hat Enterprise Linux ES 3 (IA32) | VERITAS Cluster Server 2.2[#1#2] |
| | VERITAS Cluster Server 4.0[#1#3] |
| | Red Hat Cluster Manager 1.2.9-1[#4] |
| | Oracle9i RAC Release 2(9.2.0)[#5] |
| | Oracle9i RAC Release 2 (9.2.0.5.0)[#6] |
| | Oracle9i RAC Release 2 (9.2.0.6.0)[#7] |
| | Oracle RAC 10g Release 1 (10.1.0.2)[#8] |
| | Oracle RAC 10g Release 1 (10.1.0.4)[#9] |
| | Oracle RAC 10g Release 1 (10.1.0.5)[#10] |
| | Oracle RAC 10g Release 2 (10.2.0.1)[#11] |
| | Oracle RAC 10g Release 2 (10.2.0.2)[#12] |
| | LifeKeeper 4.6.2[#13] |
| Red Hat Enterprise Linux AS 3 (IPF)<br>Red Hat Enterprise Linux ES 3 (IPF) | VERITAS Cluster Server 2.2[#1#2] |
| | Oracle9i RAC Release 2(9.2.0.6.0)[#14] |
| Red Hat Enterprise Linux AS 3 (EM64T)<br>Red Hat Enterprise Linux ES 3 (EM64T) | Red Hat Cluster Manager 1.2.22-2[#15] |
| Red Hat Enterprise Linux AS 3 (AMD64)<br>Red Hat Enterprise Linux ES 3 (AMD64) | Oracle9i RAC Release 2(9.2.0.6)[#16] |
| Red Hat Enterprise Linux AS 4 (IA32)<br>Red Hat Enterprise Linux ES 4 (IA32) | Red Hat Cluster Manager 1.0.12-1.0 |
| | Oracle9i RAC Release 2(9.2.0.8.0)[#17] |
| | Oracle RAC 10g Release 2 (10.2.0.2)[#18] |
| Red Hat Enterprise Linux AS 4 (IPF) | Red Hat Cluster Manager 1.0.12-1.0[#19] |

| OS | Related programs |
|---|---|
| Red Hat Enterprise Linux ES 4 (IPF) | Oracle RAC 10g Release 2 (10.2.0.1)#20 |
| | Oracle RAC 10g Release 2 (10.2.0.2)#20 |
| Red Hat Enterprise Linux AS 4 (EM64T/AMD64)<br>Red Hat Enterprise Linux ES 4 (EM64T/AMD64) | Red Hat Cluster Manager 1.0.12-1.0 |
| | Oracle RAC 10g Release 2(10.2.0.2)#21 |
| SUSE LINUX Enterprise Server 9 (IA32) | None |
| SUSE LINUX Enterprise Server 9 (IPF) | None |
| SUSE LINUX Enterprise Server 9 (EM64T/AMD64) | Oracle RAC 10g Release 2 (10.2.0.2.0)#22 |
| SUSE LINUX Enterprise Server10 (IA32) | None |
| SUSE LINUX Enterprise Server10 (IPF) | None |
| SUSE LINUX Enterprise Server10 (EM64T/AMD64) | None |

#1

VERITAS Cluster Server has a DiskReservation resource that reserves disks, but an environment that combines VERITAS Cluster Server using the DiskReservation resource and HDLM is not supported.

#2

The version of this program is not supported with kernels 2.4.21-27.0.1.EL or later.

#3

The version of this program is not supported with kernels 2.4.21-9.0.1.EL, 2.4.21-9.0.1.ELsmp, and 2.4.21-9.0.1.ELhugemem.

#4

The version of this program is not supported with kernels 2.4.21-37.EL or later.

#5

The version of this program is not supported with kernels 2.4.21-32.0.1.EL or later.

#6

The configuration for Oracle Cluster File System (1.0.14-1) that uses a raw device and the configuration for Oracle Cluster Management Software are supported with kernel 2.4.21-32.0.1.EL, 2.4.21-32.0.1.ELsmp, and 2.4.21-32.0.1.ELhugemem.

#7

The configuration for Oracle Cluster File System (1.0.14-1) that uses a raw device and the configuration for Oracle Cluster Management Software are supported with kernel 2.4.21-37.EL, 2.4.21-37.ELsmp, and 2.4.21-37.ELhugemem.

#8

The configuration for the Automatic Storage Management (ASM) function that uses a raw device and Cluster Ready Service (CRS), and the configuration for Oracle Cluster File System (1.0.12-1) and Cluster Ready Service (CRS) are supported on Red Hat Enterprise Linux AS3 (2.4.21-9.0.1.ELsmp).

#9

The following configurations are supported:

- The configuration for the Automatic Storage Management (ASM) function that uses a raw device and Cluster Ready Service (CRS) is supported with kernel 2.4.21-20.EL.

- The configuration for Oracle Cluster File System (1.0.14-1) and Cluster Ready Service (CRS) is supported with kernel 2.4.21-27.ELsmp.

- The configuration for the Automatic Storage Management (ASM) function that uses a raw device and Cluster Ready Service (CRS) is supported with kernel 2.4.21-32.0.1.EL, 2.4.21-32.0.1.ELsmp or 2.4.21-32.0.1.ELhugemem.

For details, see the notes in 3.4.2.

#10

The configuration for Oracle Cluster File System (1.0.14-1) and Cluster Ready Service (CRS) are supported with kernels 2.4.21-27.EL, 2.4.21-27.ELsmp, 2.4.21-27.ELhugemem, 2.4.21-32.0.1.EL, 2.4.21-32.0.1.ELsmp, 2.4.21-32.0.1.ELhugemem, 2.4.21-37.EL, 2.4.21-37.ELsmp and 2.4.21-37.ELhugemem.

#11

The configuration for Automatic Storage Management (ASM) function that uses a raw device and the configuration of Cluster Ready Service (CRS) are supported with kernel 2.4.21-27.EL, 2.4.21-27.ELsmp, 2.4.21-27.ELhugemem, 2.4.21-27.0.1.EL, 2.4.21-27.0.1.ELsmp, 2.4.21-27.0.1.ELhugemem, 2.4.21-27.0.2.EL, 2.4.21-27.0.2.ELsmp, 2.4.21-27.0.2.ELhugemem, 2.4.21-32.0.1.EL, 2.4.21-32.0.1.ELsmp, 2.4.21-32.0.1.ELhugemem, 2.4.21-37.EL, 2.4.21-37.ELsmp or 2.4.21-37.ELhugemem.

For details, see the description for notes common to OSs in 3.4.2.

#12

The configuration for Automatic Storage Management (ASM) function that uses a raw device and the configuration of Cluster Ready Service (CRS) are supported with kernel 2.4.21-27.EL or later.

For details, see the description for notes common to OSs in 3.4.2.

#13

The version of this program is supported only with kernels 2.4.21-27.EL, 2.4.21-27.ELsmp, and 2.4.21-27.ELhugemem. This program is also supported in an environment where one of these kernels is installed on a system that uses an AMD Opteron CPU.

#14

The configuration for Oracle Cluster File System (1.0.14-1) that uses a raw device and the configuration for Oracle Cluster Management Software are supported with kernel 2.4.21-37.EL.

#15

The version of this program is only supported with kernels 2.4.21-20.EL or 2.4.21-27.EL.

#16

The configuration for Oracle Cluster File System (1.0.14-1) that uses a raw device and the configuration for Oracle Cluster Management Software are supported with kernel 2.4.21-37.EL or 2.4.21-37.ELsmp.

#17

The configuration that uses a raw device and the configuration for Oracle Cluster Management Software are supported with kernel 2.6.9-11.EL, 2.6.9-11.ELsmp, 2.6.9-11.ELhugemem, 2.6.9-22.EL, 2.6.9-22.ELsmp, 2.6.9-22.ELhugemem, 2.6.9-34.EL, 2.6.9-34.ELsmp, 2.6.9-34.ELhugemem, 2.6.9-42.EL, 2.6.9-42.ELsmp, or 2.6.9-42.ELhugemem.

#18

The following configurations are supported:

- The configuration for the Automatic Storage Management (ASM) function that uses ASMLib Kernel Driver and Cluster Ready Service (CRS) is supported with kernel 2.6.9-11.ELsmp.
- The configuration for the Automatic Storage Management (ASM) function that uses a raw device and Cluster Ready Service (CRS) is supported with kernel 2.6.9-11.EL, 2.6.9-11.ELsmp, 2.6.9-11.ELhugemem, 2.6.9-22.EL, 2.6.9-22.ELsmp.

#19

The version of this program is not supported with Red Hat Enterprise Linux AS4 Update1 (2.6.9-11.EL).

#20

The configuration for Automatic Storage Management (ASM) function that uses a raw device and the configuration of Cluster Ready Service (CRS) are supported with kernel 2.6.9-11.EL, 2.6.9-22.EL, 2.6.9-34.EL, 2.6.9-42.EL.

For details, see the description for notes common to OSs in 3.4.2.

#21

The configuration for Automatic Storage Management (ASM) function that uses a raw device and the configuration of Cluster Ready Service (CRS) are supported with kernel 2.6.9-11.EL, 2.6.9-11.ELsmp, 2.6.9-22.EL, 2.6.9-22.ELsmp.

For details, see the description for notes common to OSs in 3.4.2.

#22

The configuration for Automatic Storage Management (ASM) function that uses a raw device and the configuration of Cluster Ready Service (CRS) are supported with kernel 2.6.5-7.244-default or 2.6.5-7.244-smp.

For details, see the description for notes common to OSs in 3.4.2.

### 3.1.3.2 When Using a Volume Manager

Table 3.5 list and describes the related programs when using a volume manager.

**Table 3.5 Related Programs When a Volume Manager Is Used**

| OS | Kernel | Related programs |
|---|---|---|
| Red Hat Enterprise Linux AS 3 (IA32) | 2.4.21-9.0.1.EL<br>2.4.21-9.0.1.ELsmp<br>2.4.21-9.0.1.ELhugemem<br>2.4.21-15.EL<br>2.4.21-15.ELsmp<br>2.4.21-15.ELhugemem<br>2.4.21-20.EL<br>2.4.21-20.ELsmp<br>2.4.21-20.ELhugemem | LVM 1.0.3-15 |
| | 2.4.21-27.EL<br>2.4.21-27.ELsmp<br>2.4.21-27.ELhugemem<br>2.4.21-27.0.1.EL<br>2.4.21-27.0.1.ELsmp<br>2.4.21-27.0.1.ELhugemem<br>2.4.21-27.0.2.EL<br>2.4.21-27.0.2.ELsmp<br>2.4.21-27.0.2.ELhugemem | LVM 1.0.8-9 |
| | 2.4.21-32.0.1.EL<br>2.4.21-32.0.1.ELsmp<br>2.4.21-32.0.1.ELhugemem | LVM 1.0.8-12.2 |
| | 2.4.21-37.EL<br>2.4.21-37.ELsmp<br>2.4.21-37.ELhugemem<br>2.4.21-47.EL<br>2.4.21-47.ELsmp<br>2.4.21-47.ELhugemem | LVM 1.0.8-14 |
| Red Hat Enterprise Linux AS 3 (IPF) | 2.4.21-9.0.1.EL<br>2.4.21-15.EL<br>2.4.21-20.EL | LVM 1.0.3-15 |
| | 2.4.21-27.EL<br>2.4.21-27.0.1.EL<br>2.4.21-27.0.2.EL | LVM 1.0.8-9 |
| | 2.4.21-32.0.1.EL | LVM 1.0.8-12.2 |
| | 2.4.21-37.EL<br>2.4.21-47.EL | LVM 1.0.8-14 |

| OS | Kernel | Related programs |
|---|---|---|
| Red Hat Enterprise Linux<br><br>ES 3 (IPF) | 2.4.21-20.EL | LVM 1.0.3-15 |
| | 2.4.21-27.EL<br>2.4.21-27.0.1.EL<br>2.4.21-27.0.2.EL | LVM 1.0.8-9 |
| | 2.4.21-32.0.1.EL | LVM 1.0.8-12.2 |
| | 2.4.21-37.EL<br>2.4.21-47.EL | LVM 1.0.8-14 |
| Red Hat Enterprise Linux AS/ES 3 (EM64T) | 2.4.21-20.EL | LVM 1.0.8-5 |
| | 2.4.21-27.EL | LVM 1.0.8-9 |
| | 2.4.21-32.0.1.EL | LVM1.0.8-12.2 |
| | 2.4.21-37.EL<br>2.4.21-47.EL | LVM 1.0.8-14 |
| Red Hat Enterprise Linux AS/ES 3 (AMD64) | 2.4.21-32.0.1.EL<br>2.4.21-32.0.1.ELsmp | LVM1.0.8-12.2 |
| | 2.4.21-37.EL<br>2.4.21-37.ELsmp<br>2.4.21-47.EL<br>2.4.21-47.ELsmp | LVM 1.0.8-14 |
| Red Hat Enterprise Linux AS/ES 4 (IA32) | 2.6.9-11.EL<br>2.6.9-11.ELsmp<br>2.6.9-11.ELhugemem | LVM 2.01.08-1.0.RHEL4 |
| | 2.6.9-22.EL<br>2.6.9-22.ELsmp<br>2.6.9-22.ELhugemem | LVM2.01.14-2.0.RHEL4 |
| | 2.6.9-34.EL<br>2.6.9-34.ELsmp<br>2.6.9-34.ELhugemem | LVM2.02.01-1.3.RHEL4 |
| | 2.6.9-42.EL<br>2.6.9-42.ELsmp<br>2.6.9-42.ELhugemem | LVM2.02.06-6.0.RHEL4 |
| | 2.6.9-42.0.3.EL<br>2.6.9-42.0.3.ELsmp<br>2.6.9-42.0.3.ELhugemem | LVM2.02.06-6.0.RHEL4 |
| Red Hat Enterprise Linux AS/ES 4 (IPF) | 2.6.9-11.EL | LVM 2.01.08-1.0.RHEL4 |
| | 2.6.9-22.EL | LVM2.01.14-2.0.RHEL4 |
| | 2.6.9-34.EL | LVM2.02.01-1.3.RHEL4 |

| OS | Kernel | Related programs |
|---|---|---|
| | 2.6.9-42.EL | LVM2.02.06-6.0.RHEL4 |
| | 2.6.9-42.0.3.EL | LVM2.02.06-6.0.RHEL4 |
| Red Hat Enterprise Linux AS/ES 4 (EM64T/AMD64) | 2.6.9-11.EL<br>2.6.9-11.ELsmp | LVM 2.01.08-1.0.RHEL4 |
| | 2.6.9-22.EL<br>2.6.9-22.ELsmp | LVM2.01.14-2.0.RHEL4 |
| | 2.6.9-34.EL<br>2.6.9-34.ELsmp<br>2.6.9-34.ELlargesmp | LVM2.02.01-1.3.RHEL4 |
| | 2.6.9-42.EL<br>2.6.9-42.ELsmp<br>2.6.9-42.ELlargesmp | LVM2.02.06-6.0.RHEL4 |
| | 2.6.9-42.0.3.EL<br>2.6.9-42.0.3.ELsmp<br>2.6.9-42.0.3.ELlargesmp | LVM2.02.06-6.0.RHEL4 |
| SUSE LINUX Enterprise Server 9 (IA32) | 2.6.5-7.139-default<br>2.6.5-7.139-smp<br>2.6.5-7.139-bigsmp | LVM 2.00.15-0.8 |
| | 2.6.5-7.191-default<br>2.6.5-7.191-smp<br>2.6.5-7.191-bigsmp | LVM 2.00.33-1.3 |
| | 2.6.5-7.201-smp | LVM2.00.33-1.3 |
| | 2.6.5-7.244-default<br>2.6.5-7.244-smp<br>2.6.5-7.244-bigsmp | LVM2.01.14-3.6 |
| SUSE LINUX Enterprise Server 9 (IPF) | 2.6.5-7.139-default<br>2.6.5-7.139-64k-pagesize | LVM 2.00.15-0.8 |
| | 2.6.5-7.191-default<br>2.6.5-7.191-64k-pagesize | LVM 2.00.33-1.3 |
| | 2.6.5-7.244-default<br>2.6.5-7.244-64k-pagesize | LVM2.01.14-3.6 |
| SUSE LINUX Enterprise Server 9 (EM64T/AMD64) | 2.6.5-7.191-default<br>2.6.5-7.191-smp | LVM 2.00.33-1.3 |
| | 2.6.5-7.244-default<br>2.6.5-7.244-smp | LVM2.01.14-3.6 |
| | 2.6.5-7.252-smp | LVM2.01.14-3.6 |

| OS | Kernel | Related programs |
|---|---|---|
| SUSE LINUX Enterprise Server10 (IA32) | 2.6.16.21-0.8-default<br>2.6.16.21-0.8-smp<br>2.6.16.21-0.8-bigsmp | LVM2.02.02-11.2 |
| | 2.6.16.27-0.9-default<br>2.6.16.27-0.9-smp<br>2.6.16.27-0.9-bigsmp | LVM2.02.02-11.2 |
| SUSE LINUX Enterprise Server10 (IPF) | 2.6.16.21-0.8-default | LVM2.02.02-11.2 |
| | 2.6.16.27-0.9-default | LVM2.02.02-11.2 |
| SUSE LINUX Enterprise Server10 (EM64T/AMD64) | 2.6.16.21-0.8-default<br>2.6.16.21-0.8-smp | LVM2.02.02-11.2 |
| | 2.6.16.27-0.9-default<br>2.6.16.27-0.9-smp | LVM 2.02.02-11.2 |

### 3.1.3.3 When Using the File System

Table 3.6 lists and describes the related programs when using the file system.

**Table 3.6 Related Programs When the File System Is Used**

| OS | Related programs |
|---|---|
| Red Hat Enterprise Linux AS/ES 3 (IA32) | ext2 (supplied with the OS) |
| | ext3 (supplied with the OS) |
| | VERITAS File System 4.0# |
| Red Hat Enterprise Linux AS/ES 3 (IPF) | ext2 (supplied with the OS) |
| | ext3 (supplied with the OS) |
| Red Hat Enterprise Linux AS/ES 3 (EM64T) | ext2 (supplied with the OS) |
| | ext3 (supplied with the OS) |
| Red Hat Enterprise Linux AS/ES 3 (AMD64) | ext2 (supplied with the OS) |
| | ext3 (supplied with the OS) |
| Red Hat Enterprise Linux AS/ES 4 | ext2 (supplied with the OS) |
| | ext3 (supplied with the OS) |
| SUSE LINUX Enterprise Server 9 | ext2 (supplied with the OS) |
| | ext3 (supplied with the OS) |
| | ReiserFS (supplied with the OS) |
| SUSE LINUX Enterprise Server 10 | ext2 (supplied with the OS) |

| OS | Related programs |
| --- | --- |
| | ext3 (supplied with the OS) |
| | ReiserFS (supplied with the OS) |

#

- VERITAS File System 4.0 is the file system provided with VERITAS Storage Foundation 4.0.The GUI functions provided by VERITAS Storage Foundation 4.0 are not supported.

- This file system is not supported with kernels 2.4.21-9.0.1.EL, 2.4.21-9.0.1.ELsmp, and 2.4.21-9.0.1.ELhugemem.

- This file system is not supported in an environment that also includes BladeSymphony.

### 3.1.4 Memory and Disk Space Requirements

This section describes memory and disk space requirements.

#### 3.1.4.1 Memory Requirements

Table 3.7 lists the memory requirements for a host.

**Table 3.7 Memory Requirements for a Host**

| OS | Required memory |
| --- | --- |
| Red Hat Enterprise Linux AS4/ES4, SUSE Linux Enterprise Server9, SUSE Linux Enterprise Server 10 | 26 MB |
| Red Hat Enterprise Linux AS3/ES3 | 9 MB |

#### 3.1.4.2 Disk Requirements

This section describes disk space requirements for a host.

Table 3.8 lists the disk space requirements for a host.

**Table 3.8 Disk Space Requirements for a Host**

| Directory | Disk space requirements |
| --- | --- |
| `/tmp` | 570 KB |
| `/var` | p MB[#1] + q MB[#2] + 200 KB + 10 MB |
| `/opt` | 168 MB |
| `/etc` | 3 KB + (size of the settings file)[#3] |
| `/root` | 580 KB |

#1

This size depends on the settings for the log files. The maximum size is 30 GB.

When s is the error log file size (units: KB, the default value is 9900) and m is the number of error log files (the default value is 2), this value (p) can be calculated as follows: p = (s x m) / 1024 (rounded-up integer) (units: MB). If there is a remainder, the result is rounded up to the nearest MB.

#2

This size depends on the setting for the trace files. The maximum size is 1000 MB.

When t is the trace file size (units: KB, the default value is 1000) and n is the number of trace files (the default value is 4), this value (q) can be calculated as follows:

q = (t x n) / 1024 (rounded-up integer) (units: MB). If there is a remainder, the result is rounded up to the nearest MB.

#3

The size of the settings file depends on the size of the configuration definition file and the size of the HDLM-unmanaged disk definition file. The equation below gives an idea of the size:

Size of the settings file = Configuration definition file size (620Bytes x number of the Lightning 9900 Series, the Lightning 9900V Series, TagmaStore USP, or Universal Storage Platform V connection path + 1.2 KB x number of the Thunder 9200 and the Thunder 9500V Series and the TagmaStore AMS/WMS series connection paths)

+ HDLM-unmanaged disk definition file (16Bytes x number of excluded disks)

### 3.1.5 Number of Paths Guaranteed in HDLM

Table 3.9 lists the number of LUs, number of paths per LU, and total number of paths guaranteed in HDLM.

**Table 3.9 Number of Paths Guaranteed in HDLM**

| Target | | Maximum value |
|---|---|---|
| Number of paths per LU | Red Hat Enterprise Linux AS3/ES3 or Red Hat Enterprise Linux AS4/ES4 | 64 |
| | SUSE Linux Enterprise Server 9 | 32 |
| | SUSE Linux Enterprise Server 10 | 32 |
| Maximum number of LUs from 1 port for the Lightning 9900 Series or the Lightning 9900V Series | | 64 |
| Maximum number of LUs for the Thunder 9200 or the Thunder 9500V Series | | 64 |
| Maximum number of LUs from 1 ports for the TagmaStore USP#1 | | 64 |
| Maximum number of LUs from 1 ports for the TagmaStore AMS | | 64 |
| Maximum number of LUs from 1 ports for the TagmaStore WMS | | 64 |
| Maximum number of LUs from 1 ports for the Universal Storage Platform V#1 | | 64 |

| | | |
|---|---|---|
| Total number of paths (Number of LU x Number of paths per LU) | Red Hat Enterprise Linux AS3/ES3 | 128#2 |
| | Red Hat Enterprise Linux AS4/ES4 | 128#2 |
| | SUSE LINUX Enterprise Server 9 | 128#2 |
| | SUSE LINUX Enterprise Server 10 | 128#2 |

#1

In TagmaStore USP or in Universal Storage Platform V, LUN 0 to 1023 can be assigned but the range supported by HDLM is from 0 to 255. Therefore, HDLM cannot recognize LUs from 256 to 1023.

#2

The maximum value is the total number of block devices in use on a given host.

If there are devices in use (such as internal disks and FD devices) other than those managed by HDLM, the number of available paths may be less than the total number of paths.

### 3.1.6 When HDLM Is Used in a BladeSymphony Environment

HDLM supports BladeSymphony in any of the environments shown in Table 3.10. For details on OSs (kernels) that each BladeSymphony model supports, see the BladeSymphony documentation.

**Table 3.10 Operating Environments for BladeSymphony When Red Hat Enterprise Linux AS4/ES4 Is Used**

| OS | Kernal | HBA | storage subsystems |
|---|---|---|---|
| Red Hat Enterprise Linux AS4/ES4 (IA32) | 2.6.9-11.ELsmp | hitachi | Thunder 9500V Series<br>Lightning 9900V Series#<br>TagmaStore USP<br>TagmaStore AMS<br>Universal Storage Platform V |
| | 2.6.9-34.ELsmp | | |
| Red Hat Enterprise Linux AS4/ES4 (IPF) | 2.6.9-11.EL | | |
| | 2.6.9-34.EL | | |
| Red Hat Enterprise Linux AS4/ES4 (EM64T) | 2.6.9-11.ELsmp | | |
| | 2.6.9-34.ELsmp | | |

#

The storage subsystem is not supported in an environment where an HDLM device is specified for the boot disk.

When HDLM is used in a BladeSymphony environment, the following functions cannot be used:

- HotPlug
- SystemManager linkage

***Note:*** HDLM cannot use the functions provided by ServerConductor/DeploymentManager other than the following:

- Backup and restoration

- BIOS/firmware updating

For details on the functions provided by ServerConductor/DeploymentManager, see the manuals *ServerConductor/DeploymentManager User's Guide*.

### 3.1.7 The Operation Environment When Using SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10

This subsection describes the operating environment common to SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10.

- Among the functions for SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10, HDLM only supports Complete Fair Queuing (CFQ), and the default I/O scheduler functionality.
- An HDLM device that applies EVMS (Enterprise Volume Manager) functions is not supported.
- You cannot use DRBD (Distributed Replicated Block Device) functions in an environment where HDLM is installed.
- You cannot use HDLM in a User-Mode Linux environment.
- Although disk partitioning and system backup functions are supported for disks in a YaST2 system, LVM functions are not supported.
- When you execute the `pvscan`, `vgscan`, `pvcreate`, or `lvcreate` command, a CD-ROM IOCTL command is sent to each device in the OS. At this time, a CD-ROM IOCTL command is also sent to HDLM devices. HDLM outputs an error message (KAPL05023-E) to syslog because this IOCTL command is not for SCSI devices in HDLM. However, HDLM operations are not affected.

  This message is also output in the following cases, but HDLM operations are not affected:
  - You execute the Utility for Collecting HDLM Error Information (`DLMgetras`).
  - The system is started.

### 3.1.8 Using a Boot Disk Support Utility (dlmmkinitrd)

The boot disk support utility (`dlmmkinitrd`) supports the operation that uses the SAN environments shown below to start the system.

**Table 3.11 Operating Environments for BladeSymphony When Red Hat Enterprise Linux AS4 Is Used**

| OS | Kernal | HBA | storage subsystems |
|---|---|---|---|
| Red Hat Enterprise Linux AS4/ES4 (IA32) | 2.6.9-11.ELsmp | hitachi | Thunder 9500V Series<br>TagmaStore USP<br>TagmaStore AMS<br>Universal Storage Platform V |
| | 2.6.9-34.ELsmp | | |
| Red Hat Enterprise Linux AS4/ES4 (IPF) | 2.6.9-11.EL | | |
| | 2.6.9-34.EL | | |
| Red Hat Enterprise Linux AS3/ES3 | 2.6.9-11.ELsmp | | |

| OS | Kernal | HBA | storage subsystems |
|---|---|---|---|
| (EM64T) | 2.6.9-34.ELsmp | | |

**Table 3.12 Operating Environments for Hosts That Can Be Booted from a Storage Subsystem on Red Hat Enterprise Linux AS3/ES3**

| OS | Kernal | HBA | storage subsystems |
|---|---|---|---|
| Red Hat Enterprise Linux AS3/ES3 (IA32) | 2.4.21-20.EL<br>2.4.21-20.ELsmp<br>2.4.21-20.ELhugemem | QLogic | Thunder 9200 Series<br>Thunder 9500V Series<br>Lightning 9900V Series<br>TagmaStore USP<br>TagmaStore AMS/WMS Series<br>Universal Storage Platform V |
| | 2.4.21-27.EL<br>2.4.21-27.ELsmp<br>2.4.21-27.ELhugemem | | |
| | 2.4.21-32.0.1.EL<br>2.4.21-32.0.1.ELsmp<br>2.4.21-32.0.1.ELhugemem | | |
| | 2.4.21-37.EL<br>2.4.21-37.ELsmp<br>2.4.21-37.ELhugemem | | |
| | 2.4.21-47.EL<br>2.4.21-47.ELsmp<br>2.4.21-47.ELhugemem | | |
| Red Hat Enterprise Linux AS3/ES3 (IPF) | 2.4.21-20.EL | Emulex | |
| | 2.4.21-27.EL | | |
| | 2.4.21-32.0.1.EL | | |
| | 2.4.21-37.EL | | |
| | 2.4.21-47.EL | | |
| Red Hat Enterprise Linux AS3/ES3 (EM64T) | 2.4.21-20.EL | QLogic | |
| | 2.4.21-27.EL | | |
| | 2.4.21-32.0.1.EL | | |
| | 2.4.21-37.EL | | |
| | 2.4.21-47.EL | | |
| Red Hat Enterprise Linux AS3/ES3 (AMD64) | 2.4.21-32.0.1.EL<br>2.4.21-32.0.1.ELsmp | QLogic | |
| | 2.4.21-37.EL<br>2.4.21-37.ELsmp | | |
| | 2.4.21-47.EL<br>2.4.21-47.ELsmp | | |

**Table 3.13 Operating Environments for Hosts That Can Be Booted from a Storage Subsystem on Red Hat Enterprise Linux AS4/ES4**

| OS | Kernal | HBA | storage subsystems |
|---|---|---|---|
| Red Hat Enterprise Linux AS4/ES4(IA32) | 2.6.9-11.EL<br>2.6.9-11.ELsmp<br>2.6.9-11.ELhugemem | QLogic | Thunder 9200 Series<br>Thunder 9500V Series<br>Lightning 9900V Series<br>TagmaStore USP<br>TagmaStore AMS/WMS Series<br>Universal Storage Platform V |
| | 2.6.9-22.EL<br>2.6.9-22.ELsmp<br>2.6.9-22.ELhugemem | | |
| | 2.6.9-34.EL<br>2.6.9-34.ELsmp<br>2.6.9-34.ELhugemem | | |
| | 2.6.9-42.EL<br>2.6.9-42.ELsmp<br>2.6.9-42.ELhugemem | | |
| | 2.6.9-42.0.3.EL<br>2.6.9-42.0.3.ELsmp<br>2.6.9-42.0.3.ELhugemem | | |
| Red Hat Enterprise Linux AS4/ES4(EM64T/ADM64) | 2.6.9-11.EL<br>2.6.9-11.ELsmp | QLogic | |
| | 2.6.9-22.EL<br>2.6.9-22.ELsmp | | |
| | 2.6.9-34.EL<br>2.6.9-34.ELsmp<br>2.6.9-34.ELlargesmp | | |
| | 2.6.9-42.EL<br>2.6.9-42.ELsmp<br>2.6.9-42.ELlargesmp | | |
| | 2.6.9-42.0.3.EL<br>2.6.9-42.0.3.ELsmp<br>2.6.9-42.0.3.ELlargesmp | | |

## 3.2 Flow for Creating an HDLM Environment

Set up the environment to use HDLM as shown in the following diagram.

If you are going to use volume management software, the procedure for building a file system shown in the figure is not required.

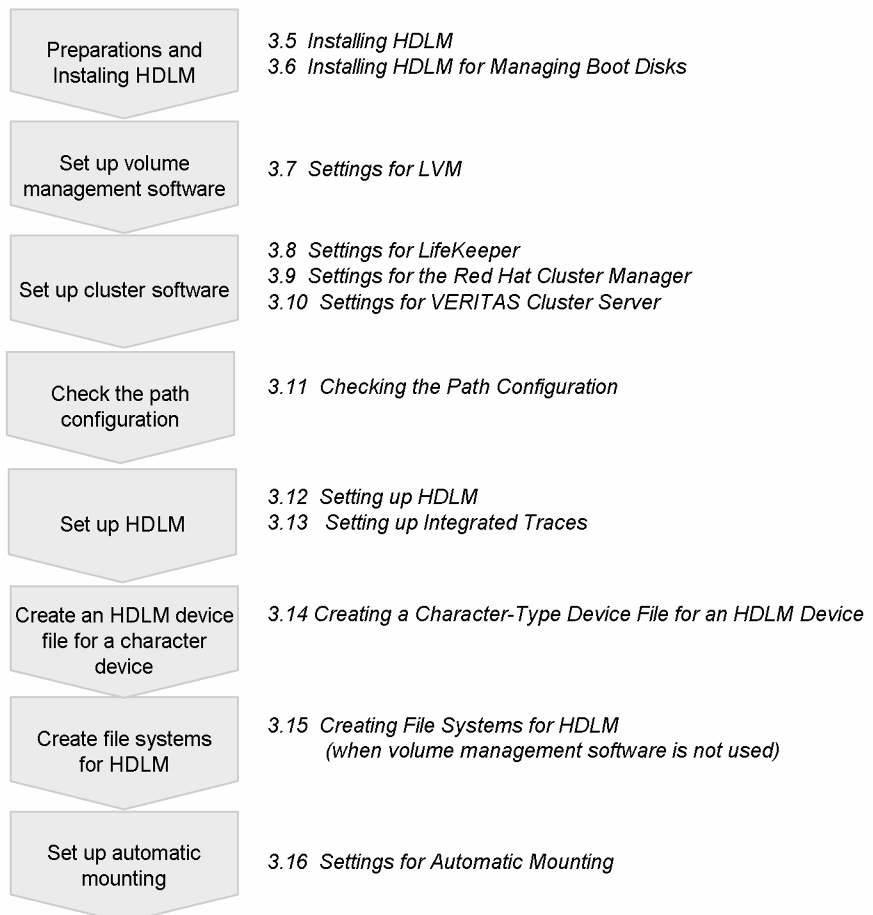


**Figure 3.1 Flow of HDLM Environment Setup**

## 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 new version of HDLM without uninstalling the existing old version is called an *upgrade installation of HDLM*.

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

Be sure to note the following when installing HDLM.

### 3.4.1 Notes on Hardware Settings

- The types of host bus adapter (HBA) that are installed on one host can be mixed. However, the HBAs connected to an HDLM-managed device must all be of the same type. Also make sure that the HBA driver versions are the same. If the types or versions of HBAs connected to the devices for the HDLM-managed SCSI disks vary, you cannot switch the path when an error occurs.
- If you use version 8.00.00, 8.00.00b21-k, or 8.00.01 of the QLogic HBA driver, a path recovered after an error might be unable to be placed in the `Online` status, such as by an `Online` operation or by automatic failback. In this case, take the following procedure:

  a) Execute the following command:

  `# echo "scsi-qlascan" > /proc/scsi/`*driver-name*`/`*host-port-number*

  *driver-name*: *qla2xxx* (the name of the driver being used)

  When the path does not recover even if this command is executed, restart the system.
- You must not change vendor ID and product ID of the storage subsystem.

  If you change these IDs, HDLM cannot recognize the storage subsystem.
- Make sure that you start the storage subsystem before starting the host.
- When a device is added to a PCI bus, and the IRQ is shared with the HBA and other devices, a problem may occur. When adding a device, make sure that the IRQ is not shared with other devices.
- In an environment where a host is connected to a storage subsystem by using a switch, if the host starts while the path between the switch and the storage subsystem is disconnected, the disconnected path might not be recoverable while the host is running because the target ID of the path before the host restarts is different from the one after the host restarts. Also, when you execute the HDLM command (dlnkmgr) to display the path status, the target ID of the path before the change might be displayed or the disconnected path might not be displayed. In such a case, you can use the target path normally when the same configuration where the disconnected path was once operated normally is physically restored and then the host is restarted. Check the HDLM device name after the host is restarted because the HDLM device name might have been changed. If the HDLM device name has been changed, reconfigure the above program.

  However, you can avoid the above situation by fixing the target ID for each LU with the HBA driver persistent binding functionality. For details on whether the persistent binding functionality is supported for your HBA driver and the setting method, see the manual provided with the HDLM driver.

- In an environment where a host is connected to storage subsystems by using a switch, if the zone setting for the switch has changed, the target ID of the path might have changed. When you execute the HDLM command (dlnkmgr) to display the path status, the target ID of the path before the change might be displayed and also the HDLM device name might be changed after the host is restarted. Check the HDLM device name after the host is restarted. If the HDLM device name has changed, reconfigure the above program.

  However, you can avoid the above situation by fixing the target ID for each LU with the HBA driver persistent binding functionality. For details on whether the persistent binding functionality is supported for your HBA driver and the setting method, see the manual provided with the HDLM driver.

- To use a QLogic HBA driver under Red Hat Enterprise Linux AS3/ES3, both of the modules qla2300.o and qla2300_conf.o (or qla2200.o and qla2200_conf.o) must be loaded into the kernel.
- When using SUSE LINUX Enterprise Server 9, SUSE LINUX Enterprise Server 10, or Red Hat Enterprise Linux AS4/ES4, you must use `modprobe` to load the filter driver.

### 3.4.2 Notes on Installing HDLM

- If you created an environment by using the QLogic HBA driver while SUSE LINUX Enterprise Server 9 is being used, before installing HDLM, make sure that the following coding is in the line for `INITRD_MODULES=` in the `/etc/sysconfig/kernel` file.
  - When using the QLA22XX HBA:

    `qla2200`
  - When using the QLA23XX HBA:

    `qla2300`

  If `qla2xxx` is appears in the file instead of `qla2200` or `qla2300`, change `qla2xxx` to `qla2200` or `qla2300` as follows:

  ```
  INITRD_MODULES="mptscsih qla2300 reiserfs"
  ```

  After editing the `/etc/sysconfig/kernel` file, execute the `mkinitrd` command to re-create an initial RAM disk image.

- Create the `/etc/opt` directory on the host where HDLM is to be installed, in the same partition as the root directory (`/`).
- Note that if the `/var` directory is set up on a different partition than the root directory (`/`), no logs will be output to the `/var/log/messages` file until the `/var` directory has been mounted during the OS start process.
- If multiple OSs exist on a disk that differs physically or logically within the same host, HDLM can be installed on each OS. Also, if multiple kernels exist on one OS, HDLM can be used only when the kernel installed with HDLM is started.
- You can install HDLM on the following disks or devices:
  - An internal disk on a host machine
  - LUs of storage subsystems that are used as a boot disk

For details about the storage subsystems on which HDLM can be installed, see section 3.1.8.

- Although the following messages are output when you load the alert driver and the filter driver respectively, HDLM operations are not affected.

```
sddlmadrv: unsupported module, tainting kernel.
sddlmadrv: module license 'Proprietary: Contact Hitachi for any bugs' taints kernel.
```

- HDLM outputs the log information for analyzing errors at the time of installation to a file. The file name, which is the output destination, is `/var/tmp/hdlminstlog/installhdlm[01-10].log`, and the maximum number of files to be created is 10. If 10 files already exist and you perform another installation, the oldest file is deleted and a new log file is output.

- The following information is inherited from the previous environment even when HDLM is installed in an upgrade installation or re-installation.
  - Information for defining the configuration of an HDLM device
  - Information for setting HDLM functions
  - Log information
  - Driver configuration

- When you begin an upgrade installation, re-installation, or uninstallation of HDLM, you must not execute the `dlmlfk` utility for supporting LifeKeeper from the command line or a custom script.

- If you perform a re-installation, upgrade installation, or uninstallation of HDLM in a Red Hat Enterprise Linux AS3/ES3 environment that uses LifeKeeper, you must stop LifeKeeper before executing that operation. After performing an upgrade installation or re-installation of HDLM in an environment that uses LifeKeeper, execute the `dlmlfk` utility (shown below).

```
# /opt/DynamicLinkManager/bin/dlmlfk -c
```

- A license key is required in the following cases:
  - When you install HDLM for the first time
  - When the license expires and then you perform an upgrade installation of HDLM

- If you install HDLM for the first time, or perform an upgrade installation of HDLM when the license expires, a license key is necessary. To update the HDLM license, execute the `dlnkmgr` command's `set -lic` operation. The expiration time of the license key is determined by the license key specified in the license key file or the input license key type. For information on license key types and `set` operation, see section 6.6.

- Set the storage subsystem, switches, and host (including HBA) before installing HDLM, and make sure that the host and storage subsystem are connected.

- Do not activate the same volume group concurrently from multiple servers.

- In a Red Hat Enterprise Linux AS3 (IPF) environment, when one of the following conditions is satisfied, the following message might be output to `/var/log/messages` file, however operations are not affected:

  `set personality to 8` or `IA32 syscall #252 issued, maybe we should implement it`

  The message might be output at the following times:

When HDLM is installed.

When the license information is updated by using the `dlnkmgr` command's `set` operation.

When the license information is viewed by using the `dlnkmgr` command's `view -sys` operation.

- The HDLM manager is stopped during an upgrade installation. Therefore, when you perform an upgrade installation, stop any applications that require HDLM manager operations, such as a log output operation.
  - Since the HDLM manager is stopped during an upgrade installation, an error log will not be output. Also, you will not be able to set HDLM functions. After the upgrade installation finishes successfully, immediately restart the host.
  - Do not cancel an upgrade installation.
- On a server that has an AMD 64-bit CPU, HDLM does not support 32-bit distributions other than the following kernel packages:
  - kernel-2.4.21-27.EL.i686.rpm
  - kernel-smp-2.4.21-27.EL.i686.rpm
  - kernel-hugemem-2.4.21-27.EL.i686.rpm
  - kernel-smp-2.6.9-42.EL.i686.rpm

  If a 32-bit distribution of a non-supported kernel package is used on a server that has an AMD 64-bit CPU, when you install HDLM, the KAPL09160-E message might be output, and the installation might fail.

  To check the kernel architecture and the CPU vendor:

  a) Execute the following command to check which kernel architecture is used:

```
# uname -m
i686
#
```

  The following shows the meaning of the execution result of the `uname` command:

  `i686`: IA32 architecture

  `ia64`: IPF architecture

  `x86_64`: AMD64/EM64T architecture

  b) Execute the following command to check the vendor of the CPU you are using:

```
# cat /proc/cpuinfo
processor       : 0
vendor_id       : AuthenticAMD
cpu family      : 15
model           : 37
model name      : AMD Opteron(tm) Processor 252
stepping        : 1
    :
```

```
:
#
```

Check the `vendor_id` line. `AuthenticAMD` is displayed for AMD CPUs, and `GenuineIntel` is displayed for Intel CPUs.

- If Red Hat Enterprise Linux AS3/ES3 is being used, the following problems might occur when you install HDLM (when you execute the `installhdlm` command) because of a problem in the `rpm` command:
  - When executing the `installhdlm` command without the `-s` option specified

    After the KAPL09093-I message is displayed, if you enter `y`, the KAPL09098-E message will be output, and the installation will terminate abnormally.
  - When executing the `installhdlm` command with the `-s` option specified

    Only the KAPL09098-E message will be output, and the installation will terminate abnormally.

  If the above problems occur, check the version of the rpm package by executing the following command:

```
# rpm -q rpm
rpm-4.2.2-0.14
```

  If the rpm package version is 4.2.1-4.4 or earlier, upgrade the rpm package to version 4.2.2-0.14 or later.

- If Red Hat Enterprise Linux AS3/ES3 or Red Hat Enterprise Linux AS4/ES4 is used and an HDLM device is not used as a boot disk, after you upgrade the initscripts package (when HDLM has been installed) and then restart the OS, you need to upgrade or re-install HDLM if all of the following conditions are satisfied:
  - There is no HDLM device file (Red Hat Enterprise Linux AS4/ES4 only):

    Confirm that there is no HDLM device file by executing the following command:

    ```
    # ls /dev/sddlm*
    ```

    If the following message is displayed, it indicates that there is no HDLM device file:

    ```
    ls: /dev/sddlm*: No such file or directory
    ```
  - The HDLM driver has not been loaded

    Confirm the load status of the HDLM driver by executing the following command:

    ```
    # lsmod | grep sddlmfdrv
    ```

    If nothing is output after command execution, this indicates that no HDLM driver has been loaded.
  - HDLM information is missing in either the `/etc/rc.d/rc.sysinit` file or the `/etc/init.d/halt` file

    Confirm that the HDLM information is missing by executing the following commands:

    ```
    # grep DLM /etc/rc.d/rc.sysinit
    # grep DLM /etc/init.d/halt
    ```

    If nothing is output after command execution, this indicates that the HDLM information is missing.

- If Red Hat Enterprise Linux AS3/ES3 or Red Hat Enterprise Linux AS4/ES4 is used and an HDLM device is used as a boot disk, after you upgrade the initscripts package (when HDLM has been installed) and then restart the OS, you need to upgrade or re-install HDLM if both of the following conditions are satisfied:
  - The HDLM manager is not running

    Confirm that the HDLM manager is not running by executing the command below.

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

    If `Dead` is output as follows, this indicates that the HDLM manager is not running.

    ```
    HDLM Manager Ver      WakeupTime
    Dead
    ```

    **KAPL01001-I The HDLM command completed normally. Operation name = view, completion time =** ***yyyy/mm/dd hh:mm:ss***
  - When HDLM information is missing in either the `/etc/rc.d/rc.sysinit` file or the `/etc/init.d/halt` file

    Confirm that the HDLM information is missing by executing the following commands:

    ```
    # grep DLM /etc/rc.d/rc.sysinit
    # grep DLM /etc/init.d/halt
    ```

    If nothing is output after command execution, this indicates that the HDLM information is missing.

  For details about upgrade installation or re-installation in an environment in which an HDLM device is used as a boot disk, see section 3.6.4.
- When you use Oracle Cluster file System in an Oracle RAC environment, set `0` to the `comm_voting` parameter in the `/etc/ocfs.conf` file.
- 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.

- 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
```

## 3.5 Installing HDLM

The following section describes how to install HDLM for a new installation, for a re-installation installation, and for an upgrade installation. Depending on the procedure, the host may or may not have to be restarted after HDLM is installed. Select the procedure suited for your environment.

When you install HDLM in a boot disk environment that uses SCSI devices, see section 3.6.

For details on the `dlmcfgmgr` utility, see section 7.3.

### 3.5.1 Preparations for a New Installation of HDLM

This section describes the preparations required before newly installing HDLM. This includes backing up devices to be managed by HDLM, and setting up the hardware and volume management software.

#### 3.5.1.1 Operations Required for Devices to Be Managed by HDLM

To perform operations required for disks to be managed by HDLM:

1. Terminate all processes of applications that are accessing devices to be managed by HDLM.
2. If necessary, back up all devices that are to be managed by HDLM.
3. Unregister the devices.

   If a device to be managed by HDLM is registered in a program, including cluster software (excluding the volume management software), release the registration.

   When accessing a device that is to be managed after HDLM installation, the setting name used until now cannot be used for access because the logical device file name for the HDLM device that HDLM creates is used.
4. Unmount the disks.
   - If the disks to be managed by HDLM were mounted by specifying SCSI devices, unmount them.

     First, check the current settings. Execute the following command:

     ```
     # mount
     ```

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

```
# mount
/dev/hda5 on / type ext3 (rw)
/dev/sda on /mntpt type ext3 (rw)
none on /proc type proc (rw)
usbdevfs on /proc/bus/usb type usbdevfs (rw)
/dev/hda1 on /boot type ext3 (rw)
none on /dev/pts type devpts (rw,gid=5,mode=620)
none on /dev/shm type tmpfs (rw)
none on /proc/sys/fs/binfmt_misc type binfmt_misc (rw)
#
```

**Figure 3.2 Execution Results of the mount Command**

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

```
# umount /mntpt
```

5. If the disks are set to be mounted automatically when the host starts, delete this setting in the /etc/fstab file.

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

```
LABEL=/1          /              ext3     defaults                1 1
/dev/hda1         /boot          ext3     defaults                1 2
/dev/sda          /mntpt         ext3     defaults,noauto         0 0
none              /dev/pts       devpts   gid=5,mode=620          0 0
none              /proc          proc     defaults                0 0
none              /dev/shm       tmpfs    defaults                0 0
/dev/hda3         swap           swap     defaults                0 0
/dev/cdrom        /mnt/cdrom     iso9660  noauto,owner,kudzu,ro   0 0
/dev/fd0          /mnt/floppy    auto     noauto,owner,kudzu      0 0
```

**Figure 3.3 Example of How to Edit the /etc/fstab File**

Comment out the shaded portions by placing a hash mark (#) at the beginning of each line. The Linux functionality that adds LABEL= to a SCSI device is not supported in HDLM. Do not use this functionality.

### 3.5.1.2 Checking the Volume Group

If you have already created a physical volume, volume group, or logical volume by using LVM, you can use the procedure shown in 3.7 only when all of the conditions below are satisfied. Migration is not affected even when a logical volume or file system has been created.

- A physical volume is created for only one of the logical device files on any one path for each SCSI device to be managed by HDLM. In addition, a volume group is created for only the physical volume.

  This subsection describes how to check whether this condition exists.
- The logical volume is unmounted.

To check the physical volumes that belong to the volume group `vg02`, execute the following command:

```
# vgdisplay -v
```

The following figures show examples of command execution when the volume group consists of only one physical volume (condition satisfied) and when the volume group consists of two physical volumes (condition not satisfied).

```
# vgdisplay -v
--- Volume group ---
VG Name               vg02
VG Access             read/write
VG Status             available/resizable
VG #                  2
MAX LV                256
Cur LV                0
Open LV               0
MAX LV Size           255.99 GB
Max PV                256
Cur PV                1
Act PV                1
VG Size               2.29 GB
PE Size               4 MB
Total PE              585
Alloc PE / Size       0 / 0
Free  PE / Size       585 / 2.29 GB
VG UUID               SCaKcF-17i2-0jxy-m2Xw-YIxj-XE1h-WuyCkO

--- No logical volumes defined in "vg02" ---


--- Physical volumes ---
PV Name (#)           /dev/sde (1)
PV Status             available / allocatable
Total PE / Free PE    585 / 585
```

**Figure 3.4 Result of Executing vgdisplay -v (When There Is One Physical Volume)**

Check the shaded portion and confirm that `vg02` consists of `/dev/sde`.

```
# vgdisplay -v
--- Volume group ---
VG Name             vg02
VG Access           read/write
VG Status           available/resizable
VG #                2
MAX LV              256
Cur LV              0
Open LV             0
MAX LV Size         255.99 GB
Max PV              256
Cur PV              2
Act PV              2
VG Size             4.57 GB
PE Size             4 MB
Total PE            1170
Alloc PE / Size     0 / 0
Free  PE / Size     1170 / 4.57 GB
VG UUID

--- No logical volumes defined in "vg02" ---


--- Physical volumes ---
PV Name (#)         /dev/sdu (2)
PV Status           available / allocatable
Total PE / Free PE 585 / 585

PV Name (#)         /dev/sde (1)
PV Status           available / allocatable
Total PE / Free PE 585 / 585
```

**Figure 3.5 Result of Executing vgdisplay -v (When There Are Two Physical Volumes)**

Check the shaded portion and confirm that `vg02` consists of `/dev/sde` and `/dev/sdu`, both of which have been defined for the same device to be managed by HDLM.

### 3.5.1.3 Setting in the syslogd settings File

The log upon executing the installation is output to syslog. In Red Hat Enterprise Linux, check the syslogd settings file (system log configuration definition file) and make sure that it specifies that messages with an Error level or higher are to be output.

Make sure that the file contains the code shown in Figure 3.6.

```
# cat /etc/syslog.conf | grep /var/log/messages
*.info;mail.none;news.none;authpriv.none;cron.none             /var/log/messages
#
```

**Figure 3.6 Example of the Contents of the syslogd settings File (In Red Hat Enterprise Linux)**

### 3.5.2 Performing a New Installation of HDLM

This subsection describes how to perform a new installation of HDLM. The preparations before newly installing HDLM are required. For details about the preparations before installing HDLM, see section 3.5.1.

To perform a new installation of HDLM:

1. Log on to Linux as the root user.
2. Store the license key file (`*.plk`), using the name `hdlm_license`, directly under `/var/tmp`.

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

   Alternatively, create the `/etc/opt/DynamicLinkManager` directory and create the license key file (`dlm.lic_key`) in that directory. Execute the following commands:

   ```
   # mkdir /etc/opt/DynamicLinkManager
   # echo "license-key" >
   /etc/opt/DynamicLinkManager/dlm.lic_key
   ```

3. Insert the CD-ROM and mount it.

   If the CD-ROM is not mounted automatically, execute the `mount` command to mount the CD-ROM device to the given mount point. The following shows the given mount point for each distribution:

   – Mount point for Red Hat Enterprise Linux AS3 or Red Hat Enterprise Linux ES3

     `/mnt/cdrom`

     The following is an example of executing the command:

     ```
     # mount /dev/cdrom /mnt/cdrom
     ```

   – Mount point for Red Hat Enterprise Linux AS4 or Red Hat Enterprise Linux ES4

     When using CD-R, DVD-R, CD-ROM, or DVD-ROM:

     `/media/cdrom`

     When using CD-RW or DVD-RW:

     `/media/cdrecorder`

     The following example shows an example of executing the command when the mount point is `/media/cdrom`:

     ```
     # mount /dev/cdrom /media/cdrom
     ```

   – Mount point for SUSE LINUX Enterprise Server 9

     When using SUSE LINUX Enterprise Server 9, the mount point differs depending on the device (for example, CD-ROM, CD-R, or CD-RW, or DVD), which is installed on the host machine, used for performing an installation. The following shows the mount point for each CD/DVD device:

     Mount point when the CD/DVD device is a CD-ROM:

`/media/cdrom`

Mount point when the CD/DVD device is a CD-R or CD-RW:

`/media/cdrecorder`

Mount point when the CD/DVD device is a DVD:

`/media/dvd`, `/media/dvdrecorder`, or `/media/dvdram`

The following example shows an example of executing the command when the mount point is `/media/cdrom`:

```
# mount /dev/cdrom /media/cdrom
```

– Mount point for SUSE LINUX Enterprise Server 10

`/media/cdrom`

The following shows the automatic mount point for the CD-ROM:

`/media/`*media-volume-ID*

*media-volume-ID* indicates the volume name of media (such as a CD-ROM) that was formatted with the ISO-9660 filesystem. To check the value of *media-volume-ID*, execute the `volname` command.

The following example shows an example of executing the command when *media-volume-ID* is `VOL01234`:

```
# volname  /dev/cdrom
VOL01234
#
```

***Note***: you cannot install HDLM if you mount the CD-ROM to a mount point other than those above.

If you copy the contents of the CD-ROM to a directory and install HDLM from that directory, specify the directory name so that the name will be the same as the above mount point.

***Note:*** you cannot install HDLM if you specify a different directory name.

4. Check the version of HDLM.

   Execute the following command and check the version of HDLM to be installed. The following example shows how to execute the command when the mount point for the CD-ROM is `/media/cdrom`:

   ```
   # /media/cdrom/installhdlm -v
   ```

   A command execution example is described below. *xx–xx* is the version of HDLM.

   `KAPL09177-I HDLM version:` ***xx–xx***

5. Perform the installation.

Perform the installation by specifying `installhdlm` under the mount point. The following shows an installation example for each distribution:

- For Red Hat Enterprise Linux AS3 or Red Hat Enterprise Linux ES3

  ```
  # /mnt/cdrom/installhdlm
  ```

- For Red Hat Enterprise Linux AS4 or Red Hat Enterprise Linux ES4

  For example, when the CD device on the host machine is CD-R, execute the following command:

  ```
  # /media/cdrom/installhdlm
  ```

- For SUSE LINUX Enterprise Server 9

  The following example shows the command execution when the CD/DVD device on the host is a CD-ROM:

  ```
  # /media/cdrom/installhdlm
  ```

- For SUSE LINUX Enterprise Server 10

  The following example shows an example of executing the command when *media-volume-ID* is `VOL01234`:

  ```
  # /media/VOL01234/installhdlm
  ```

6. A message appears confirming the execution of the new installation. Enter `y` in response to the message.

   The license key file will be deleted when installation finishes.

   ```
   KAPL09093-I HDLM xx-xx will be installed. Is this OK ? [y/n]: y
   Preparing packages for installation...
   KAPL09076-I The permanent license was installed.
   HDLM-x.xx.x.xxx-xx
   KAPL09043-I The installation of HDLM-x.xx.x.xxx-xx completed successfully.
   ```

7. Make sure that HDLM is installed.

   Execute the following command to display detailed information about the installed package. If `HDLM Version` is *xx*-*xx*, the installed version of HDLM is correct. *xx*-*xx* indicates the HDLM version that was installed. The following is the information output for a Red Hat Enterprise Linux AS 3 package.

```
# rpm -qi HDLM
Name        : HDLM                         Relocations:  (not relocateable)
Version     : x.xx.x.xxx                        Vendor:  Hitachi, Ltd.
Release     : xx                            Build Date:  date/time
Install date : date/time                    Build Host:  inspire.hitachi.co.jp
Group       : System Environment/Driver     Source RPM:  HDLM-x.xx.x.xxx-xx.src.rpm
Size        : xxxxxxx                          License:  All Rights Reserved. Copyright (C)
2003, yyyy, Hitachi, Ltd.
Signature   : (none)
Packager    : Hitachi, Ltd.
Summary     : I/O Path Management Software
Description :
HDLM manages paths between a host and storage subsystem.
HDLM evenly distributes the load across paths and switches to another path if there is a
failure in a path being used, thus improving system reliability.
#
```

**Figure 3.7 Display of an Installed HDLM Package**

8. If you want to set the HDLM driver options, run the `dlmsetopt` utility.

   For details on this utility, see section 7.7.

   If you performed this step, proceed to step 12.

9. When you want to install HDLM without restarting the host, go to step 9. When you want to install HDLM with restarting the host, then go to step 12.

10. Load the HDLM alert driver and the filter driver.

    Enter the following command to load the HDLM alert driver (`sddlmadrv`) and the filter driver (`sddlmfdrv`):

    – For Red Hat Enterprise Linux AS 3 or Red Hat Enterprise Linux ES 3

    ```
    # /opt/DynamicLinkManager/bin/dlminsadrv
    # insmod sddlmfdrv
    ```

    – For SUSE LINUX Enterprise Server 9, SUSE LINUX Enterprise Server 10, Red Hat Enterprise Linux AS4, or Red Hat Enterprise Linux ES4

    ```
    # /opt/DynamicLinkManager/bin/dlminsadrv
    # modprobe sddlmfdrv
    ```

11. Execute the HDLM-configuration definition utility (`dlmcfgmgr`) to create an HDLM device.

    The utility displays a message asking whether you want to continue the environment setup (creation) processing of the HDLM device. Enter `y` in this message to continue the processing.

    Enter the utility as follows:

    ```
    # /sbin/dlmcfgmgr -r
    KAPL10339-I This operation will change the configuration of HDLM devices. Do you want to
    continue? [y/n]: y
    ```

```
KAPL10302-I /sbin/dlmcfgmgr completed normally.
```

For details on this utility, see section 7.3

12. Start the HDLM manager.

    Enter the following command to start the HDLM manager:

```
# /etc/init.d/DLMManager start
```

13. Add `/opt/DynamicLinkManager/bin` in the environment variable `PATH` in the environment setup file of the root user.

    When the BourneAgain shell or Korn shell is used:

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

    When the C shell is used:

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

    The path is temporarily added to the environment variable `PATH` so that the command can be executed in a simple form. To execute an HDLM command or HDLM utility without setting the environment variable `PATH`, specify the absolute path to execute the command or utility.

    If you have performed steps 9 to 11, go to step 14.

14. Restart the host.

    Execute the following command to restart the host:

```
# shutdown -r now
```

    A path will be established in the HDLM device and the HDLM manager will start.

15. If you have edited the `/etc/fstab` file as described in 3.5.1.1, restore the file to its status before it was edited.

    Delete the hash marks (`#`) from the lines that you commented out. The Linux functionality that adds `LABEL=` to a SCSI device is not supported in HDLM. Do not use this functionality.

    An example of how to edit the `/etc/fstab` file is shown in the following figure:

```
LABEL=/1        /             ext3     defaults                 1 1
/dev/hda1       /boot         ext3     defaults                 1 2
#/dev/sda       /mntpt        ext3     defaults,noauto          0 0
none            /dev/pts      devpts   gid=5,mode=620           0 0
none            /proc         proc     defaults                 0 0
none            /dev/shm      tmpfs    defaults                 0 0
/dev/hda3       swap          swap     defaults                 0 0
/dev/cdrom      /mnt/cdrom    iso9660  noauto,owner,kudzu,ro    0 0
/dev/fd0        /mnt/floppy   auto     noauto,owner,kudzu       0 0
```

    Delete the hash mark (`#`) at the beginning of each line in the shaded portions shown in the above figure.

16. Specify the settings required for using the volume management software.

    If the logical volume created with the volume management software is already being used, replace it with a logical volume whose physical volume is an HDLM device.

For details on setting the volume management software, see section 3.7.

17. Specify the settings required for operations in a cluster configuration.

    For a cluster configuration, change the logical device file names of SCSI devices specified in the cluster to logical device file names of HDLM devices.

    For details on setting the cluster software, see section 3.8, 3.9, or 3.10.

    Also, in the following cases, you must change the settings for Oracle RAC 10g:

    – You use a release of 10.1.0.3.0 or later of Oracle RAC 10g, and connect a host and a voting disk by using multiple paths (in a multi-path configuration).
    – You use a release of 10.2.0.2.0 or later of Oracle RAC 10g, and connect a host and a voting disk by using multiple paths (in a multi-path configuration).

    For details, see section 3.4.2.

18. Execute the `dlnkmgr` command's `view` operation to check the HDLM settings and status of each program.

    The following is 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
Reservation Status           :
Intermittent Error Monitor : off
HDLM Manager Ver      WakeupTime
Alive        xx-xx    yyyy/mm/dd hh:nn:ss
HDLM Alert Driver Ver      WakeupTime          ElogMem Size
Alive             xx-xx    yyyy/mm/dd hh:nn:ss 1000
HDLM Driver Ver      WakeupTime
Alive       xx-xx    yyyy/mm/dd hh:nn:ss
License Type     Expiration
Permanent        -
KAPL01001-I The HDLM command completed normally. Operation name = view,
completion time = yyyy/mm/dd hh:nn:ss
#
```

**Figure 3.8 Executing the dlnkmgr Command's View Operation**

    Even when cluster software is being used, `Support Cluster` is blank. However, the cluster function is operating normally.

19. Check that the correct version of HDLM is installed.

    If `HDLM version` *xx-xx* is displayed, the installed HDLM version is correct. *xx-xx* is the version of the installed HDLM.

20. Check that the programs are running properly.

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

If the programs are not running properly, see section 5.4, and then take actions.

After the installation, check the path configuration according to the procedure indicated in 3.11.

### 3.5.3 Preparations for an Upgrade Installation or Re-installation of HDLM

This section describes the preparations required before performing an upgrade installation or re-installation of HDLM. Before performing such an installation, unmount and back up the HDLM-managed HDLM devices.

To perform the preparations required before performing an update or overwrite installation of HDLM:

1. Terminate the processes of all applications that access the disks managed by HDLM.
2. Unmount the HDLM devices.

   If HDLM-managed devices have been mounted by specifying HDLM devices, unmount them.
3. Back up the HDLM-managed disks (for example, to tape).

   Perform this operation if necessary.

For details on how to re-install HDLM, see section 3.5.4. For details on how to upgrade-install HDLM, see section 3.5.5.

### 3.5.4 Performing a Re-installation of HDLM

Reinstalling HDLM without uninstalling an already installed HDLM is called an *overwrite installation*. If an HDLM file that is installed on a host is damaged, you can perform an overwrite installation to restore the HDLM file without changing any other settings. For an overwrite installation, an HDLM version that is the same as the installed version can be installed.

The following is the procedure for a re-installation. Some preparations before installing HDLM are required. For details about such preparations, see section 3.5.3.

The license key is required when you update HDLM after the valid license period has expired.

To perform an overwrite-installation of HDLM:

1. Log on to Linux as the root user.
2. Store the license key file, using the name hdlm_license, directly under /var/tmp.

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

   Notes on the license key for a re-installation:

When a permanent license key has been entered, you do not need to create the license key file.

When a temporary or emergency license key has been entered, the expiration date for the key is displayed.

If the temporary license key or the emergency license key is expired, you need to update the license key file (`dlm.lic_key`) in the `/etc/opt/DynamicLinkManager` directory.

When the license key file is not found during installation, the following message appears, and the process continues: `KAPL09090-W This operation will now be continued without updating the license.`

For details about how to create the license key, see section 3.5.2.

For details about the license key, see *Appendix B*.

3. Insert a CD-ROM and mount it.

   If the CD-ROM is not mounted automatically, execute the `mount` command to mount the CD-ROM to the given mount point. The following shows the given mount point for each distribution:

   – Mount point for Red Hat Enterprise Linux AS3 or Red Hat Enterprise Linux ES3

     `/mnt/cdrom`

     The following is an example of executing the command:

     ```
     # mount /dev/cdrom /mnt/cdrom
     ```

   – Mount point for Red Hat Enterprise Linux AS4 or Red Hat Enterprise Linux ES4

     When using CD-R, DVD-R, CD-ROM, or DVD-ROM:

     `/media/cdrom`

     When using CD-RW or DVD-RW:

     `/media/cdrecorder`

     The following example shows an example of executing the command when the mount point is `/media/cdrom`:

     ```
     # mount /dev/cdrom /media/cdrom
     ```

   – Mount point for SUSE LINUX Enterprise Server 9

     When using SUSE LINUX Enterprise Server 9, the mount point differs depending on the CD/DVD device (for example, CD-ROM, CD-R, or CD-RW, or DVD), which is installed on the host machine, used for performing an installation. The following shows the mount point for each device:

     Mount point when the CD/DVD device is a CD-ROM:

     `/media/cdrom`

     Mount point when the CD/DVD device is a CD-R/CD-RW:

     `/media/cdrecorder`

Mount point when the CD/DVD device is a DVD:

`/media/dvd`, `/media/dvdrecorder`, or `/media/dvdram`

The following example shows an example of executing the command when the mount point is `/media/cdrom`:

```
# mount /dev/cdrom /media/cdrom
```

– Mount point for SUSE LINUX Enterprise Server 10

`/media/cdrom`

The following shows the automatic mount point for the CD-ROM:

`/media/`*`media-volume-ID`*

*media-volume-ID* indicates the volume name of media (such as a CD-ROM) that was formatted with the ISO-9660 filesystem. To check the value of *media-volume-ID*, execute the `volname` command.

The following example shows an example of executing the command when *media-volume-ID* is `VOL01234`:

```
# volname  /dev/cdrom

VOL01234

#
```

***Note:*** you cannot install HDLM if you mount the CD-ROM to a mount point other than those above.

If you copy the contents of the CD-ROM to a directory and install HDLM from that directory, specify the directory name so that the name will be the same as the above mount point.

***Note:*** you cannot install HDLM if you specify a different directory name.

4. Check the version of HDLM.

   Execute the following command and check the version of HDLM to be installed. The following example shows how to execute the command when the mount point for the CD-ROM is `/media/cdrom`:

```
# /media/cdrom/installhdlm -v
```

   A command execution example is described below. *xx–xx* is the version of HDLM.

   `KAPL09177-I HDLM version:` ***xx–xx***

5. Perform the installation.

   Perform the installation by specifying `installhdlm` under the mount point. The following shows the installation example for each distribution:

   – For Red Hat Enterprise Linux AS3 or Red Hat Enterprise Linux ES3

```
# /mnt/cdrom/installhdlm
```

– For Red Hat Enterprise Linux AS4 or Red Hat Enterprise Linux ES4

For example, when the CD device on the host machine is CD-R, execute the following command:

```
# /media/cdrom/installhdlm
```

– For SUSE LINUX Enterprise Server 9

The following example shows the command execution when the CD/DVD device on the host is a CD-ROM:

```
# /media/cdrom/installhdlm
```

– For SUSE LINUX Enterprise Server 10

The following example shows an example of executing the command when *media-volume-ID* is `VOL01234`:

```
# /media/VOL01234/installhdlm
```

6. A message appears confirming the execution of the re-installation. Enter `y` in response to the message.
7. Make sure that HDLM is installed.

   Execute the following command to display the detailed information on the installed package. If `HDLM Version` is *x.xx.x.xxx*, the installed version of HDLM is correct. *x.xx.x.xxx* indicates the HDLM version that was installed. The following is the information output for a Red Hat Enterprise Linux AS 3 package.

```
# rpm -qi HDLM
Name        : HDLM                        Relocations:  (not relocateable)
Version     : x.xx.x.xxx                      Vendor:  Hitachi, Ltd.
Release     : xx                          Build Date:  date/time
Install date: date/time                   Build Host:  inspire.hitachi.co.jp
Group       : System Environment/Driver   Source RPM:  HDLM-x.xx.x.xxx-xx.src.rpm
Size        : xxxxxxx                        License:  All Rights Reserved. Copyright (C)
2003, yyyy, Hitachi, Ltd.
Signature   : (none)
Packager    : Hitachi, Ltd.
Summary     : I/O Path Management Software
Description :
HDLM manages paths between a host and storage subsystem.
HDLM evenly distributes the load across paths and switches to another path if there is a
failure in a path being used, thus improving system reliability.
#
```

**Figure 3.9 Display of an Installed HDLM Package**

8. When you want to set the HDLM driver options, run the `dlmsetopt` utility for setting HDLM driver option.

For details on this utility, see section 7.7.

If you performed this step, proceed to step 12.

9. When you want to install HDLM without restarting the host, go to step 9. When you want to install HDLM with restarting the host, then go to step 12.

10. Load the HDLM alert driver and the filter driver.

    Enter the following command to load the HDLM alert driver (`sddlmadrv`) and the filter driver (`sddlmfdrv`):

    – For Red Hat Enterprise Linux AS 3 or Red Hat Enterprise Linux ES 3

    ```
    # /opt/DynamicLinkManager/bin/dlminsadrv
    # insmod sddlmfdrv
    ```

    – For SUSE LINUX Enterprise Server 9, SUSE LINUX Enterprise Server 10, Red Hat Enterprise Linux AS4, or Red Hat Enterprise Linux ES4

    ```
    # /opt/DynamicLinkManager/bin/dlminsadrv
    # modprobe sddlmfdrv
    ```

11. Execute the HDLM-configuration definition utility (`dlmcfgmgr)` to create an HDLM device.

    The utility displays a message asking whether you want to continue the environment setup (creation) processing of the HDLM device. Enter `y` in this message to continue the processing.

    Enter the utility as follows:

    ```
    # /sbin/dlmcfgmgr -r
    KAPL10339-I This operation will change the configuration of HDLM devices. Do you want to
    continue? [y/n]: y
    KAPL10302-I /sbin/dlmcfgmgr completed normally.
    ```

    For details on this utility, see section 8.3

12. Start the HDLM manager.

    Enter the following command to start the HDLM manager:

    ```
    # /etc/init.d/DLMManager start
    ```

13. Add `/opt/DynamicLinkManager/bin` in the environment variable `PATH` in the environment setup file for the shell used by the root user.

    Add `/opt/DynamicLinkManager/bin` in the environment variable `PATH` in the environment setup file of the root user.

    When the BourneAgain shell or Korn shell is used:

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

    When the C shell is used:

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

The path is temporarily added to the environment variable `PATH` so that the command can be executed in a simple form. To execute an HDLM command or HDLM utility without setting the environment variable `PATH`, specify the absolute path to execute the command or utility.

If you have performed steps 9 to 11, go to step 14.

14. Restart the host.

    Execute the following command to restart the host:

```
# shutdown -r now
```

    A path will be established in the HDLM device and the HDLM manager will start.

15. If you have edited the `/etc/fstab` file as described in 3.5.1.1, restore the file to its status before it was edited.

    If a device to be managed by HDLM has been mounted by defining a SCSI device in the `/etc/fstab` file, change the definition in the `/etc/fstab` file to the logical device name of the HDLM device. The Linux functionality that adds `LABEL=` to a SCSI device is not supported in HDLM. Do not use this functionality.

    An example of how to edit the `/etc/fstab` file is shown in the following figure:

```
LABEL=/1            /                  ext3     defaults                  1 1
/dev/hda1           /boot              ext3     defaults                  1 2
#/dev/sda           /mntpt             ext3     defaults,noauto           0 0
none                /dev/pts           devpts   gid=5,mode=620            0 0
none                /proc              proc     defaults                  0 0
none                /dev/shm           tmpfs    defaults                  0 0
/dev/hda3           swap               swap     defaults                  0 0
/dev/cdrom          /mnt/cdrom         iso9660  noauto,owner,kudzu,ro     0 0
/dev/fd0            /mnt/floppy        auto     noauto,owner,kudzu        0 0
```

16. Specify the settings required for using the volume management software.

    If the logical volume created with the volume management software is already being used, replace it with a logical volume whose physical volume is an HDLM device.

    For details on setting the volume management software, see section 3.7.

17. Specify the settings required for operations in a cluster configuration.

    For a cluster configuration, change the logical device file names of SCSI devices specified in the cluster to logical device file names of HDLM devices.

    For details on setting the cluster software, see section 3.8, 3.9, or 3.10.

### 3.5.5 Performing an Upgrade Installation of HDLM

The following is the procedure for an upgrade installation. The preparations before installing HDLM are required. For details about the preparations before installing HDLM, see section 3.5.3.

A license key file is required when install HDLM 5.9.1 by overwriting HDLM 5.4 or earlier or when install HDLM 5.4 or later after the valid license period has expired.

To perform an upgrade installation of HDLM:

1. Log on to Linux as the root user.
2. Store the license key file, using the name `hdlm_license`, directly under `/var/tmp`.

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

   Alternatively, create a license key file (`dlm.lic_key`) in the `/etc/opt/DynamicLinkManager` directory. If the `/etc/opt/DynamicLinkManager` directory does not exist, create it and create the license key file in the directory. Execute the following command:

   ```
   # mkdir /etc/opt/DynamicLinkManager
   # echo "license-key" > /etc/opt/DynamicLinkManager/dlm.lic_key
   ```

3. Insert and mount a CD-ROM.

   If the CD-ROM is not mounted automatically, execute the `mount` command to mount the CD-ROM to the given mount point. The following shows the given mount point for each distribution:

   – Mount point for Red Hat Enterprise Linux AS3 or Red Hat Enterprise Linux ES3

     `/mnt/cdrom`

     The following is an example of executing the command:

     ```
     # mount /dev/cdrom /mnt/cdrom
     ```

   – Mount point for Red Hat Enterprise Linux AS4 or Red Hat Enterprise Linux ES4

     When using CD-R, DVD-R, CD-ROM, or DVD-ROM:

     `/media/cdrom`

     When using CD-RW or DVD-RW:

     `/media/cdrecorder`

     The following example shows an example of executing the command when the mount point is `/media/cdrom`:

     ```
     # mount /dev/cdrom /media/cdrom
     ```

   – Mount point for SUSE LINUX Enterprise Server 9

     When using SUSE LINUX Enterprise Server 9, the mount point differs depending on the CD/DVD device (for example, CD-ROM, CD-R, or CD-RW, or DVD), which is installed on the host machine, used for performing an installation. The following shows the mount point for each CD/DVD device:

     Mount point when the CD/DVD device is a CD-ROM:

     `/media/cdrom`

Mount point when the CD/DVD device is a CD-R/CD-RW:

`/media/cdrecorder`

Mount point when the CD/DVD device is a DVD:

`/media/dvd`, `/media/dvdrecorder`, or `/media/dvdram`

The following example shows an example of executing the command when the mount point is `/media/cdrom`:

```
# mount /dev/cdrom /media/cdrom
```

– Mount point for SUSE LINUX Enterprise Server 10

`/media/cdrom`

The following shows the automatic mount point for the CD-ROM:

`/media/`*`media-volume-ID`*

*media-volume-ID* indicates the volume name of media (such as a CD-ROM) that was formatted with the ISO-9660 filesystem. To check the value of *media-volume-ID*, execute the `volname` command.

The following example shows an example of executing the command when *media-volume-ID* is `VOL01234`:

```
# volname  /dev/cdrom

VOL01234

#
```

***Note***: you cannot install HDLM if you mount the CD-ROM to a mount point other than those above.

If you copy the contents of the CD-ROM to a directory and install HDLM from that directory, specify the directory name so that the name will be the same as the above mount point.

***Note***: you cannot install HDLM if you specify a different directory name.

4. Check the version of HDLM.

   Execute the following command and check the version of HDLM to be installed. The following example shows how to execute the command when the mount point for the CD-ROM is `/media/cdrom`:

   ```
   # /media/cdrom/installhdlm -v
   ```

   A command execution example is described below. *xx–xx* is the version of HDLM.

   `KAPL09177-I HDLM version: `***`xx–xx`***

5. Perform the installation.

   Perform the installation by specifying `installhdlm` under the mount point. The following shows the installation example for each distribution:

– For Red Hat Enterprise Linux AS3 or Red Hat Enterprise Linux ES3

```
# /mnt/cdrom/installhdlm
```

– For Red Hat Enterprise Linux AS4 or Red Hat Enterprise Linux ES4

For example, when the CD device on the host machine is CD-R, execute the following command:

```
# /media/cdrom/installhdlm
```

– For SUSE LINUX Enterprise Server 9

The following example shows the command execution when the CD/DVD device on the host is a CD-ROM:

```
# /media/cdrom/installhdlm
```

– For SUSE LINUX Enterprise Server 10

The following example shows an example of executing the command when *media-volume-ID* is `VOL01234`:

```
# /media/VOL01234/installhdlm
```

6. A message appears confirming the execution of the upgrade installation. Enter `y` in response to the message.
7. Make sure that HDLM is installed.

   Execute the following command to display the detailed information on the installed package. If `HDLM Version` is ***x.xx.x.xxx***, the installed version of HDLM is correct. ***x.xx.x.xxx*** indicates the HDLM version that was installed. The following is the information output for a Red Hat Enterprise Linux AS 3 package.

```
# rpm -qi HDLM
Name        : HDLM                      Relocations:  (not relocateable)
Version     : x.xx.x.xxx                     Vendor:  Hitachi, Ltd.
Release     : xx                         Build Date:  date/time
Install date : date/time                 Build Host:  inspire.hitachi.co.jp
Group       : System Environment/Driver  Source RPM:  HDLM-x.xx.x.xxx-xx.src.rpm
Size        : xxxxxxx                       License:  All Rights Reserved. Copyright (C)
2003, yyyy, Hitachi, Ltd.
Signature   : (none)
Packager    : Hitachi, Ltd.
Summary     : I/O Path Management Software
Description :
HDLM manages paths between a host and storage subsystem.
HDLM evenly distributes the load across paths and switches to another path if there is a
failure in a path being used, thus improving system reliability.
#
```

**Figure 3.10 Display of an Installed HDLM Package**

8. When you want to set the HDLM driver options, run the `dlmsetopt` utility for setting HDLM driver option.

   For details on this utility, see section 7.7.

   If you performed this step, proceed to step 12.

9. When you want to install HDLM without restarting the host, go to step 9. When you want to install HDLM with restarting the host, then go to step 12.

10. Load the HDLM alert driver and the filter driver.

    Enter the following command to load the HDLM alert driver (`sddlmadrv`) and the filter driver (`sddlmfdrv`):

    – For Red Hat Enterprise Linux AS 3 or Red Hat Enterprise Linux ES 3

    ```
    # /opt/DynamicLinkManager/bin/dlminsadrv

    # insmod sddlmfdrv
    ```

    – For SUSE LINUX Enterprise Server 9, SUSE LINUX Enterprise Server 10, Red Hat Enterprise Linux AS4, or Red Hat Enterprise Linux ES4

    ```
    # /opt/DynamicLinkManager/bin/dlminsadrv

    # modprobe sddlmfdrv
    ```

11. Execute the HDLM-configuration definition utility (`dlmcfgmgr)` to create an HDLM device.

    The utility displays a message asking whether you want to continue the environment setup (creation) processing of the HDLM device. Enter `y` in this message to continue the processing.

    Enter the utility as follows:

    ```
    # /sbin/dlmcfgmgr -r
    KAPL10339-I This operation will change the configuration of HDLM devices. Do you want to continue? [y/n]: y
    KAPL10302-I /sbin/dlmcfgmgr completed normally.
    ```

    For details on this utility, see section 7.3

12. Start the HDLM manager.

    Enter the following command to start the HDLM manager:

    ```
    # /etc/init.d/DLMManager start
    ```

13. Add `/opt/DynamicLinkManager/bin` in the environment variable `PATH` in the environment setup file for the shell used by the root user.

    Add `/opt/DynamicLinkManager/bin` in the environment variable `PATH` in the environment setup file of the root user.

    When the BourneAgain shell or Korn shell is used:

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

    When the C shell is used:

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

The path is temporarily added to the environment variable PATH so that the command can be executed in a simple form. To execute an HDLM command or HDLM utility without setting the environment variable PATH, specify the absolute path to execute the command or utility.

If you have performed steps 9 to 11, go to step 14.

14. Restart the host.

    Execute the following command to restart the host:

    ```
    # shutdown -r now
    ```

    A path will be established in the HDLM device and the HDLM manager will start.

15. If you have edited the /etc/fstab file as described in 3.5.1.1, restore the file to its status before it was edited.

    If a device to be managed by HDLM is mounted by defining a SCSI device in the /etc/fstab file, change the definition in the /etc/fstab file to a logical device name of an HDLM device. The Linux functionality that adds LABEL= to a SCSI device is not supported in HDLM. Do not use this functionality.

    An example of how to edit the /etc/fstab file is shown in the following figure:

```
LABEL=/1       /             ext3     defaults                1 1
/dev/hda1      /boot         ext3     defaults                1 2
#/dev/sda      /mntpt        ext3     defaults,noauto         0 0
none           /dev/pts      devpts   gid=5,mode=620          0 0
none           /proc         proc     defaults                0 0
none           /dev/shm      tmpfs    defaults                0 0
/dev/hda3      swap          swap     defaults                0 0
/dev/cdrom     /mnt/cdrom    iso9660  noauto,owner,kudzu,ro   0 0
/dev/fd0       /mnt/floppy   auto     noauto,owner,kudzu      0 0
```

16. Specify the settings required for using the volume management software.

    If the logical volume created with the volume management software is already being used, replace it with a logical volume whose physical volume is an HDLM device.

    For details on setting the volume management software, see section 3.7.

17. Specify the settings required for operations in a cluster configuration.

    For a cluster configuration, change the logical device file names of SCSI devices specified in the cluster to logical device file names of HDLM devices.

    For details on setting the cluster software, see section 3.8, 3.9, or 3.10.

## 3.6 Installing HDLM for Managing Boot Disks

This section explains the procedure for installing HDLM in an environment in which the boot disk is on a SCSI device in the storage subsystem, and for creating an HDLM device. This section also explains the procedure for setting up an environment in which the HDLM device is the boot disk.

### 3.6.1 Notes on Installing HDLM in a Boot Disk Environment

Note the following when using HDLM on a boot disk:

- The storage subsystem for which the boot disk will be created must satisfy the following conditions:
  - The storage subsystem supports the boot disk functionality even if HDLM is not used.
  - HBAs support startup from the storage subsystem.
- When you install the OS, create the root directory (`/`) and the `/boot` directory in separate partitions.
- Since an initial RAM disk image for HDLM is mounted in `/initrd` directory when using Red Hat Enterprise Linux AS3/ES3, the following message indicating that unmounting is impossible appears at startup. However, there are no operational problems.

  ```
  Unmounting initrd :  umount: initrd: device is busy        [FAILED]
  ```
- In an environment where multiple HBAs are used, enable the BIOS setting for only one HBA, and disable the BIOS settings for other HBAs.
- When Red Hat Enterprise Linux AS3/ES3 is used, HDLM does not support an environment in which the boot disk is an LVM or md device.
- When Red Hat Enterprise Linux AS4/ES4 is used, HDLM supports an environment that uses LVM2 as a boot disk.
- If you changed the configuration during startup or if the `/etc/fstab` file settings are incorrect, the system might not start.
- If the root directory (`/`) has not been mounted at startup, you cannot collect a boot error log (`/etc/opt/DynamicLinkManager/hdlmboot.log`). In this case, if a problem occurs, use the log displayed on the console to analyze the problem.
- When you start the OS with HDLM installed on the boot disk, you must change the `/etc/fstab` file root directory (`/`) to a SCSI device (as stated in the HDLM uninstallation procedure) before shutting down the system. In this case, the following messages, indicating that an attempt to unload the HDLM driver failed, appear. However, there are no operational problems.

  ```
  Unloading HDLM Filter driver sddlmfdrv: Device or resource busy    [FAILED]

  Unloading HDLM Alert driver sddlmadrv: Device or resource busy     [FAILED]

  Removing HDLM devices                                              [FAILED]
  ```
- The `dlmcfgmgr` log created at startup is output to `/etc/opt/DynamicLinkManager/hdlmboot.log`.

- Once you have built an environment for starting the OS from an HDLM device, you cannot change the boot disk.
- If you do not use the storage subsystems disk set for the boot disk, before you install HDLM, edit the `/etc/fstab` file in the following manner to disable the `LABEL` specification and enable the sd device specification.

  a) Check the installation directory of the OS specified in the `LABEL` setting.

  ```
  # cat /etc/fstab
  LABEL=/    /    ext3    defaults    1 1
  LABEL=/boot    /boot    ext3    defaults    1 2
  ```

  Confirm that LABEL indicates the root directory and the `/boot` directory.

  Check the correspondence between the `LABEL` setting and the sd device.

  ```
  # mount
  /dev/sda2 on / type ext3 (rw)
  none on /proc type proc (rw)
  none on /dev/pts type devpts (rw,gid=5,mode=620)
  usbdevfs on /proc/bus/usb type usbdevfs (rw)
  /dev/sda1 on /boot type ext3 (rw)
  none on /dev/shm type tmpfs (rw)
  ```

  Confirm that the root directory is `/dev/sda2`, and the /`boot` directory is `/dev/sda1`.

  Using an editor such as vi, change the `LABEL` specification to an sd device specification.

  (before)

  ```
  LABEL=/    /    ext3    defaults    1 1
  LABEL=/boot    /boot    ext3    defaults    1 2
  ```

  (after)

  ```
  /dev/sda2    /    ext3    defaults        1 1
  /dev/sda1    /boot    ext3    defaults        1 2
  ```

  Install HDLM.

  After configuring the HDLM device, execute the `dlmcfgmgr -o` command to exclude the HDLM device that corresponds to the boot disk as a management target.

- If the SCSI device name has been changed by using the `udev` function, see section 3.6.3 and 3.6.4 and perform operations while reading the SCSI device names in those subsections as the new SCSI device names (`udev` name) that were changed by using the `udev` function.

  In an environment where HDLM has been installed, you can use the `dlmcfgmgr` utility to check the correspondence between a SCSI device and an `udev` name.

  Figure 3.11 shows an example of executing the `dlmcfgmgr` utility with the `-v` and `-udev` parameters specified.

```
# dlmcfgmgr -v -udev
HDevName     Management  Device       Host  Channel Target  Lun   Udev
/dev/sddlmaa configured  /dev/sda        0     0       0      0   /dev/aaaaaaaa
KAPL10302-I /sbin/dlmcfgmgr completed normally.
```

**Figure 3.11 Example of Executing the dlmcfgmgr Utility with the -v and -udev Parameters Specified**

Check the shaded portion, which shows the correspondence between the SCSI device and the HDLM device.

### 3.6.2 Overview of the Procedure for Installing HDLM in a Boot Disk Environment

The following explains how to perform a new installation of HDLM in a boot disk environment that uses a SCSI device. (For details on the following procedure, see section 3.6.3.)

To prepare for the installation:

1. Check whether `LABEL` is used. If it is, remove it.

   To use HDLM, you must remove the `LABEL` setting, since a SCSI device cannot be identified from the value set in `LABEL` in an HDLM environment. See steps 1 to 14 in 3.6.3.

2. Install HDLM.

   See step 15 in 3.6.3.

3. Edit the configuration file to change the definition so that the OS will start with HDLM specified. After editing the configuration file, restart the host to activate the setting. See steps 16 to 33 in 3.6.3.

After completing the above procedure, you can perform a new installation of HDLM in a boot disk environment that uses a SCSI device.

The following explains how to perform an upgrade installation of HDLM in a boot disk environment that uses a SCSI device. (For details on the following procedure, see section 3.6.4.)

To perform such an installation:

1. Edit the configuration file to change the definition so that the OS will start with a SCSI device specified.

   In the definition before the change, the OS starts with HDLM specified. See steps 1 to 16 in 3.6.4.

2. Perform an upgrade installation of HDLM.

   See step 17 in 3.6.4.

3. Change the definition so that the OS will start with HDLM specified.

   Change the definition to restore the configuration file to its state before the upgrade installation was performed. See steps 18 to 33 in 3.6.4.

After completing the above procedure, you can perform an upgrade installation of HDLM in a boot disk environment that uses a SCSI device.

The following explains how to migrate a single-path boot disk environment that uses a logical volume (LVM2) on a SCSI device to a multi-path boot disk environment that uses a logical volume on an HDLM device. (For details on the following procedure, see section 3.6.3.)

To prepare for the installation:

1. Check whether `LABEL` is used. If it is, remove it.

   To use HDLM, you must remove the `LABEL` setting, since a SCSI device cannot be identified from the value set in `LABEL` in an HDLM environment. See steps 1 to 7 in 3.6.3.

2. Install HDLM.

   See step 8 in 3.6.3.

3. Edit the configuration file to change the definition so that the OS will start with HDLM specified.

   After editing the configuration file, restart the host to activate the setting. See steps 9 to 21 in 3.6.3.

After completing the above procedure, you can migrate to a multi-path boot disk environment that uses a logical volume on an HDLM device.

The following explains how to perform an upgrade installation of HDLM in a multi-path boot disk environment that uses a logical volume (LVM2) on an HDLM device. (For details on the following procedure, see section 3.6.6.)

To perform such an installation:

1. Edit the configuration file to change the definition so that the OS will start with a SCSI device specified.

   In the definition before the change, the OS starts with HDLM specified. See steps 1 to 13 in 3.6.6.

2. Perform an upgrade installation of HDLM.

   See step 14 in 3.6.6.

3. Change the definition so that the OS will start with HDLM specified.

   Change the definition to restore the configuration file to its state before the upgrade installation was performed. See steps 15 to 27 in 3.6.6.

After completing the above procedure, you can perform an upgrade installation of HDLM in a multi-path boot disk environment that uses a logical volume (LVM2) on an HDLM device.

### 3.6.3 Settings for Using an HDLM Device as a Boot Disk

This subsection describes how to perform a new installation of HDLM in a single-path boot disk environment that uses a SCSI device and how to set up the environment. Note that if the settings are incorrect, the OS might not start. For details about what action to take if the OS cannot be started from an HDLM device, see 3.6.7.

To install HDLM in a boot disk environment that uses a SCSI device, and set up the environment:

1. Log in to Linux as a user with root privileges.

   In an HDLM environment, a SCSI device cannot be identified from the value set in `LABEL`. Before installing HDLM, therefore, you must remove the `LABEL` setting. Perform steps 2 and 3 to check whether `LABEL` is set.

2. Check whether the definition in the `/etc/fstab` file contains `LABEL`.

   Check the contents of the `/etc/fstab` file.

   – When the boot loader is LILO or GRUB

     Figure 3.12 shows a content example of `/etc/fstab` file.

| | | | | |
|---|---|---|---|---|
| LABEL=/ | / | ext2 | defaults | 1 1 |
| LABEL=/boot | /boot | ext2 | defaults | 1 2 |
| LABEL=/tmp | /tmp | ext2 | defaults | 1 2 |
| LABEL=/var | /var | ext2 | defaults | 1 2 |
| LABEL=/usr | /usr | ext2 | defaults | 1 2 |
| none | /dev/pts | devpts | gid=5,mode=620 | 0 0 |
| none | /proc | proc | defaults | 0 0 |
| none | /dev/shm | tmpfs | defaults | 0 0 |
| /dev/sda3 | swap | swap | defaults | 0 0 |
| /dev/cdrom | /mnt/cdrom / | udf,iso9660 | noauto,owner,kudzu,ro | 0 0 |
| /dev/fd0 | mnt/floppy | auto | noauto,owner,kudzu | 0 0 |

**Figure 3.12 Example Content of /etc/fstab file When the Boot Loader Is LILO or GRUB**

   `LABEL` is set in the shaded portion.

   If Red Hat Enterprise Linux AS4/ES4 is being used, you must also check the following line:

```
LABEL=SWAP-sda3    swap      swap         defaults            0  0
```

   – When the boot loader is ELILO installed on an IPF host

     Figure 3.13 shows example content of `/etc/fstab` file.

| | | | | |
|---|---|---|---|---|
| LABEL=/ | / | ext3 | defaults | 1 1 |
| /dev/sda1 | /boot/efi | vfat | defaults | 0 0 |
| LABEL=/tmp | /tmp | ext3 | defaults | 1 2 |
| LABEL=/var | /var | ext3 | defaults | 1 2 |
| LABEL=/usr | /usr | ext3 | defaults | 1 2 |
| none | /dev/pts | devpts | gid=5,mode=620 | 0 0 |
| none | /proc | proc | defaults | 0 0 |
| none | /dev/shm | tmpfs | defaults | 0 0 |
| /dev/sda3 | swap | swap | defaults | 0 0 |
| /dev/cdrom | /mnt/cdrom | udf,iso9660 | noauto,owner,kudzu,ro | 0 0 |

**Figure 3.13 Example Content of /etc/fstab file When the Boot Loader Is ELILO**

   `LABEL` is set in the shaded portion.

   If Red Hat Enterprise Linux AS4/ES4 is being used, you must also check the following line:

```
LABEL=SWAP-sda3    swap      swap         defaults            0  0
```

3. Check whether the definition in the boot loader configuration file contains `LABEL`.

   Separate examples are provided for when the boot loader is LILO, GRUB, or ELILO running on an IPF host.

   – When the boot loader is LILO

     Figure 3.14 shows example content of the configuration file.

```
boot=/dev/sda
map=/boot/map
install=/boot/boot.b
prompt
timeout=50
linear
default=linux

image=/boot/vmlinuz-2.4.21-32.0.1.ELsmp
     label=linux
     initrd=/boot/initrd-2.4.21-20-32.0.1.ELsmp
     read-only
     append="root=LABEL=/"
```

**Figure 3.14 Example Content of /etc/lilo.conf file**

`LABEL` is set in the shaded portion.

   – When the boot loader is GRUB

   Figure 3.15 shows example content of the configuration file.

```
default=0
timeout=10
splashimage=(hd0,0)/grub/splash.xpm.gz

title Red Hat Enterprise Linux AS 3.0 (2.4.21-32.0.1.ELsmp)
   root (hd0,0)
   kernel /vmlinuz-2.4.21-32.0.1.ELsmp ro root=LABEL=/
   initrd /initrd- 2.4.21-32.0.1.ELsmp
```

**Figure 3.15 Example Content of /etc/grub.conf file**

`LABEL` is set in the shaded portion.

   – When the boot loader is ELILO running on an IPF host

     Figure 3.16 shows example content of the configuration file.

```
prompt
timeout=50
linear
default=2.4.21-27.EL

image=vmlinuz-2.4.21-27.EL
    label=2.4.21-27.EL
    initrd=initrd-2.4.21-27.EL.img
    read-only
    append="root=LABEL=/"
```

**Figure 3.16 Example Content of /etc/elilo.conf file**

`LABEL` is set in the shaded portion.

In steps 2 and 3, you checked the `/etc/fstab` file and the boot loader configuration file. If `LABEL` is not set in these files, go to step 15 and perform a new installation of HDLM. If `LABEL` is set in either or both of these files, perform steps 4 to 14 to check the mounted file system, and then to remove the `LABEL` setting.

4. Execute the `mount` command to check the mounted file system.

   Check the correspondence between the SCSI device and `LABEL` that you checked in steps 2 and 3. You will need this information to remove the `LABEL` setting.

   – When the boot loader is LILO or GRUB

     Figure 3.17 shows an example of executing the `mount` command.

```
# mount
/dev/sda2 on / type ext2 (rw)
none on /proc type proc (rw)
none on /dev/pts type devpts (rw,gid=5,mode=620)
usbdevfs on /proc/bus/usb type usbdevfs (rw)
/dev/sda1 on /boot type ext2 (rw)
/dev/sda4 on /tmp type ext2 (rw)
/dev/sda5 on /var type ext2 (rw)
/dev/sda6 on /usr type ext2 (rw)
none on /dev/shm type tmpfs (rw)
```

**Figure 3.17 Example of Executing the mount Command When the Boot Loader Is LILO or GRUB**

     The shaded portion indicates the SCSI device for which `LABEL` is set.

   – When the boot loader is ELILO running on an IPF host

     Figure 3.18 shows an example of executing the `mount` command.

```
# mount
/dev/sda2 on / type ext3 (rw)
none on /proc type proc (rw)
none on /dev/pts type devpts (rw,gid=5,mode=620)
usbdevfs on /proc/bus/usb type usbdevfs (rw)
/dev/sda1 on /boot/efi type vfat (rw)
/dev/sda4 on /tmp type ext3 (rw)
/dev/sda5 on /var type ext3 (rw)
/dev/sda6 on /usr type ext3 (rw)
none on /dev/shm type tmpfs (rw)
```

**Figure 3.18 Example of Executing the mount Command When the Boot Loader Is ELILO**

The shaded portion indicates the SCSI device for which `LABEL` is set.

Remove the `LABEL` setting in the `/etc/fstab` file and the boot loader configuration file according to the information for the mounted file system. First perform steps 5 to 7 to remove the `LABEL` setting in the `/etc/fstab` file.

If Red Hat Enterprise Linux AS4/ES4 is being used, check the SCSI device name by referring to the `/proc/swaps` file.

```
Filename          Type       Size      Used    Priority
/dev/sda3         partition  2096448   0       -1
```

**Figure 3.19 Example Content of /proc/swaps file**

5. Edit the `/etc/fstab` file to remove the `LABEL` setting.

   The detailed procedure is explained in steps 6 and 7.

   – When the boot loader is LILO or GRUB

     Figure 3.20 shows an example of editing the `/etc/fstab` file.

```
#LABEL=/       /            ext2          defaults                 1 1
/dev/sda2      /            ext2          defaults                 1 1
#LABEL=/boot   /boot        ext2          defaults                 1 2
/dev/sda1      /boot        ext2          defaults                 1 2
#LABEL=/tmp    /tmp         ext2          defaults                 1 2
/dev/sda4      /tmp         ext2          defaults                 1 2
#LABEL=/var    /var         ext2          defaults                 1 2
/dev/sda5      /var         ext2          defaults                 1 2
#LABEL=/usr    /usr         ext2          defaults                 1 2
/dev/sda6      /usr         ext2          defaults                 1 2
none           /dev/pts     devpts        gid=5,mode=620           0 0
none           /proc        proc          defaults                 0 0
none           /dev/shm     tmpfs         defaults                 0 0
/dev/sda3      swap         swap          defaults                 0 0
/dev/cdrom     /mnt/cdrom   udf,iso9660   noauto,owner,kudzu,ro    0 0
/dev/fd0       /mnt/floppy  auto          noauto,owner,kudzu       0 0
```

**Figure 3.20 Example of Editing /etc/fstab file When the Boot Loader Is LILO or GRUB**

Comment out the `LABEL` line, and then add the shaded line.

If Red Hat Enterprise Linux AS4/ES4 is being used, you must also comment out the following LABEL line, and then add the line specified by the SCSI device.

```
# LABEL=SWAP-sda3    swap      swap        defaults            0 0
/dev/sda3            swap      swap        defaults            0 0
```

– When the boot loader is ELILO running on an IPF host

Figure 3.21 shows an example of editing the /etc/fstab file.

```
#LABEL=/      /            ext3          defaults                  1 1
/dev/sda2     /            ext3          defaults                  1 1
/dev/sda1     /boot/efi    vfat          defaults                  0 0
#LABEL=/tmp   /tmp         ext3          defaults                  1 2
/dev/sda4     /tmp         ext3          defaults                  1 2
#LABEL=/var   /var         ext3          defaults                  1 2
/dev/sda5     /var         ext3          defaults                  1 2
#LABEL=/usr   /usr         ext3          defaults                  1 2
/dev/sda6     /usr         ext3          defaults                  1 2
none          /dev/pts     devpts        gid=5,mode=620            0 0
none          /proc        proc          defaults                  0 0
none          /dev/shm     tmpfs         defaults                  0 0
/dev/sda3     swap         swap          defaults                  0 0
/dev/cdrom    /mnt/cdrom   udf,iso9660   noauto,owner,kudzu,ro     0 0
/dev/fd0      /mnt/floppy  auto          noauto,owner,kudzu         0 0
```

**Figure 3.21 Example of Editing /etc/fstab file When the Boot Loader Is ELILO**

Comment out the LABEL line, and then add the shaded line.

If Red Hat Enterprise Linux AS4/ES4 is being used, you must also comment out the following LABEL line, and then add the line specified by the SCSI device.

If Red Hat Enterprise Linux AS4/ES4 is being used, you must also comment out the following line:

```
# LABEL=SWAP-sda3    swap      swap        defaults            0 0
/dev/sda3            swap      swap        defaults            0 0
```

6. If necessary, back up the /etc/fstab file.
7. Edit the /etc/fstab file.

   Refer to the correspondence between LABEL and the SCSI device that you checked in step 4.

   Copy the registration of the LABEL specification.

   Comment out the existing settings.

   Edit the file to disable the LABEL specification and enable the SCSI device specification. After editing the /etc/fstab file, perform steps 8 to 13 to remove the LABEL setting contained in the boot loader configuration file.
8. Edit the boot loader configuration file to remove the LABEL setting.

   Edit the boot loader configuration file as shown in Figure 3.22, Figure 3.23, and Figure 3.24.

   The detailed procedure is explained in steps 9 to 13.

   Separate examples are provided for when the boot loader is LILO, GRUB, or ELILO running on an IPF host.

   – When the boot loader is LILO

Figure 3.22 shows an example of editing the configuration file. After changing the configuration file, execute `/sbin/lilo` command to activate the new setting.

```
boot=/dev/sda
map=/boot/map
install=/boot/boot.b
prompt
timeout=50
linear
#default=linux
default=linux-sd

image=/boot/vmlinuz-2.4.21-32.0.1.ELsmp
      label=linux
      initrd=/boot/initrd-2.4.21-32.0.1.ELsmp
      read-only
      append="option root=LABEL=/ option"

image=/boot/vmlinuz-2.4.21-32.0.1.ELsmp
      label=linux-sd
      initrd=/boot/initrd-2.4.21-32.0.1.ELsmp
      read-only
      append="option"
      root=/dev/sda2

 Legend:
   option:   option that depends on the user environment is specified
```

**Figure 3.22 Example of Editing /etc/lilo.conf file**

Edit the shaded portion.

– When the boot loader is GRUB

Figure 3.23 shows an example of editing the configuration file.

```
#default=0
default=1
timeout=10
splashimage=(hd0,0)/grub/splash.xpm.gz

title Red Hat Enterprise Linux AS 3.0 (2.4.21-32.0.1.ELsmp)
    root (hd0,0)
    kernel /vmlinuz-2.4.21-32.0.1.ELsmp ro option root=LABEL=/ option
    initrd /initrd- 2.4.21-32.0.1.ELsmp

title Red Hat Enterprise Linux AS 3.0 (2.4.21-32.0.1.ELsmp-sd)
    root (hd0,0)
#    kernel /vmlinuz-2.4.21-32.0.1.ELsmp ro option root=LABEL=/ option
    kernel /vmlinuz-2.4.21-32.0.1.ELsmp ro option root=/dev/sda2 option
    initrd /initrd- 2.4.21-32.0.1.ELsmp

 Legend:
   option:   option that depends on the user environment is specified
```

**Figure 3.23 Example of Editing /etc/grub/grub.conf file**

Edit the shaded portion.

– When the boot loader is ELILO running on an IPF host

Figure 3.24 shows an example of editing the configuration file.

```
prompt
timeout=50
linear
#default=2.4.21-27.EL
default=2.4.21-27.EL-sd

image=vmlinuz-2.4.21-27.EL
    label=vmlinuz-2.4.21-27.EL
    initrd=initrd-2.4.21-27.EL.img
    read-only
    append="option root=LABEL=/ option"

image=vmlinuz-2.4.21-27.EL
    label=2.4.21-27.EL-sd
    initrd=initrd-2.4.21-27.EL.img
    read-only
    append="option"
    root=/dev/sda2
```

Legend:
option: option that depends on the user environment is specified

**Figure 3.24 Example of Editing /etc/elilo.conf file**

Edit the shaded portion.

9. If necessary, back up the boot loader configuration file.
10. Copy the startup configuration that has the `LABEL` specification.
11. Assign a name to the copied configuration.

    You can assign any name that indicates that the configuration is used for specifying the SCSI device.

    – If the boot loader is LILO, or ELILO running on an IPF host

      Change `label`.

    – If you will use GRUB as the boot loader

      Change `title`.

12. Convert the copied configuration from the `LABEL` specification to the SCSI device specification according to the correspondence between `LABEL` and the SCSI device that you checked in step 4.

    Edit the file to disable the `LABEL` specification and enable the SCSI device specification.

    When editing the file, be careful of the following:

    – When an option that depends on the user environment is specified:

      Be careful not to delete the option.

    – When the boot loader is LILO, or ELILO running on an IPF host, and an option that depends on the user environment is not specified:

      Comment out or delete `append`.

13. Specify the configuration to be used for booting the system.

– When the boot loader is LILO, or ELILO running on an IPF host

  Specify `label` for `default`. If LILO is the boot loader, you must then execute `/sbin/lilo` command to activate the new setting.

– When GRUB is the boot loader

  Specify the position of the valid `title` for `default`.

14. Restart the host to make sure that booting from the SCSI device is possible.

    This completes the removal of the `LABEL` setting.

15. Install HDLM.

    To install HDLM, see section 3.5.2 perform steps 2 and 3 to 11 to create an HDLM device. You will also need to perform step 12 if you are only restarting the host. Note that the procedure differs depending on whether the host is restarted after installation.

    In steps 16 to 33 that follow, you will edit the configuration file to change the definition so that the OS starts with HDLM specified.

16. Execute the HDLM-configuration definition utility (`dlmcfgmgr)` with the `-v` parameter specified to check the correspondence between the SCSI device and the HDLM device.

    Figure 3.25 shows an example of executing the `dlmcfgmgr` utility with the `-v` parameter specified.

```
# dlmcfgmgr -v
HDevName       Management   Device          Host    Channel Target  Lun
/dev/sddlmaa   configured   /dev/sda           0     0          0      0
KAPL10302-I /sbin/dlmcfgmgr completed normally.
```

**Figure 3.25 Example of Executing the dlmcfgmgr Utility with the -v Parameter Specified**

Check the shaded portion, which shows the correspondence between the SCSI device and the HDLM device.

17. Execute the `dlmmkinitrd` utility for supporting a boot disk to create an initial RAM disk image for HDLM.

    For the initial RAM disk image file specified in the parameter, assign a name that is different from all currently used names.

    For details on the `dlmmkinitrd` utility, see section 7.5.

    – When using Red Hat Enterprise Linux AS3/ES3

      When the `dlmmkinitrd` utility terminates normally, a message indicating the RAM disk size required for the created initial RAM disk image file appears. Write down the displayed value, since you will need it when you edit the boot loader configuration file. Figure 3.26 and Figure 3.27 show examples of executing the `dlmmkinitrd` utility.

```
# /opt/DynamicLinkManager/bin/dlmmkinitrd /boot/initrd-hdlm-2.4.21-32.0.1.EL.gz
`uname -r`
KAPL12329-I The utility for supporting the boot disk started.
KAPL12315-I A compressed initial ramdisk image /boot/initrd-hdlm-2.4.21-32.0.1.EL.gz
was created with ramdisk size = 8517 KB.
KAPL12330-I The utility for supporting the boot disk completed.
```

**Figure 3.26 Example of Executing the dlmmkinitrd Utility When Using Red Hat Enterprise Linux AS3/ES3 (IA32)**

```
# /opt/DynamicLinkManager/bin/dlmmkinitrd /boot/efi/efi/redhat/initrd-hdlm-2.4.21-
27.EL.img `uname -r`
KAPL12329-I The utility for supporting the boot disk started.
KAPL12315-I A compressed initial ramdisk image /boot/efi/efi/redhat/initrd-hdlm-2.4.21-
27.EL.img was created with ramdisk size = 8517 KB.
KAPL12330-I The utility for supporting the boot disk completed.
```

**Figure 3.27 Example of Executing the dlmmkinitrd Utility When Using Red Hat Enterprise Linux AS3/ES3 (IPF)**

The shaded portion indicates the RAM disk size.

– When using Red Hat Enterprise Linux AS4/ES4

Figure 3.28 and Figure 3.29 show examples of executing the dlmmkinitrd utility.

```
# /opt/DynamicLinkManager/bin/dlmmkinitrd /boot/initrd-hdlm-2.6.9-11.EL.gz `uname -r`
KAPL12329-I The utility for supporting the boot disk started.
KAPL12344-I A compressed initial ramdisk image /boot/initrd-hdlm-2.6.9-11.EL.gz was
created.
KAPL12330-I The utility for supporting the boot disk completed.
```

**Figure 3.28 Example of Executing the dlmmkinitrd Utility When Using Red Hat Enterprise Linux AS4/ES4 (IA32)**

```
# /opt/DynamicLinkManager/bin/dlmmkinitrd /boot/efi/efi/redhat/initrd-hdlm-2.6.9-11.EL.gz
`uname -r`
KAPL12329-I The utility for supporting the boot disk started.
KAPL12344-I A compressed initial ramdisk image /boot/initrd-hdlm-2.6.9-11.EL.gz was
created.
KAPL12330-I The utility for supporting the boot disk completed.
```

**Figure 3.29 Example of Executing the dlmmkinitrd Utility When Using Red Hat Enterprise Linux AS4/ES4 (IPF)**

18. Edit the `/etc/fstab` file.

    Change the definition in the mount point for HDLM management-target devices other than `/boot` directory or `/boot/efi` directory so that the HDLM device will be mounted rather than the SCSI device. Do not specify the `/boot` or `/boot/efi` directory as the mount point for the HDLM device. Also, change the definition so that the swap partition uses the HDLM device. Edit the `/etc/fstab` file as shown in Figure 3.30 or Figure 3.31.

The detailed procedure is explained in steps 19 to 22.

– When the boot loader is LILO or GRUB

  Figure 3.30 shows an example of editing the `/etc/fstab` file.

```
#LABEL=/         /            ext2          defaults                  1 1
#/dev/sda2       /            ext2          defaults                  1 1
/dev/sddlmaa2    /            ext2          defaults                  1 1
#LABEL=/boot     /boot        ext2          defaults                  1 2
/dev/sda1        /boot        ext2          defaults                  1 2
#LABEL=/tmp      /tmp         ext2          defaults                  1 2
#/dev/sda4       /tmp         ext2          defaults                  1 2
/dev/sddlmaa4    /tmp         ext2          defaults                  1 2
#LABEL=/var      /var         ext2          defaults                  1 2
#/dev/sda5       /var         ext2          defaults                  1 2
/dev/sddlmaa5    /var         ext2          defaults                  1 2
#LABEL=/usr      /usr         ext2          defaults                  1 2
#/dev/sda6       /usr         ext2          defaults                  1 2
/dev/sddlmaa6    /usr         ext2          defaults                  1 2
none             /dev/pts     devpts        gid=5,mode=620            0 0
none             /proc        proc          defaults                  0 0
none             /dev/shm     tmpfs         defaults                  0 0
#/dev/sda3       swap         swap          defaults                  0 0
/dev/sddlmaa3    swap         swap          defaults                  0 0
/dev/cdrom       /mnt/cdrom   udf,iso9660   noauto,owner,kudzu,ro     0 0
/dev/fd0         /mnt/floppy  auto          noauto,owner,kudzu        0 0
```

**Figure 3.30 Example of Editing /etc/fstab When the Boot Loader Is LILO or GRUB**

  Edit the shaded portion.

– When the boot loader is ELILO running on an IPF host

  Figure 3.31 shows an example of editing the `/etc/fstab` file.

```
#LABEL=/          /           ext3        defaults                 1 1
#/dev/sda2        /           ext3        defaults                 1 1
/dev/sddlmaa2     /           ext3        defaults                 1 1
/dev/sda1         /boot/efi   vfat        defaults                 0 0
#LABEL=/tmp       /tmp        ext3        defaults                 1 2
#/dev/sda4        /tmp        ext3        defaults                 1 2
/dev/sddlmaa4     /tmp        ext3        defaults                 1 2
#LABEL=/var       /var        ext3        defaults                 1 2
#/dev/sda5        /var        ext3        defaults                 1 2
/dev/sddlmaa5     /var        ext3        defaults                 1 2
#LABEL=/usr       /usr        ext3        defaults                 1 2
#/dev/sda6        /usr        ext3        defaults                 1 2
/dev/sddlmaa6     /usr        ext3        defaults                 1 2
none              /dev/pts    devpts      gid=5,mode=620           0 0
none              /proc       proc        defaults                 0 0
none              /dev/shm    tmpfs       defaults                 0 0
#/dev/sda3        swap        swap        defaults                 0 0
/dev/sddlmaa3     swap        swap        defaults                 0 0
/dev/cdrom        /mnt/cdrom  udf,iso9660 noauto,owner,kudzu,ro    0 0
```

**Figure 3.31 Example of Editing /etc/fstab When the Boot Loader Is ELILO**

  Edit the shaded portion.

19. If necessary, back up the file.
20. Copy the registration of the SCSI device specification.
21. Comment out the existing SCSI device specification by placing a hash mark (`#`) at the beginning of each line.
22. Convert the copied registration from a SCSI device specification to an HDLM device specification according to the correspondence between the SCSI device and the HDLM device that you checked in step 16.
23. Edit the boot loader configuration file so that the OS will start with HDLM specified.

    Edit the boot loader configuration file as shown in Figure 3.32, Figure 3.33, and Figure 3.34.

    The detailed procedure is explained in steps 24 to 29.

    Separate examples are provided for when the boot loader is LILO, when the boot loader is GRUB, and when the boot loader is ELILO running on an IPF host.

    – When the boot loader is LILO

      Figure 3.32 shows an example of editing the configuration file. After changing the configuration file, execute `/sbin/lilo` command to activate the new setting.

```
boot=/dev/sda
map=/boot/map
install=/boot/boot.b
prompt
timeout=50
linear
#default=linux
#default=linux-sd
default=HDLM_32.0.1.ELsmp

image=/boot/vmlinuz-2.4.21-20.ELsmp
    label=linux
    initrd=/boot/initrd-2.4.21-20.ELsmp
    read-only
    append="option root=LABEL=/ option"

image=/boot/vmlinuz-2.4.21-32,0,1.ELsmp
    label=linux-sd
    initrd=/boot/initrd-2.4.21-32.0.1.ELsmp
    read-only
   append="option"
    root=/dev/sda2

image=/boot/vmlinuz-2.4.21-32.0.1.ELsmp
    label=HDLM_32.0.1.ELsmp
    initrd=/boot/initrd-hdlm-2.4.21-32.0.1.EL.gz
    read-only
   append="opution ramdisk_size=8517"
#     root=/dev/sda2
```

Legend:
  option: option that depends on the user environment is specified

**Figure 3.32 Example of Editing /etc/lilo.conf file**

Edit the shaded portion.

– When the boot loader is GRUB

Figure 3.33 shows an example of editing the configuration file.

```
#default=0
#default=1
default=2
timeout=10
splashimage=(hd0,0)/grub/splash.xpm.gz

title Red Hat Enterprise Linux AS 3.0 (2.4.21-32.0.1.ELsmp)
   root (hd0,0)
   kernel /vmlinuz-2.4.21-32.0.1.ELsmp ro option root=LABEL=/ option
   initrd /initrd- 2.4.21-32.0.1.ELsmp

title Red Hat Enterprise Linux AS 3.0 (2.4.21-32.0.1.ELsmp-sd)
   root (hd0,0)
#   kernel /vmlinuz-2.4.21-32.0.1.ELsmp ro option root=LABEL=/ option
   kernel /vmlinuz-2.4.21-32.0.1.ELsmp ro option root=/dev/sda2 option
   initrd /initrd- 2.4.21-32.0.1.ELsmp

title Red Hat Enterprise Linux AS 3.0 (HDLM_32.0.1.ELsmp)
   root (hd0,0)
#   kernel /vmlinuz-2.4.21-32.0.1.ELsmp ro option root=LABEL=/ option
   kernel /vmlinuz-2.4.21-32.0.1.ELsmp ro option ramdisk_size=8517
   initrd /initrd-hdlm-2.4.21-32.0.1.EL.gz


Legend:
  option:  option that depends on the user environment is specified
```

**Figure 3.33 Example of Editing /etc/grub.conf file**

Edit the shaded portion.

– When the boot loader is ELILO running on an IPF host

Figure 3.34 shows an example of editing the configuration file.

```
prompt
timeout=50
#default=2.4.21-27.EL
#default=2.4.21-27.EL-sd
default=HDLM_2.4.21-27.EL

image=vmlinuz-2.4.21-27.EL
     label=2.4.21-27.EL
     initrd=initrd-2.4.21-27.EL.img
     read-only
     append="option root=LABEL=/ option"

image=/boot/vmlinuz-2.4.21-27.EL
     label=2.4.21-27.EL-sd
     initrd=initrd-2.4.21-27.EL.img
     read-only
     append="option"
     root=/dev/sda2

image=vmlinuz-2.4.21-27.EL
     label=HDLM_2.4.21-27.EL
     initrd=initrd-hdlm-2.4.21-27.EL.img
     read-only
     append="option ramdisk_size=1600"
#     root=/dev/sda2

Legend:
  option:  option that depends on the user environment is specified
```

**Figure 3.34 Example of Editing /etc/elilo.conf file**

Edit the shaded portion.

24. If necessary, back up the file.
25. Copy the configuration that was used for startup from the SCSI device.
26. Assign a name to the copied configuration that indicates that the configuration is used for specifying the HDLM device.
    - If the boot loader is LILO, or ELILO running on an IPF host

      Change `label`.
    - If you will use GRUB as the boot loader

      Change `title`.
27. Delete the `root` specification.
28. In `initrd`, specify the initial RAM disk image file you created in step 17, and specify the RAM disk size in `ramdisk_size` if Red Hat Enterprise Linux AS3/ES3 is being used.

    ***Note:*** how you specify `initrd` and `ramdisk_size` differs depending on the boot loader.

    When editing the file, be careful of the following:
    - When an option that depends on the user environment is specified:

      Be careful not to delete the option.
    - When the boot loader is LILO, or ELILO running on an IPF host, and `append` is deleted or commented out:

      Add `append` and specify `ramdisk_size`.

29. Specify the configuration that will be used for booting.
    - When the boot loader is LILO, or ELILO running on an IPF host

      Specify `label` for `default`. If LILO is the boot loader, you must then execute `/sbin/lilo` command to activate the new setting.
    - When GRUB is the boot loader

      Specify the position of the valid `title` for `default`.
30. Shut down the host.

    Execute the following command to shut down the host:

```
# shutdown -h now
```

31. Change the configuration from a single-path configuration to a multi-path configuration.
32. Start the host.
33. Make sure that the HDLM device is used by the file system and the swap partition.

    Make sure that the HDLM device is mounted and the HDLM device is assigned to the swap partition.
    - When the boot loader is LILO or GRUB

      Figure 3.35 shows an example of executing the `mount` command.

```
# mount
/dev/sddlmaa2 on / type ext2 (rw)
none on /proc type proc (rw)
none on /dev/pts type devpts (rw,gid=5,mode=620)
usbdevfs on /proc/bus/usb type usbdevfs (rw)
/dev/sda1 on /boot type ext2 (rw)
/dev/sddlmaa4 on /tmp type ext2 (rw)
/dev/sddlmaa5 on /var type ext2 (rw)
/dev/sddlmaa6 on /usr type ext2 (rw)
none on /dev/shm type tmpfs (rw)
```

**Figure 3.35 Example of Executing the mount Command When the Boot Loader Is LILO or GRUB**

Check the shaded portion.

- When the boot loader is ELILO running on an IPF host

  Figure 3.36 shows an example of executing the `mount` command.

```
# mount
/dev/sddlmaa2 on / type ext3 (rw)
none on /proc type proc (rw)
none on /dev/pts type devpts (rw,gid=5,mode=620)
usbdevfs on /proc/bus/usb type usbdevfs (rw)
/dev/sda1 on /boot/efi type vfat (rw)
/dev/sddlmaa4 on /tmp type ext3 (rw)
/dev/sddlmaa5 on /var type ext3 (rw)
/dev/sddlmaa6 on /usr type ext3 (rw)
none on /dev/shm type tmpfs (rw)
```

**Figure 3.36 Example of Executing the mount Command When the Boot Loader Is ELILO**

Check the shaded portion.

Refer to the `/proc/swaps` file to make sure that the specified HDLM device name is displayed in the `Filename` column.

```
Filename            Type       Size      Used     Priority
/dev/sddlmaa3       partition  6008300   0        -1
```

**Figure 3.37 Example Content of /proc/swaps file**

### 3.6.4 Upgrade Installation in an Environment Where an HDLM Device Is Used as a Boot Disk

This subsection describes how to perform an upgrade installation of HDLM in a multi-path boot disk environment that uses an HDLM device and how to set up the environment. Note that if settings are incorrect, the OS might not start. For details about what action to take if the OS cannot be started from an HDLM device, see section 3.6.7.

To perform such an installation:

1. Log in to Linux as a user with root privileges.
2. Make sure that the HDLM device is specified in the definition in the `/etc/fstab` file.

   Make sure that the mount points for HDLM management-target devices other than `/boot` directory and `/boot/efi` directory are defined so that the HDLM device will be mounted.

   – When the boot loader is LILO or GRUB

     Figure 3.38 shows example content of `/etc/fstab` file.

```
/dev/sddlmaa2   /             ext2          defaults                 1 1
/dev/sda1       /boot         ext2          defaults                 1 2
/dev/sddlmaa4   /tmp          ext2          defaults                 1 2
/dev/sddlmaa5   /var          ext2          defaults                 1 2
/dev/sddlmaa6   /usr          ext2          defaults                 1 2
none            /dev/pts      devpts        gid=5,mode=620           0 0
none            /proc         proc          defaults                 0 0
none            /dev/shm      tmpfs         defaults                 0 0
/dev/sddlmaa3   swap          swap          defaults                 0 0
/dev/cdrom      /mnt/cdrom/   udf,iso9660   noauto,owner,kudzu,ro    0 0
/dev/fd0        mnt/floppy    auto          noauto,owner,kudzu        0 0
```

**Figure 3.38 Example Content of /etc/fstab file When the Boot Loader Is LILO or GRUB**

   Check the shaded portion.

   – When the boot loader is ELILO running on an IPF host

     Figure 3.39 shows example content of `/etc/fstab` file.

```
/dev/sddlmaa2   /             ext3        defaults                 1 1
/dev/sda1       /boot/efi     vfat        defaults                 0 0
/dev/sddlmaa4   /tmp          ext3        defaults                 1 2
/dev/sddlmaa5   /var          ext3        defaults                 1 2
/dev/sddlmaa6   /usr          ext3        defaults                 1 2
none            /dev/pts      devpts      gid=5,mode=620           0 0
none            /proc         proc        defaults                 0 0
none            /dev/shm      tmpfs       defaults                 0 0
/dev/sddlmaa3   swap          swap        defaults                 0 0
/dev/cdrom      /mnt/cdrom    udf ,iso9660 noauto,owner,kudzu,ro   0 0
```

**Figure 3.39 Example Content of /etc/fstab file When the Boot Loader Is ELILO**

Check the shaded portion.

3. Execute the HDLM-configuration definition utility (`dlmcfgmgr`) with the `-v` parameter specified to check the correspondence between the HDLM device and the SCSI device.

   Figure 3.40 shows an example of executing the `dlmcfgmgr` utility with the `-v` parameter specified.

```
# dlmcfgmgr -v
HDevName      Management    Device      Host   Channel Target  Lun
/dev/sddlmaa  configured    /dev/sda       0     0         0     0
                            /dev/sdb       0     0         1     0
KAPL10302-I /sbin/dlmcfgmgr completed normally.
```

**Figure 3.40 Example of Executing the dlmcfgmgr Utility with the -v Parameter Specified**

Check the shaded portion.

4. Edit the `/etc/fstab` file.

   Change the mount point definition for HDLM management-target devices so that the SCSI device will be mounted rather than the HDLM device. Also, change the definition so that the swap partition uses the SCSI device.

   Edit the `/etc/fstab` file as shown in Figure 3.41 or Figure 3.42.

   The detailed procedure is explained in steps 5 to 7.

   – When the boot loader is LILO or GRUB

     Figure 3.41 shows an example of editing the `/etc/fstab` file.

```
#/dev/sddlmaa2   /             ext2          defaults                  1 1
/dev/sda2        /             ext2          defaults                  1 1
/dev/sda1        /boot         ext2          defaults                  1 2
#/dev/sddlmaa4   /tmp          ext2          defaults                  1 2
/dev/sda4        /tmp          ext2          defaults                  1 2
#/dev/sddlmaa5   /var          ext2          defaults                  1 2
/dev/sda5        /var          ext2          defaults                  1 2
#/dev/sddlmaa6   /usr          ext2          defaults                  1 2
/dev/sda6        /usr          ext2          defaults                  1 2
none             /dev/pts      devpts        gid=5,mode=620            0 0
none             /proc         proc          defaults                  0 0
none             /dev/shm      tmpfs         defaults                  0 0
#/dev/sddlmaa3   swap          swap          defaults                  0 0
/dev/sda3        swap          swap          defaults                  0 0
/dev/cdrom       /mnt/cdrom/   udf ,iso9660  noauto,owner,kudzu,ro     0 0
/dev/fd0         mnt/floppy    auto          noauto,owner,kudzu        0 0
```

**Figure 3.41 Example of Editing /etc/fstab file When the Boot Loader Is LILO or GRUB**

Edit the shaded portion.

– When the boot loader is ELILO running on an IPF host

Figure 3.42 shows an example of editing the `/etc/fstab` file.

```
#/dev/sddlmaa2   /             ext3          defaults                  1 1
/dev/sda2        /             ext3          defaults                  1 1
/dev/sda1        /boot/efi     vfat          defaults                  0 0
#/dev/sddlmaa4   /tmp          ext3          defaults                  1 2
/dev/sda4        /tmp          ext3          defaults                  1 2
#/dev/sddlmaa5   /var          ext3          defaults                  1 2
/dev/sda5        /var          ext3          defaults                  1 2
#/dev/sddlmaa6   /usr          ext3          defaults                  1 2
/dev/sda6        /usr          ext3          defaults                  1 2
none             /dev/pts      devpts        gid=5,mode=620            0 0
none             /proc         proc          defaults                  0 0
none             /dev/shm      tmpfs         defaults                  0 0
#/dev/sddlmaa3   swap          swap          defaults                  0 0
/dev/sda3        swap          swap          defaults                  0 0
/dev/cdrom       /mnt/cdrom    udf ,iso9660  noauto,owner,kudzu,ro     0 0
```

**Figure 3.42 Example of Editing /etc/fstab file When the Boot Loader Is ELILO**

Edit the shaded portion.

5. Copy the registration of the HDLM device specification.
6. Comment out the existing HDLM device specification by placing a hash mark (`#`) at the beginning of each line.
7. Convert the copied registration from an HDLM device specification to a SCSI device specification according to the correspondence between the HDLM device and the SCSI device that you checked in step 3.
8. Edit the boot loader configuration file so that the OS will start with a SCSI device specified.

   Edit the boot loader configuration file as shown in Figure 3.43, Figure 3.44, and Figure 3.45.

   The detailed procedure is explained in steps 9 to 12.

Separate examples are provided for when the boot loader is LILO, GRUB, or ELILO running on an IPF host.

Separate examples are provided for when the boot loader is LILO, when the boot loader is GRUB, and when the boot loader is ELILO running on an IPF host.

– When the boot loader is LILO

  Figure 3.43 shows an example of editing the configuration file. Specify the `label` setting in `default`, and then execute `/sbin/lilo` command to activate the new setting.

```
boot=/dev/sda
map=/boot/map
install=/boot/boot.b
prompt
timeout=50
linear
#default=HDLM_32.0.1.ELsmp
default=linux

image=/boot/vmlinuz-2.4.21-32.0.1.ELsmp
    label=HDLM_32.0.1.ELsmp
    initrd=/boot/initrd-hdlm-2.4.21-32.0.1.ELsmp.gz
    read-only
    append="option ramdisk_size=8517 option"

image=/boot/vmlinuz-2.4.21-32.0.1.ELsmp
    label=linux
    initrd=/boot/initrd-2.4.21-32.0.1.ELsmp
    read-only
    append="option"
    root=/dev/sda2

Legend:
  option:  option that depends on the user environment is specified
```

**Figure 3.43 Example of Editing /etc/lilo.conf file**

  Edit the shaded portion.

– When the boot loader is GRUB

  Figure 3.44 shows an example of editing the configuration file.

```
#default=0
default=1
timeout=10
splashimage=(hd0,0)/grub/splash.xpm.gz

title Red Hat Enterprise Linux AS 3.0 (HDLM-32.0.1.ELsmp)
   root (hd0,0)
   kernel /vmlinuz-2.4.21-32.0.1.ELsmp ro option ramdisk_size=8517 option
   initrd /initrd-hdlm-2.4.21-32.0.1.ELsmp.gz

title Red Hat Enterprise Linux AS 3.0 (2.4.21-32.0.1.ELsmp)
   root (hd0,0)
   kernel /vmlinuz-2.4.21-32.0.1.ELsmp ro root=/dev/sda2 option
   initrd /initrd- 2.4.21-32.0.1.ELsmp

 Legend:
   option:  option that depends on the user environment is specified
```

**Figure 3.44 Example of Editing /etc/grub.conf file**

Edit the shaded portion.

- When the boot loader is ELILO running on an IPF host

  Figure 3.45 shows an example of editing the configuration file.

```
prompt
timeout=50
#default= HDLM_2.4.21-27.EL
default= 2.4.21-27.EL-sd

image=vmlinuz-2.4.21-27.EL
     label=HDLM_2.4.21-27.EL
     initrd=initrd-hdlm-2.4.21-27.EL.img
     read-only
     append="option ramdisk_size=16000 option"

image=vmlinuz-2.4.21-27.EL
     label=2.4.21-27.EL-sd
     initrd=initrd-2.4.21-27.EL.img
     read-only
     append="option"
     root=/dev/sda2

Legend:
  option: option that depends on the user environment is specified
```

**Figure 3.45 Example of Editing /etc/elilo.conf file**

Edit the shaded portion.

9. Copy the configuration that was used for startup from the HDLM device.
10. Assign a name to the copied configuration.

    You can assign any name that indicates that the configuration is used for specifying the SCSI device.

    - If the boot loader is LILO, or ELILO running on an IPF host

      Change `label`.

If LILO is the boot loader, you must then execute `/sbin/lilo` command to activate the new setting.

- If you will use GRUB as the boot loader

  Change `title`.

11. In `root`, specify the SCSI device you checked in step 3.
12. In `initrd`, specify the initial RAM disk image file for the SCSI device. Delete ramdisk_size that was used for startup from the HDLM device if Red Hat Enterprise Linux AS3/ES3 is being used.

    When editing the file, be careful of the following:

    - When an option that depends on the user environment is specified:

      Be careful not to delete the option.

    - When the boot loader is LILO, or ELILO running on an IPF host, and an option that depends on the user environment is not specified:

      Comment out or delete `append`.

    Create the initial RAM disk image file for the SCSI device according to the `mkinitrd` command documentation for the OS.

13. Shut down the host.

    Execute the following command to restart the host:

```
# shutdown -h now
```

14. Change the configuration from a multi-path configuration to a single-path configuration.
15. Start the host.
16. Make sure that the SCSI device is used by the file system and the swap partition.

    Make sure that the SCSI device is mounted and the SCSI device is assigned to the swap partition.

    - When the boot loader is LILO or GRUB

      Figure 3.46 shows an example of executing the `mount` command.

```
# mount
/dev/sda2 on / type ext2 (rw)
none on /proc type proc (rw)
none on /dev/pts type devpts (rw,gid=5,mode=620)
usbdevfs on /proc/bus/usb type usbdevfs (rw)
/dev/sda1 on /boot type ext2 (rw)
/dev/sda4 on /tmp type ext2 (rw)
/dev/sda5 on /var type ext2 (rw)
/dev/sda6 on /usr type ext2 (rw)
none on /dev/shm type tmpfs (rw)
```

**Figure 3.46 Example of Executing the mount Command When the Boot Loader Is LILO or GRUB**

Check the shaded portion.

- When the boot loader is ELILO running on an IPF host

Figure 3.47 shows an example of executing the `mount` command.

```
# mount
/dev/sda2 on / type ext3 (rw)
none on /proc type proc (rw)
none on /dev/pts type devpts (rw,gid=5,mode=620)
usbdevfs on /proc/bus/usb type usbdevfs (rw)
/dev/sda1 on /boot/efi type vfat (rw)
/dev/sda4 on /tmp type ext3 (rw)
/dev/sda5 on /var type ext3 (rw)
/dev/sda6 on /usr type ext3 (rw)
none on /dev/shm type tmpfs (rw)
```

**Figure 3.47 Example of Executing the mount Command When the Boot Loader Is ELILO**

Check the shaded portion.

Refer to the `/proc/swaps` file to make sure that the specified SCSI device name is displayed in the `Filename` column.

17. Perform an upgrade installation of HDLM.

    To perform an upgrade installation of HDLM, see section 3.5.5 perform steps 2 and 3 to 11 to create an HDLM device.

18. Execute the HDLM-configuration definition utility (`dlmcfgmgr`) with the `-v` parameter specified to check the correspondence between the SCSI device and the HDLM device.

    Figure 3.48 shows an example of executing the `dlmcfgmgr` utility with the `-v` parameter specified.

```
# dlmcfgmgr -v
HDevName      Management   Device          Host    Channel Target  Lun
/dev/sddlmaa  configured   /dev/sda           0       0         0    0
KAPL10302-I /sbin/dlmcfgmgr completed normally.
```

**Figure 3.48 Example of Executing the dlmcfgmgr Utility with the -v Parameter Specified**

Check the shaded portion.

19. Execute the `dlmmkinitrd` utility for supporting a boot disk to create an initial RAM disk image for HDLM

    If the initial RAM disk image file to be created already exists, specify the `-f` option to overwrite the initial RAM disk image file.

    For details on the dlmmkinitrd utility, see section 7.5.

    – When using Red Hat Enterprise Linux AS3/ES3

      When the `dlmmkinitrd` utility terminates normally, a message indicating the RAM disk size required for the created initial RAM disk image file appears. Write down the displayed value, since you will need it when you edit the boot loader configuration file. Figure 3.49 and Figure 3.50 show examples of executing the `dlmmkinitrd` utility.

```
# /opt/DynamicLinkManager/bin/dlmmkinitrd -f /boot/initrd-hdlm-2.4.21-32.0.1.ELsmp.gz 'uname -r'
KAPL12329-I The utility for supporting the boot disk started.
KAPL12315-I A compressed initial ramdisk image /boot/initrd-hdlm-2.4.21-32.0.1.ELsmp.gz was
created with ramdisk size = 8517 KB.
KAPL12330-I The utility for supporting the boot disk completed.
```

**Figure 3.49 Example of Executing the dlmmkinitrd Utility When Using Red Hat Enterprise Linux AS3/ES3 (IA32)**

The shaded portion indicates the RAM disk size.

```
# /opt/DynamicLinkManager/bin/dlmmkinitrd /boot/efi/efi/redhat/initrd-hdlm-2.4.21-27.EL.gz `uname -r`
KAPL12329-I The utility for supporting the boot disk started.
KAPL12315-I A compressed initial ramdisk image /boot/initrd-hdlm-2.4.21-27.EL.gz was created with ramdisk
size = 8517 KB.
KAPL12330-I The utility for supporting the boot disk completed.
```

**Figure 3.50 Example of Executing the dlmmkinitrd Utility When Using Red Hat Enterprise Linux AS3/ES3 (IPF)**

The shaded portion indicates the RAM disk size.

- When using Red Hat Enterprise Linux AS4/ES4

  Figure 3.51 and Figure 3.52 show examples of executing the dlmmkinitrd utility.

```
# /opt/DynamicLinkManager/bin/dlmmkinitrd -f /boot/initrd-hdlm-2.6.9-11.EL.gz 'uname -r'
KAPL12329-I The utility for supporting the boot disk started.
KAPL12344-I A compressed initial ramdisk image /boot/initrd-hdlm-2.6.9-11.EL.gz was created.
KAPL12330-I The utility for supporting the boot disk completed.
```

**Figure 3.51 Example of Executing the dlmmkinitrd Utility When Using Red Hat Enterprise Linux AS4/ES4 (IA32)**

```
# /opt/DynamicLinkManager/bin/dlmmkinitrd /boot/efi/efi/redhat/initrd-hdlm-2.6.9-11.EL.gz
`uname -r`
KAPL12329-I The utility for supporting the boot disk started.
KAPL12344-I A compressed initial ramdisk image /boot/initrd-hdlm-2.6.9-11.EL.gz was
created.
KAPL12330-I The utility for supporting the boot disk completed.
```

**Figure 3.52 Example of Executing the dlmmkinitrd Utility When Using Red Hat Enterprise Linux AS4/ES4 (IPF)**

20. When you create a new initial RAM disk image file rather than overwriting the existing one, delete the existing initial RAM disk image file that is no longer required.

    Execute the following command:

    - When an IA32 host is used:

```
# rm /boot/initrd-hdlm-2.4.21-27.EL.gz
```

– When an IPF host is used:

```
# rm /boot/efi/efi/redhat/initrd-hdlm-2.4.21-27.EL.gz
```

21. Edit the boot loader configuration file so that the OS will start with HDLM specified.

    Edit the boot loader configuration file as shown in Figure 3.53, Figure 3.54, and Figure 3.55.

    The detailed procedure is explained in steps 22 to 25.

    Separate examples are provided for when the boot loader is LILO, when the boot loader is GRUB, and when the boot loader is ELILO running on an IPF host.

    – When the boot loader is LILO

      Figure 3.53 shows an example of editing the configuration file. Change the configuration file, and then execute `/sbin/lilo` command to activate the new setting.

```
boot=/dev/sda
map=/boot/map
install=/boot/boot.b
prompt
timeout=50
linear
#default=HDLM_32.0.1.ELsmp
#default=linux
default=HDLM_UP_32.0.1.ELsmp

image=/boot/vmlinuz-2.4.21-32.0.1.ELsmp
    label=HDLM_32.0.1.ELsmp
    initrd=/boot/initrd-hdlm-2.4.21-32.0.1.ELsmp.gz
    read-only
    append="option ramdisk_size=8517 option"

image=/boot/vmlinuz-2.4.21-32.0.1.ELsmp
    label=linux
    initrd=/boot/initrd-2.4.21-32.0.1.ELsmp
    read-only
    append="option"
    root=/dev/sda2

image=/boot/vmlinuz-2.4.21-32.0.1.ELsmp
    label=HDLM_UP_32.0.1.ELsmp
    initrd=/boot/initrd-hdlm-2.4.21-32.0.1.ELsmp.gz
    read-only
    append="option ramdisk_size=8517"
```

Legend:
option: option that depends on the user environment is specified

**Figure 3.53 Example of Editing /etc/lilo.conf file**

Edit the shaded portion.

– When the boot loader is GRUB

Figure 3.54 shows an example of editing the configuration file.

```
#default=0
#default=1
default=2
timeout=10
splashimage=(hd0,0)/grub/splash.xpm.gz

title Red Hat Enterprise Linux AS 3.0 (HDLM-32.0.1.ELsmp)
    root (hd0,0)
    kernel /vmlinuz-2.4.21-32.0.1.ELsmp ro option ramdisk_size=8517 option
    initrd /initrd-hdlm-2.4.21-32.0.1.ELsmp.gz

title Red Hat Enterprise Linux AS 3.0 (2.4.21-32.0.1.ELsmp)
    root (hd0,0)
    kernel /vmlinuz-2.4.21-32.0.1.ELsmp ro root=/dev/sda2 option
    initrd /initrd- 2.4.21-32.0.1.ELsmp

title Red Hat Enterprise Linux AS 3.0 (HDLM_UP_32.0.1.ELsmp)
    root (hd0,0)
    kernel /vmlinuz-2.4.21-32.0.1.ELsmp ro option ramdisk_size=8517
    initrd /initrd-hdlm-2.4.21-32.0.1.ELsmp.gz

Legend:
  option:  option that depends on the user environment is specified
```

**Figure 3.54 Example of Editing /etc/grub.conf file**

Edit the shaded portion.

- When the boot loader is ELILO running on an IPF host

  Figure 3.55 shows an example of editing the configuration file.

```
prompt
timeout=50
#default= HDLM_2.4.21-20.EL
default= HDLM_UP_2.4.21-27.EL

image=vmlinuz-2.4.21-27.EL
     label=HDLM_2.4.21-27.EL
     initrd=initrd-hdlm-2.4.21-27.EL.img
     read-only
     append="option ramdisk_size=16000 option"

image=vmlinuz-2.4.21-27.EL
     label=2.4.21-27.EL-sd
     initrd=initrd-2.4.21-27.EL.img
     read-only
     append="option"
     root=/dev/sda2

image=vmlinuz-2.4.21-27.EL
     label=HDLM_UP_2.4.21-27.EL
     initrd=initrd-hdlm-2.4.21-27.EL.img
     read-only
     append="option ramdisk_size=16000"

Legend:
  option:  option that depends on the user environment is specified
```

**Figure 3.55 Example of Editing /etc/elilo.conf file**

Edit the shaded portion.

22. Copy the configuration that was used for startup from the HDLM device.

23. Assign a name to the copied configuration that indicates that the configuration is used for specifying the HDLM device after update.
    - If the boot loader is LILO, or ELILO running on an IPF host

      Change `label`.
    - If you will use GRUB as the boot loader

      Change `title`.
24. In `initrd`, specify the initial RAM disk image file you created in step 19, and specify the RAM disk size in `ramdisk_size` if Red Hat Enterprise Linux AS3/ES3 is being used.

    ***Note***: how you specify `initrd` and `ramdisk_size` differs depending on the boot loader.

    When editing the file, be careful of the following:
    - When an option that depends on the user environment is specified:

      Be careful not to delete the option.
    - When the boot loader is LILO, or ELILO running on an IPF host, and `append` is deleted or commented out:

      Add `append` and specify `ramdisk_size`.

      If Red Hat Enterprise Linux AS4/ES4 is being used. you do not have to specify `ramdisk_size`.
25. Specify the configuration that will be used for booting.
    - When the boot loader is LILO, or ELILO running on an IPF host

      Specify `label` for `default`. If LILO is the boot loader, you must then execute `/sbin/lilo` command to activate the new setting.
    - When GRUB is the boot loader

      Specify the position of the valid `title` for `default`.
26. Edit the `/etc/fstab` file.

    Change the definitions of the mount points for HDLM management-target devices other than `/boot` directory and `/boot/efi` directory so that the HDLM device will be mounted rather than the SCSI device. Do not specify the /boot or /boot/efi directory as the mount point for the HDLM device. Also, change the definition so that the swap partition uses the HDLM device. Edit the `/etc/fstab` file as shown in Figure 3.56 and Figure 3.57.

    The detailed procedure is explained in steps 27 to 29.
    - When the boot loader is LILO or GRUB

      Figure 3.56 shows an example of editing `/etc/fstab` file.

```
#/dev/sddlmaa2  /              ext2           defaults                  1 1
#/dev/sda2      /              ext2           defaults                  1 1
/dev/sddlmaa2   /              ext2           defaults                  1 1
/dev/sda1       /boot          ext2           defaults                  1 2
#/dev/sddlmaa4  /tmp           ext2           defaults                  1 2
#/dev/sda4      /tmp           ext2           defaults                  1 2
/dev/sddlmaa4   /tmp           ext2           defaults                  1 2
#/dev/sddlmaa5  /var           ext2           defaults                  1 2
#/dev/sda5      /var           ext2           defaults                  1 2
/dev/sddlmaa5   /var           ext2           defaults                  1 2
#/dev/sddlmaa6  /usr           ext2           defaults                  1 2
#/dev/sda6      /usr           ext2           defaults                  1 2
/dev/sddlmaa6   /usr           ext2           defaults                  1 2
none            /dev/pts       devpts         gid=5,mode=620            0 0
none            /proc          proc           defaults                  0 0
none            /dev/shm       tmpfs          defaults                  0 0
#/dev/sddlmaa3  swap           swap           defaults                  0 0
#/dev/sda3      swap           swap           defaults                  0 0
/dev/sddlmaa3   swap           swap           defaults                  0 0
/dev/cdrom      /mnt/cdrom     udf,iso9660    noauto,owner,kudzu,ro     0 0
/dev/fd0        /mnt/floppy    auto           noauto,owner,kudzu        0 0
```

**Figure 3.56 Example of Editing /etc/fstab file When the Boot Loader Is LILO or GRUB**

Edit the shaded portion.

– When the boot loader is ELILO running on an IPF host

Figure 3.57 shows an example of editing `/etc/fstab` file.

```
#/dev/sddlmaa2   /           ext3         defaults                  1 1
#/dev/sda2       /           ext3         defaults                  1 1
/dev/sddlmaa2    /           ext3         defaults                  1 1
/dev/sda1        /boot/efi   vfat         defaults                  0 0
#/dev/sddlmaa4   /tmp        ext3         defaults                  1 2
#/dev/sda4       /tmp        ext3         defaults                  1 2
/dev/sddlmaa4    /tmp        ext3         defaults                  1 2
#/dev/sddlmaa5   /var        ext3         defaults                  1 2
#/dev/sda5       /var        ext3         defaults                  1 2
/dev/sddlmaa5    /var        ext3         defaults                  1 2
#/dev/sddlmaa6   /usr        ext3         defaults                  1 2
#/dev/sda6       /usr        ext3         defaults                  1 2
/dev/sddlmaa6    /usr        ext3         defaults                  1 2
none             /dev/pts    devpts       gid=5,mode=620            0 0
none             /proc       proc         defaults                  0 0
none             /dev/shm    tmpfs        defaults                  0 0
#/dev/sddlmaa3   swap        swap         defaults                  0 0
#/dev/sda3       swap        swap         defaults                  0 0
/dev/sddlmaa3    swap        swap         defaults                  0 0
/dev/cdrom       /mnt/cdrom  udf,iso9660  noauto,owner,kudzu,ro     0 0
/dev/fd0         /mnt/floppy auto         noauto,owner,kudzu        0 0
```

**Figure 3.57 Example of Editing /etc/fstab file When the Boot Loader Is ELILO**

Edit the shaded portion.

27. Copy the registration of the SCSI device specification.
28. Comment out the existing SCSI device specification by placing a hash mark (`#`) at the beginning of each line.

29. Convert the copied registration from a SCSI device specification to an HDLM device specification according to the correspondence between the SCSI device and the HDLM device that you checked in step 18.
30. Shut down the host.

    Execute the following command to shut down the host:

```
# shutdown -h now
```

31. Change the configuration from a single-path configuration to a multi-path configuration.
32. Start the host.
33. Make sure that the HDLM device is used by the file system and the swap partition.

    Make sure that the HDLM device is mounted and the HDLM device is assigned to the swap partition.

    – When the boot loader is LILO or GRUB

    Figure 3.58 shows an example of editing the `mount` command.

```
# mount
/dev/sddlmaa2 on / type ext2 (rw)
none on /proc type proc (rw)
none on /dev/pts type devpts (rw,gid=5,mode=620)
usbdevfs on /proc/bus/usb type usbdevfs (rw)
/dev/sda1 on /boot type ext2 (rw)
/dev/sddlmaa4 on /tmp type ext2 (rw)
/dev/sddlmaa5 on /var type ext2 (rw)
/dev/sddlmaa6 on /usr type ext2 (rw)
none on /dev/shm type tmpfs (rw)
```

**Figure 3.58 Example of Executing the mount Command When the Boot Loader Is LILO or GRUB**

Check the shaded portion.

– When the boot loader is ELILO

Figure 3.59 shows an example of executing the `mount` command.

```
# mount
/dev/sddlmaa2 on / type ext3 (rw)
none on /proc type proc (rw)
none on /dev/pts type devpts (rw,gid=5,mode=620)
usbdevfs on /proc/bus/usb type usbdevfs (rw)
/dev/sda1 on /boot/efi type vfat (rw)
/dev/sddlmaa4 on /tmp type ext3 (rw)
/dev/sddlmaa5 on /var type ext3 (rw)
/dev/sddlmaa6 on /usr type ext3 (rw)
none on /dev/shm type tmpfs (rw)
```

**Figure 3.59 Example of Executing the mount Command When the Boot Loader Is ELILO**

Check the shaded portion.

Refer to the `/proc/swaps` file to make sure that the specified HDLM device name is displayed in the `Filename` column.

```
Filename          Type       Size     Used    Priority
/dev/sddlmaa3     partition  6008300  0       -1
```

**Figure 3.60 Example Content of /proc/swaps file**

### 3.6.5 Settings for Using a Logical Volume (LVM2) on an HDLM Device as a Boot Disk

This subsection describes how to migrate a single-path boot disk environment that uses a logical volume (LVM2) on a SCSI device to a multi-path boot disk environment that uses a logical volume on an HDLM device. Note that if the settings are incorrect, the OS might not start.

In this subsection, the procedure assumes that `GRUB` or `ELILO` is used as the default boot loader in Red Hat Enterprise Linux.

To perform such an installation:

1. Log in to Linux as a user with root privileges.

   In an HDLM environment, a SCSI device cannot be identified from the value set in `LABEL`. Before installing HDLM, therefore, you must remove the `LABEL` setting. Perform steps 2 to check whether `LABEL` is set.

2. Check whether the definition in the `/etc/fstab` file contains `LABEL`.

   Check the contents of the `/etc/fstab` file.

   Figure 3.61 shows example content of `/etc/fstab` file.

```
/dev/VolGroup00/LogVol00   /          ext3    defaults          1 1
LABEL=/boot                /boot      ext3    defaults          1 2
none                       /dev/pts   devpts  gid=5,mode=620    0 0
none                       /dev/shm   tmpfs   defaults          0 0
none                       /proc      proc    defaults          0 0
none                       /sys       sysfs   defaults          0 0
/dev/VolGroup00/LogVol02   /tmp       ext3    defaults          1 2
/dev/VolGroup00/LogVol03   /var       ext3    defaults          1 2
/dev/VolGroup00/LogVol01   swap       swap    defaults          0 0
```

**Figure 3.61 Example Content of /etc/fstab file When using LVM2**

   `LABEL` is set in the shaded portion.

   This example assumes that IA32 uses `GRUB` as the boot loader. When you use IPF that uses `ELILO` as the boot loader, the mount point is displayed as `/boot/efi` instead of `/boot` (in the line where `/boot` is coded), and the file system type is displayed as `vfat` instead of `ext3`.

   If `LABEL` is not used in the `/etc/fstab` file, go to step 8 and install HDLM. If `LABEL` is used in the `/etc/fstab` file, perform steps 3 to 6 to check the mounted file system, and then remove the `LABEL` setting.

3. Execute the `mount` command to check the mounted file system.

Check the correspondence between the SCSI device and `LABEL` that you checked in steps 2. You will need this information to remove the `LABEL` setting.

Figure 3.62 shows an example of executing the mount command.

```
# mount
/dev/mapper/VolGroup00-LogVol00 on / type ext3 (rw)
none on /proc type proc (rw)
none on /sys type sysfs (rw)
none on /dev/pts type devpts (rw,gid=5,mode=620)
usbfs on /proc/bus/usb type usbfs (rw)
/dev/sda1 on /boot type ext3 (rw)
none on /dev/shm type tmpfs (rw)
/dev/mapper/VolGroup00-LogVol02 on /tmp type ext3 (rw)
/dev/mapper/VolGroup00-LogVol03 on /var type ext3 (rw)
none on /proc/sys/fs/binfmt_misc type binfmt_misc (rw)
sunrpc on /var/lib/nfs/rpc_pipefs type rpc_pipefs (rw)
```

**Figure 3.62 Example of Executing the mount Command**

The shaded portion indicates the SCSI device for which `LABEL` is set.

This example assumes that IA32 uses `GRUB` as the boot loader. When you use IPF that uses `ELILO` as the boot loader, the mount point is displayed as `/boot/efi` instead of `/boot` (in the line where `/boot` is coded), and the file system type is displayed as `vfat` instead of `ext3`.

Remove the `LABEL` setting in the `/etc/fstab` file according to the information for the mounted file system. Perform steps 4 to 6 to remove the `LABEL` setting in the `/etc/fstab` file.

4. Edit the `/etc/fstab` file to remove the `LABEL` setting.

   The detailed procedure is explained in steps 5 and 6.

   Figure 3.63 shows an example of editing the `/etc/fstab` file.

```
/dev/VolGroup00/LogVol00   /          ext3    defaults           1 1
#LABEL=/boot               /boot      ext3    defaults           1 2
/dev/sda1                  /boot      ext3    defaults           1 2
none                       /dev/pts   devpts  gid=5,mode=620     0 0
none                       /dev/shm   tmpfs   defaults           0 0
none                       /proc      proc    defaults           0 0
none                       /sys       sysfs   defaults           0 0
/dev/VolGroup00/LogVol02   /tmp       ext3    defaults           1 2
/dev/VolGroup00/LogVol03   /var       ext3    defaults           1 2
/dev/VolGroup00/LogVol01   swap       swap    defaults           0 0
```

**Figure 3.63 Example of Editing /etc/fstab file**

Comment out the `LABEL` line, and then add the shaded line.

This example assumes that IA32 uses `GRUB` as the boot loader. When you use IPF that uses `ELILO` as the boot loader, the mount point is displayed as `/boot/efi` instead of `/boot` (in the line where `/boot` is coded), and the file system type is displayed as `vfat` instead of `ext3`.

5. If necessary, back up the `/etc/fstab` file.

6. Edit the `/etc/fstab` file.

Refer to the correspondence between `LABEL` and the SCSI device that you checked in step 3.

Copy the registration of the `LABEL` specification.

Comment out the existing settings.

Edit the file to disable the `LABEL` specification and enable the SCSI device specification.

7. Restart the host to make sure that booting from the SCSI device is possible.

   This completes the removal of the `LABEL` setting.

8. Install HDLM.

   To install HDLM, see section 3.5.2, perform steps 2 and 3 to 11 to create an HDLM device. You will also need to perform step 12 if you are only restarting the host. Note that the procedure differs depending on whether the host is restarted after installation.

   In steps 9 to 22 that follow, you will edit the configuration file to change the definition so that the OS starts with HDLM specified.

9. Edit the `/etc/lvm/lvm.conf` file so that the LVM2 recognizes the HDLM device, not the SCSI device.

   Figure 3.64 shows examples of editing the `/etc/lvm/lvm.conf` file.

```
linux: /etc/lvm # vi lvm.conf
# This section allows you to configure which block devices should
# be used by the LVM system.
devices {

                    :
                    :

     # filter = [ "a/.*/" ]
       filter = [ "a|sddlm*|", "r|/dev/sd|" ]

                    :
                    :

     # types = [ "fd", 16 ]
       types = [ "sddlmfdrv", 16 ]

                    :
                    :

     md_component_detection=0
}
```

**Figure 3.64 Example of Editing /etc/fstab file**

Comment out the existing `filter` and `types` lines, and add the shaded lines in the figure.

Set 0 in `md_component_detection`.

Record the original `md_component_detection` value, since it is necessary for uninstallation or upgrade installation of HDLM.

10. Execute the `dlmmkinitrd` utility for supporting a boot disk to create an initial RAM disk image for HDLM.

For the initial RAM disk image file specified in the parameter, assign a name that is different from all currently used names.

For details on the `dlmmkinitrd` utility, see section 7.5.

Figure 3.65 and Figure 3.66 shows examples of executing the `dlmmkinitrd` utility.

```
# /opt/DynamicLinkManager/bin/dlmmkinitrd /boot/initrd-hdlm-2.6.9-11.EL.gz `uname -r`
KAPL12329-I The utility for supporting the boot disk started.
KAPL12344-I A compressed initial ramdisk image /boot/initrd-hdlm-2.6.9-11.EL.gz was
created.
KAPL12330-I The utility for supporting the boot disk completed.
```

**Figure 3.65 Example of Executing the dlmmkinitrd Utility When GRUB as the Boot Loader Is Used**

```
# /opt/DynamicLinkManager/bin/dlmmkinitrd /boot/efi/efi/redhat/initrd-hdlm-2.6.9-11.EL.gz
`uname -r`
KAPL12329-I The utility for supporting the boot disk started.
KAPL12344-I A compressed initial ramdisk image /boot/efi/efi/redhat/initrd-hdlm-2.6.9-
11.EL.gz was created.
KAPL12330-I The utility for supporting the boot disk completed.
```

**Figure 3.66 Example of Executing the dlmmkinitrd Utility When IPF (with ELILO as the Boot Loader) Is Used**

11. Edit the boot loader configuration file so that the OS will start with HDLM specified.

    Edit the boot loader configuration file as shown in Figure 3.67 and Figure 3.68.

    The detailed procedure is explained in steps 12 to 17.

    Figure 3.67 and Figure 3.68 shows an example of editing the configuration file.

```
#default=0
default=1
timeout=5
splashimage=(hd1,0)/grub/splash.xpm.gz
Hiddenmenu

title Red Hat Enterprise Linux AS (2.6.9-11.EL)
   root (hd1,0)
   kernel /vmlinuz-2.6.9-11.EL ro root=/dev/VolGroup00/
LogVol00 rhgb quiet
   initrd /initrd-2.6.9-11.EL.img

title Red Hat Enterprise Linux AS (HDLM 2.6.9-11.EL)
   root (hd1,0)
#   kernel /vmlinuz-2.6.9-11.EL ro root=/dev/VolGroup00/
LogVol00 rhgb quiet
   kernel /vmlinuz-2.6.9-11.EL ro rhgb quiet
#   initrd /initrd-2.6.9-11.EL.img
   initrd /boot/initrd-hdlm-2.6.9-11.EL.gz
```

**Figure 3.67 Example of Editing /etc/grub.conf file**

Edit the shaded portion.

```
prompt
timeout=10
#default=as4u1
default=as4u1-hdlm
relocatable

image=vmlinuz-2.6.9-11.EL
        label=as4u1
        initrd=initrd-2.6.9-11.EL.img
        read-only
        append="rhgb quiet root=/dev/VolGroup00/LogVol00"

image=vmlinuz-2.6.9-11.EL
#       label=as4u1
        label=as4u1-hdlm
#       initrd=initrd-2.6.9-11.EL.img
        initrd=initrd-hdlm-2.6.9-11.EL.gz
        read-only
#       append="rhgb quiet root=/dev/VolGroup00/LogVol00"
        append="rhgb quiet"
```

**Figure 3.68 Example of Editing /etc/elilo.conf file**

Edit the shaded portion.

12. If necessary, back up the file.
13. Copy the configuration that was used for startup from the SCSI device.
14. Assign a name to the copied configuration that indicates that the configuration is used for specifying the HDLM device.
    - If you will use GRUB as the boot loader

      Change `title`.
    - If you will use ELILO as the boot loader

      Change `label`.
15. Delete the `root` specification.
16. In initrd, specify the initial RAM disk image file you created in step 10.

    When an option that depends on the user environment is specified, be careful not to delete the option.
17. Specify the configuration that will be used for booting.
    - When GRUB is the boot loader

      Specify the position of the valid `title` for `default`.
    - When ELILO is the boot loader

      Specify `label` for `default`.
18. Shut down the host.

    Execute the following command to shut down the host:

```
# shutdown -h now
```

19. Change the configuration from a single-path configuration to a multi-path configuration.
20. Start the host.
21. Check the relationship between the volume group and HDLM devices.

    Execute the following command to make sure that the physical volumes that make up the volume group are HDLM devices.

```
# vgdisplay -v
    Finding all volume groups
    Finding volume group "VolGroup00"
  --- Volume group ---
  VG Name               VolGroup00
  System ID
  Format                lvm2
  Metadata Areas        1
  Metadata Sequence No  95
  VG Access             read/write
  VG Status             resizable
  MAX LV                0
  Cur LV                4
  Open LV               4
  Max PV                0
  Cur PV                1
  Act PV                1
  VG Size               19.88 GB
  PE Size               32.00 MB
  Total PE              636
  Alloc PE / Size       635 / 19.84 GB
  Free  PE / Size       1 / 32.00 MB
  VG UUID               Lt5T32-OOoC-d56S-KRQn-yt0e-wZuu-
xJreLx

  --- Logical volume ---
  LV Name                /dev/VolGroup00/LogVol00
  VG Name                VolGroup00
  LV UUID                Bo6xD9-vmBz-qIDF-Mw0o-grbd-Gvmh-
NHZlaX
  LV Write Access        read/write
  LV Status              available
  # open                 1
  LV Size                9.78 GB
  Current LE             313
  Segments               1
  Allocation             inherit
  Read ahead sectors     0
  Block device           252:0

                      :
                      :

   --- Physical volumes ---
  PV Name               /dev/sddlmaa2
  PV UUID               XmOok4-uPGl-Sirs-7Xx0-pVrm-nJrz-
ATm2pg
  PV Status             allocatable
  Total PE / Free PE    636 / 1
```

### 3.6.6 Upgrade Installation in an Environment Where a Logical Volume (LVM2) on an HDLM Device Is Used as a Boot Disk

This subsection describes how to perform an upgrade installation of HDLM in a multi-path boot disk environment that uses a logical volume (LVM2) on an HDLM device. Note that if settings are incorrect, the OS might not start.

In this subsection, the procedure assumes that GRUB or ELILO is used as the default boot loader in Red Hat Enterprise Linux.

To perform such an installation:

1. Log in to Linux as a user with root privileges.
2. Check the relationship between the volume group and HDLM devices.

   Execute the following command to make sure that the physical volumes that make up the volume group are HDLM devices.

```
# vgdisplay -v
    Finding all volume groups
    Finding volume group "VolGroup00"
  --- Volume group ---
  VG Name               VolGroup00
  System ID
  Format                lvm2
  Metadata Areas        1
  Metadata Sequence No  95
  VG Access             read/write
  VG Status             resizable
  MAX LV                0
  Cur LV                4
  Open LV               4
  Max PV                0
  Cur PV                1
  Act PV                1
  VG Size               19.88 GB
  PE Size               32.00 MB
  Total PE              636
  Alloc PE / Size       635 / 19.84 GB
  Free  PE / Size       1 / 32.00 MB
  VG UUID               Lt5T32-OOoC-d56S-KRQn-yt0e-wZuu-
xJreLx

  --- Logical volume ---
  LV Name                /dev/VolGroup00/LogVol00
  VG Name                VolGroup00
  LV UUID                Bo6xD9-vmBz-qIDF-Mw0o-grbd-Gvmh-
NHZlaX
  LV Write Access        read/write
  LV Status              available
  # open                 1
  LV Size                9.78 GB
  Current LE             313
  Segments               1
  Allocation             inherit
  Read ahead sectors     0
  Block device           252:0

                      :
                      :

   --- Physical volumes ---
  PV Name               /dev/sddlmaa2
  PV UUID               XmOok4-uPGI-Sirs-7Xx0-pVrm-nJrz-
ATm2pg
  PV Status             allocatable
  Total PE / Free PE    636 / 1
```

3. Edit the `/etc/lvm/lvm.conf` file so that the LVM2 recognizes the SCSI device, not the HDLM device.

   Figure 3.69 shows an example of editing the `/etc/lvm/lvm.conf` file.

```
linux: /etc/lvm # vi lvm.conf
# This section allows you to configure which block devices should
# be used by the LVM system.
devices {
                        :
                        :
     # filter = [ "a|sddlm*|", "r|/dev/sd|" ]
     filter = [ "r|sddlm*|", "a|/dev/sd|" ]
                        :
                        :
      types = [ "fd", 16 ]
      #types = [ "sddlmfdrv", 16 ]
                        :
                        :
      md_component_detection=0
}
```

**Figure 3.69 Example of Editing /etc/fstab file**

Comment out the existing `filter` and `types` lines, and add the shaded lines in the figure.

In addition, set `md_component_detection` to its original value before installation.

4. Edit the boot loader configuration file so that the OS will start with a SCSI device specified.

   Edit the boot loader configuration file as shown in Figure 3.70 and Figure 3.71.

   The detailed procedure is explained in steps 6 to 9.

   Figure 3.70 and Figure 3.71 shows an example of editing the configuration file.

```
#default=0
default=1
timeout=5
splashimage=(hd1,0)/grub/splash.xpm.gz
Hiddenmenu

title Red Hat Enterprise Linux AS (HDLM 2.6.9-11.EL)
   root (hd1,0)
   kernel /vmlinuz-2.6.9-11.EL ro rhgb quiet
   initrd /initrd-hdlm-2.6.9-11.EL.img

title Red Hat Enterprise Linux AS (2.6.9-11.EL)
   root (hd1,0)
   kernel /vmlinuz-2.6.9-11.EL ro root=/dev/VolGroup00/LogVol00 rhgb quiet
#    initrd /initrd-hdlm-2.6.9-11.EL.img
   initrd /initrd-2.6.9-11.EL.img
```

**Figure 3.70 Example of Editing /etc/grub.conf file**

Edit the shaded portion.

```
prompt
timeout=10
#default=as4u1-hdlm
default=as4u1
relocatable

image=vmlinuz-2.6.9-11.EL
        label=as4u1-hdlm
        initrd=initrd-hdlm-2.6.9-11.EL.gz
        read-only
        append="rhgb quiet"

image=vmlinuz-2.6.9-11.EL
#       label=as4u1-hdlm
        label=as4u1
#       initrd=initrd-hdlm-2.6.9-11.EL.gz
        initrd=initrd-2.6.9-11.EL.img
        read-only
#       append="rhgb quiet"
        append="rhgb quiet root=/dev/VolGroup00/LogVol00"
```

**Figure 3.71 Example of Editing /etc/elilo.conf file**

Edit the shaded portion.

5. Copy the configuration that was used for startup from the HDLM device.
6. Assign a name to the copied configuration.

   You can assign any name that indicates that the configuration is used for specifying the SCSI device.

   – If you will use GRUB as the boot loader

     Change `title`.

   – If you will use ELILO as the boot loader

     Change `label`.

7. Specify `root` as the name of the device mounted to `/` in the `/etc/fstab` file.
8. In `initrd`, specify the initial RAM disk image file for the SCSI device.

   When an option that depends on the user environment is specified, be careful not to delete the option.

   Create the initial RAM disk image file for the SCSI device according to the `mkinitrd` command documentation for the OS.

9. Specify the configuration that will be used for booting.

   – When GRUB is the boot loader

     Specify the position of the valid `title` for `default`.

   – When ELILO is the boot loader

     Specify `label` for `default`.

10. Shut down the host.

    Execute the following command to shut down the host:

```
# shutdown -h now
```

11. Change the configuration from a multi-path configuration to a single-path configuration.
12. Start the host.
13. Check the relationship between the volume group and HDLM devices.

    Execute the following command to make sure that the physical volumes that make up the volume group are not HDLM devices.

```
# vgdisplay -v
    Finding all volume groups
    Finding volume group "VolGroup00"
  --- Volume group ---
  VG Name               VolGroup00
  System ID
  Format                lvm2
  Metadata Areas        1
  Metadata Sequence No  95
  VG Access             read/write
  VG Status             resizable
  MAX LV                0
  Cur LV                4
  Open LV               4
  Max PV                0
  Cur PV                1
  Act PV                1
  VG Size               19.88 GB
  PE Size               32.00 MB
  Total PE              636
  Alloc PE / Size       635 / 19.84 GB
  Free  PE / Size       1 / 32.00 MB
  VG UUID               Lt5T32-OOoC-d56S-KRQn-yt0e-wZuu-
xJreLx

  --- Logical volume ---
  LV Name                /dev/VolGroup00/LogVol00
  VG Name                VolGroup00
  LV UUID                Bo6xD9-vmBz-qIDF-Mw0o-grbd-Gvmh-
NHZlaX
  LV Write Access        read/write
  LV Status              available
  # open                 1
  LV Size                9.78 GB
  Current LE             313
  Segments               1
  Allocation             inherit
  Read ahead sectors     0
  Block device           252:0

                    :
                    :

   --- Physical volumes ---
  PV Name               /dev/sda2
  PV UUID               XmOok4-uPGl-Sirs-7Xx0-pVrm-nJrz-
ATm2pg
  PV Status             allocatable
  Total PE / Free PE    636 / 1
```

14. Perform an upgrade installation of HDLM.

    To perform an upgrade installation of HDLM, see section 3.5.5 perform steps 2 and 3 to 11 to create an HDLM device.

15. Edit the `/etc/lvm/lvm.conf` file so that the LVM2 recognizes the HDLM device, not the SCSI device.

    Figure 3.72 shows examples of editing the `/etc/lvm/lvm.conf` file.

```
linux: /etc/lvm # vi lvm.conf
# This section allows you to configure which block devices should
# be used by the LVM system.
devices {

                        :
                        :

     # filter = [ "r|sddlm*|", "a|/dev/sd|" ]
     filter = [ "a|sddlm*|", "r|/dev/sd|" ]

                        :
                        :

     # types = [ "fd", 16 ]
     types = [ "sddlmfdrv", 16 ]

                        :
                        :

     md_component_detection=0
}
```

**Figure 3.72    Example of Editing /etc/fstab file**

Comment out the existing `filter` and `types` lines, and add the shaded lines in the figure.

Set 0 in `md_component_detection`.

Record the original `md_component_detection` value, since it is necessary for uninstallation or upgrade installation of HDLM.

16. Execute the `dlmmkinitrd` utility for supporting a boot disk to create an initial RAM disk image for HDLM

    If the initial RAM disk image file to be created already exists, specify the `-f` option to overwrite the initial RAM disk image file.

    For details on the `dlmmkinitrd` utility, see section 7.5.

    Figure 3.73 and Figure 3.74 shows examples of executing the `dlmmkinitrd` utility.

```
# /opt/DynamicLinkManager/bin/dlmmkinitrd -f /boot/initrd-hdlm-2.6.9-11.EL.gz 'uname -r'
KAPL12329-I The utility for supporting the boot disk started.
KAPL12344-I A compressed initial ramdisk image /boot/initrd-hdlm-2.6.9-11.EL.gz was created.
KAPL12330-I The utility for supporting the boot disk completed.
```

**Figure 3.73    Example of Executing the dlmmkinitrd Utility When GRUB as the Boot Loader Is Used**

```
# /opt/DynamicLinkManager/bin/dlmmkinitrd /boot/efi/efi/redhat/initrd-hdlm-2.6.9-11.EL.gz `uname -r`
KAPL12329-I The utility for supporting the boot disk started.
KAPL12344-I A compressed initial ramdisk image /boot/initrd-hdlm-2.6.9-11.EL.gz was created.
KAPL12330-I The utility for supporting the boot disk completed.
```

**Figure 3.74 Example of Executing the dlmmkinitrd Utility When IPF (with ELILO as the Boot Loader) Is Used**

17. When you create a new initial RAM disk image file rather than overwriting the existing one, delete the existing initial RAM disk image file that is no longer required.

    Execute the following command:

    – Execution example of the command when GRUB as the boot loader is used

    ```
    # rm /boot/initrd-hdlm-2.6.9-11.EL.gz
    ```

    – Execution example of the command when IPF (with ELILO as the boot loader) is used

    ```
    # rm /boot/efi/efi/redhat/initrd-hdlm-2.6.9-11.EL.gz
    ```

18. Edit the boot loader configuration file so that the OS will start with HDLM specified.

    Edit the boot loader configuration file as shown in Figure 3.75 and Figure 3.76.

    The detailed procedure is explained in steps 21 to 25.

    Figure 3.75 and Figure 3.76 shows an example of editing the configuration file.

```
#default=0
#default=1
default=2
timeout=5
splashimage=(hd1,0)/grub/splash.xpm.gz
Hiddenmenu

title Red Hat Enterprise Linux AS (HDLM 2.6.9-11.EL)
   root (hd1,0)
   kernel /vmlinuz-2.6.9-11.EL ro rhgb quiet
   initrd /initrd-hdlm-2.6.9-11.EL.img

title Red Hat Enterprise Linux AS (2.6.9-11.EL)
   root (hd1,0)
   kernel /vmlinuz-2.6.9-11.EL ro root=/dev/VolGroup00/LogVol00 rhgb quiet
#    initrd /initrd-hdlm-2.6.9-11.EL.img
   initrd /initrd-2.6.9-11.EL.img

title Red Hat Enterprise Linux AS (HDLM UP 2.6.9-11.EL)
   root (hd1,0)
   kernel /vmlinuz-2.6.9-11.EL ro rhgb quiet
   initrd /initrd-hdlm-2.6.9-11.EL.img
```

**Figure 3.75 Example of Editing /etc/grub.conf file**

Edit the shaded portion.

```
prompt
timeout=10
#default=as4u1-hdlm
#default=as4u1
default=as4u1-up-hdlm
relocatable

image=vmlinuz-2.6.9-11.EL
        label=as4u1-hdlm
        initrd=initrd-hdlm-2.6.9-11.EL.gz
        read-only
        append="rhgb quiet"

image=vmlinuz-2.6.9-11.EL
#       label=as4u1-hdlm
        label=as4u1
#       initrd=initrd-hdlm-2.6.9-11.EL.gz
        initrd=initrd-2.6.9-11.EL.img
        read-only
#       append="rhgb quiet"
        append="rhgb quiet root=/dev/VolGroup00/LogVol00"

image=vmlinuz-2.6.9-11.EL
#       label=as4u1-hdlm
#       label=as4u1
        label=as4u1-up-hdlm
#       initrd=initrd-hdlm-2.6.9-11.EL.gz
#       initrd=initrd-2.6.9-11.EL.img
        initrd=initrd-hdlm-2.6.9-11.EL.gz
        read-only
#       append="rhgb quiet"
#       append="rhgb quiet root=/dev/VolGroup00/LogVol00"
        append="rhgb quiet"
```

**Figure 3.76 Example of Editing /etc/elilo.conf file**

Edit the shaded portion.

19. Copy the configuration that was used for startup from the HDLM device.
20. Assign a name to the copied configuration that indicates that the configuration is used for specifying the HDLM device after update.
    - If you will use GRUB as the boot loader

      Change `title`.
    - If you will use ELILO as the boot loader

      Change `label`.
21. Delete the `root` specification.
22. In `initrd`, specify the initial RAM disk image file you created in step 18.

    When an option that depends on the user environment is specified, be careful not to delete the option.
23. Specify the configuration that will be used for booting.
    - When GRUB is the boot loader

      Specify the position of the valid `title` for `default`.
    - When ELILO is the boot loader

      Specify `label` for `default`.

24. Shut down the host.

    Execute the following command to shut down the host:

    ```
    # shutdown -h now
    ```

25. Change the configuration from a single-path configuration to a multi-path configuration.
26. Start the host.
27. Check the relationship between the volume group and HDLM devices.

    Execute the following command to make sure that the physical volumes that make up the volume group are HDLM devices.

```
# vgdisplay -v
    Finding all volume groups
    Finding volume group "VolGroup00"
  --- Volume group ---
  VG Name               VolGroup00
  System ID
  Format                lvm2
  Metadata Areas        1
  Metadata Sequence No  95
  VG Access             read/write
  VG Status             resizable
  MAX LV                0
  Cur LV                4
  Open LV               4
  Max PV                0
  Cur PV                1
  Act PV                1
  VG Size               19.88 GB
  PE Size               32.00 MB
  Total PE              636
  Alloc PE / Size       635 / 19.84 GB
  Free  PE / Size       1 / 32.00 MB
  VG UUID               Lt5T32-OOoC-d56S-KRQn-yt0e-wZuu-
xJreLx

  --- Logical volume ---
  LV Name                /dev/VolGroup00/LogVol00
  VG Name                VolGroup00
  LV UUID                Bo6xD9-vmBz-qIDF-Mw0o-grbd-Gvmh-
NHZlaX
  LV Write Access        read/write
  LV Status              available
  # open                 1
  LV Size                9.78 GB
  Current LE             313
  Segments               1
  Allocation             inherit
  Read ahead sectors     0
  Block device           252:0

                        :
                        :

   --- Physical volumes ---
  PV Name               /dev/sddlmaa2
  PV UUID               XmOok4-uPGl-Sirs-7Xx0-pVrm-nJrz-
ATm2pg
  PV Status             allocatable
  Total PE / Free PE    636 / 1
```

28. Check `IO-Count`.

    Execute the `sync` command, and then the `dlnkmgr view -path` command.

    Execute the above operation twice or more, and confirm that the `IO-Count` value increases after execution.

    In the following example, check the shaded portion.

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -path
Paths:000008 OnlinePaths:000008
PathStatus  IO-Count  IO-Errors
Online      46        0

PathID  PathName                        DskName                iLU   ChaPort Status  Type IO-Count IO-Errors DNum HDevName
000000  0000.0000.0000000000000000.0000 HITACHI .DF600F .0375 0017  0D      Online  Own  11       0         0    sddlmaa
000001  0000.0000.0000000000000000.0001 HITACHI .DF600F .0375 0033  0D      Online  Own  0        0         0    sddlmab
000002  0000.0000.0000000000000001.0000 HITACHI .DF600F .0375 0017  1D      Online  Non  0        0         0    sddlmaa
000003  0000.0000.0000000000000001.0001 HITACHI .DF600F .0375 0033  1D      Online  Non  0        0         0    sddlmab
000004  0001.0000.0000000000000000.0000 HITACHI .DF600F .0375 0017  0D      Online  Own  35       0         0    sddlmaa
000005  0001.0000.0000000000000001.0000 HITACHI .DF600F .0375 0017  1D      Online  Non  0        0         0    sddlmaa
000006  0001.0000.0000000000000000.0001 HITACHI .DF600F .0375 0033  0D      Online  Own  0        0         0    sddlmab
000007  0001.0000.0000000000000001.0001 HITACHI .DF600F .0375 0033  1D      Online  Non  0        0         0    sddlmab

KAPL01001-I The HDLM command completed normally. Operation name = view, completion time = yyyy/mm/dd  hh:mm:ss
#

# /opt/DynamicLinkManager/bin/dlnkmgr view -path
Paths:000008 OnlinePaths:000008
PathStatus  IO-Count  IO-Errors
Online      88        0

PathID  PathName                        DskName                iLU   ChaPort Status  Type IO-Count IO-Errors DNum
000000  0000.0000.0000000000000000.0000 HITACHI .DF600F .0375 0017  0D      Online  Own  23       0         0
000001  0000.0000.0000000000000000.0001 HITACHI .DF600F .0375 0033  0D      Online  Own  0        0         0
000002  0000.0000.0000000000000001.0000 HITACHI .DF600F .0375 0017  1D      Online  Non  0        0         0
000003  0000.0000.0000000000000001.0001 HITACHI .DF600F .0375 0033  1D      Online  Non  0        0         0
000004  0001.0000.0000000000000000.0000 HITACHI .DF600F .0375 0017  0D      Online  Own  65       0         0
000005  0001.0000.0000000000000001.0000 HITACHI .DF600F .0375 0017  1D      Online  Non  0        0         0
000006  0001.0000.0000000000000000.0001 HITACHI .DF600F .0375 0033  0D      Online  Own  0        0         0
000007  0001.0000.0000000000000001.0001 HITACHI .DF600F .0375 0033  1D      Online  Non  0        0         0

KAPL01001-I The HDLM command completed normally. Operation name = view, completion time = yyyy/mm/dd  hh:mm:ss
#
```

### 3.6.7 Countermeasures for Unsuccessful Startup of the OS from an HDLM Device

This subsection explains what actions you should take if the OS cannot be started from an HDLM device. If startup of the OS fails, either of the following happens:

- An error message appears and the OS stops.
- A prompt for a password for starting maintenance appears, and the OS stops.

The following explains the action to be taken for each case.

#### 3.6.7.1 When an Error Message Appears and the OS Stops

The probable causes of this error are as follows:

- The initial RAM disk image file was specified incorrectly in the boot loader settings.
- The initial RAM disk image file is corrupted.
- The `ramdisk_size` is set incorrectly in Red Hat Enterprise Linux AS3/ES3.

  If `ramdisk_size` was not specified in the boot loader settings, or the specified value is smaller than the value that was displayed when you created an initial RAM disk image for HDLM (by executing the `dlmmkinitrd` boot disk support utility), the following message appears:

```
RAMDISK: Compressed image found at block 0
Freeing initrd memory: 1994k freed
VFS: Mounted root (ext2 filesystem).
attempt to access beyond end of device
01:00: rw=0, want=3622, limit=1990
EXT2-fs error (device ramdisk(1,0)): ext2_read_inode: unable to read inode block
- inode=6508, block=3621
attempt to access beyond end of device
01:00: rw=0, want=6279, limit=1990
attempt to access beyond end of device
01:00: rw=0, want=6280, limit=1990
attempt to access beyond end of device
01:00: rw=0, want=6281, limit=1990
attempt to access beyond end of device
01:00: rw=0, want=6282, limit=1990
kmod: failed to exec /sbin/modprobe -s -k block-major-8, errno = 2
VFS: Cannot open root device "" or 08:03
Please append a correct "root=" boot option
Kernel panic: VFS: Unable to mount root fs on 08:03
```

- The startup disk was no longer recognized first because the system environment configuration was modified.

  If the startup disk is not found due to the changes in the system environment configuration, the following message appears:

```
KAPL10302-I Multiple instances of the utility for supporting the boot disk cannot
be executed concurrently.
KAPL10325-I The command started. Command name = /sbin/dlmcfgmgr -v
/usr/bin/expr: syntax error
/bin/echo: write error: Invalid argument
KAPL10328-I Execution of linuxrc completed.
/bin/mv: cannot move `/etc/dlmtemplog' to `/opt/DynamicLinkManager/Root/etc/opt/
DynamicLinkManager/hdlmboot.log': No such file or directory
VFS: Cannot open root device "" or 08:03
Please append a correct "root=" boot option
Kernel panic: VFS: Unable to mount root fs on 08:03
```

To correct the above error:

1. Restart the host.
2. In the window for setting information about boot loader startup, select the option for startup from a SCSI device.

   If startup from the SCSI device fails, enable only the path to the LU on which the OS is installed, and then start the OS.
3. Check the definitions in the `/etc/lilo.conf` file, `/etc/grub.conf` file, or `/etc/elilo.conf` file to make sure that the initial RAM disk image file and `ramdisk_size` are specified correctly. Modify any incorrect specification.

If startup from the HDLM device still fails after you correctly specify the initial RAM disk image file and `ramdisk_size`, the initial RAM disk image file might be corrupted. Execute the boot disk support utility to re-create the initial RAM disk image file.

4. Restart the host to start the OS from the HDLM device.

#### 3.6.7.2 When a Prompt for a Password for Starting Maintenance Appears, and the OS Stops

The probable causes of this error are as follows:

- Startup from a SCSI device was selected in the boot loader settings, but `/etc/fstab` file was defined to mount an HDLM device.

To correct the above error:

1. Enter the root user password.
2. Execute `mount -o remount rw /` to mount the root directory again.
3. In the `/etc/fstab` file, change the mount destination of `root` to the SCSI device.
4. Restart the host to start the OS from the SCSI device.
5. To restart the OS from an HDLM device, set up the environment according to the following procedure:

   Follow the procedure in 3.6.3 or 3.6.4.

## 3.7 Settings for LVM

This section explains LVM settings in an environment using Red Hat Enterprise Linux AS3/ES3, and LVM2 settings in an environment using SUSE LINUX Enterprise Server 9, SUSE LINUX Enterprise Server 10, Red Hat Enterprise Linux AS4/ES4, in the following cases:

Specify LVM settings:

- When using an HDLM device to create a new logical volume
- When moving a logical volume created on a SCSI device to an HDLM device
- When moving a logical volume created on md devices that use a SCSI device, to md devices that use an HDLM device

When using HDLM in an environment for LVMs other than LVM2, create md devices on the HDLM device, and use them as LVM physical volumes.

Specify LVM2 settings:

- When using an HDLM device to create a new logical volume
- When moving a logical volume created on a SCSI device in a single-path environment, to an HDLM device

When using HDLM in an LVM2 environment, use the HDLM device as an LVM2 physical volume.

When using Red Hat Enterprise Linux AS3/ES3 (for LVM setup), see section 3.7.1. When using SUSE LINUX Enterprise Server 9, SUSE LINUX Enterprise Server 10, Red Hat Enterprise Linux AS4/ES4 (for setting LVM2), see section 3.7.2.

***Note:***

- When changing an LVM environment to an LVM2 environment or when changing an LVM2 environment to an LVM environment, back up the disk and create an environment at the destination, and then restore the backup data.

### 3.7.1 When using Red Hat Enterprise Linux AS3/ES3

#### 3.7.1.1 Notes on Using LVM

Note the following when using LVM:

- If you do so, and then execute the `dlmcfgmgr` utility with the `-r` parameter specified, the error message `KAPL05023-E` is output to the syslog file. However, HDLM operations are not affected.
- If one LU exists as multiple physical volumes in one volume group, the last physical volume that is detected will be recognized as a physical disk because of the specification of the LVM `vgscan` command. When HDLM is uninstalled, change the settings for using an HDLM device to the settings for using a SCSI device.

- When you execute an LVM command after a logical volume that uses md devices is created, specify the name of the md device files. Do not specify the name of the logical device file of an HDLM-managed SCSI device or the logical device file name of an HDLM device.
- When multiple partitions used as physical volumes exist on the HDLM device, create an md device for each partition.
- When creating an md device on an HDLM device, in the `/etc/raidtab` file specify the following entries for the md device:
  - `raid-level`: `linear`
  - `chunk-size`: Any value (required)
  - `persistent-superblock`: `1`

  If you specify a partitioned HDLM device, make sure that there is a one-to-one correspondence between md devices and HDLM devices. See the following example.

  `/dev/md0`: `/dev/sddlmaa1`

  `/dev/md1`: `/dev/sddlmaa2`

  `/dev/md2`: `/dev/sddlmab1`

  `/dev/md3`: `/dev/sddlmab2`

  For the above md devices and HDLM devices, edit the `/etc/raidtab` file as shown below.

```
raiddev /dev/md0
raid-level      linear
chunk-size      32
nr-raid-disks   1
persistent-superblock 1
device          /dev/sddlmaa1
raid-disk       0
raiddev /dev/md1
raid-level      linear
chunk-size      32
nr-raid-disks   1
persistent-superblock 1
device          /dev/sddlmaa2
raid-disk       0
raiddev /dev/md2
raid-level      linear
chunk-size      32
nr-raid-disks   1
persistent-superblock 1
device          /dev/sddlmab1
raid-disk       0
raiddev /dev/md3
raid-level      linear
chunk-size      32
nr-raid-disks   1
persistent-superblock 1
device          /dev/sddlmab2
raid-disk       0
```

**Figure 3.77 Example of Editing the /etc/raidtab File**

- When you execute the `mkraid` command with the `-R` option specified, the command outputs the following message: `DESTROYING the contents of /dev/md0 in 5 seconds, Ctrl-C if unsure!` This message is output because the device management information that already includes the RAID management information is overwritten. However, no operating problem occurs even if the device management information is overwritten.

#### 3.7.1.2 When Using an HDLM Device to Create a New Logical Volume

This subsection explains the procedure for using an HDLM device to create a new logical volume. Use the following procedure to create the environment in Figure 3.78.

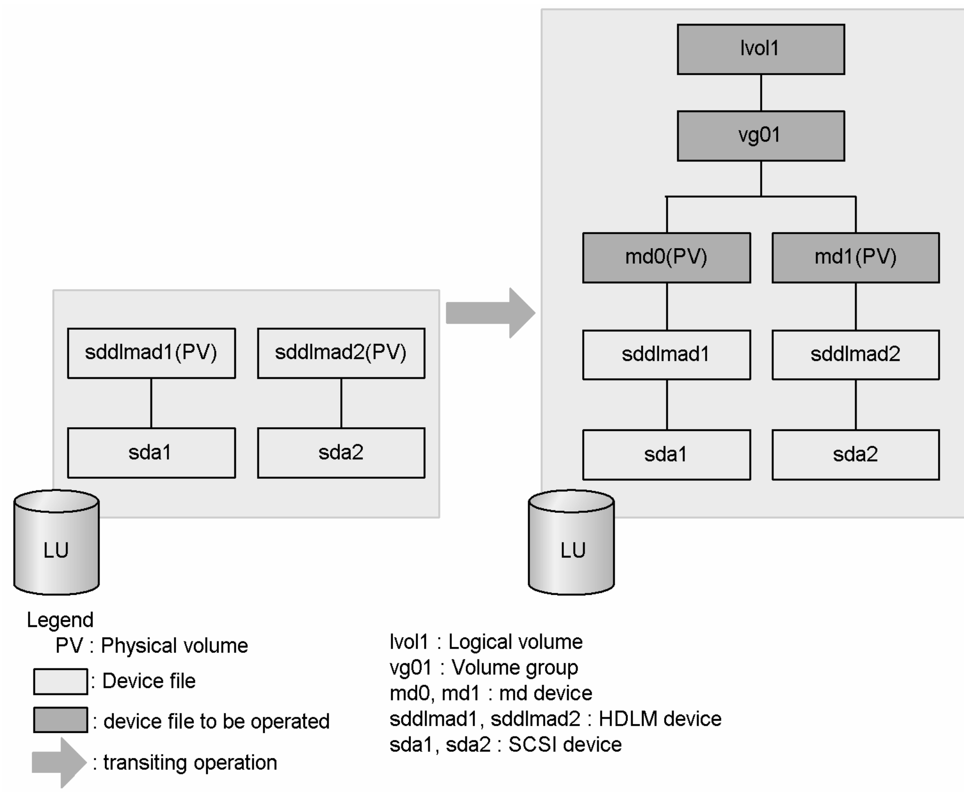


**Figure 3.78 Device Configuration When Creating a Logical Volume on an HDLM Device**

To create a new logical volume on an HDLM device:

1. Stop operation of the HDLM device.
2. Set the system ID for the disk partition to `Linux` (`83` in hexadecimal representation).
3. Edit the `/etc/raidtab` file.

   The following example shows entries when `/dev/md0` and `/dev/md1` are defined as md devices:

```
raiddev /dev/md0
raid-level      linear
chunk-size      32
nr-raid-disks   1
persistent-superblock 1
device          /dev/sddlmad1
raid-disk       0

raiddev /dev/md1
raid-level      linear
chunk-size      32
nr-raid-disks   1
persistent-superblock 1
device          /dev/sddlmad2
raid-disk       0
```

**Figure 3.79 Example of Editing the /etc/raidtab File**

4. Create the md devices on the HDLM device.

   The following example shows entries when /`dev/md0` and `/dev/md1` are created as md devices:

```
# mkraid -R /dev/md0 /dev/md1
DESTROYING the contents of /dev/md0 in 5 seconds, Ctrl-C if unsure!
handling MD device /dev/md0
analyzing super-block
disk 0: /dev/sddlmad1, 401593kB, raid superblock at 401472kB
DESTROYING the contents of /dev/md1 in 5 seconds, Ctrl-C if unsure!
handling MD device /dev/md1
analyzing super-block
disk 0: /dev/sddlmad2, 401625kB, raid superblock at 401536kB
```

   ***Note:***

   Always specify the `-R` option.

5. Execute the `vgscan` command.

   To re-create the volume group, execute the `vgscan` command as shown in the following example:

```
# vgscan
vgscan -- reading all physical volumes (this may take a while...)
vgscan -- "/etc/lvmtab" and "/etc/lvmtab.d" successfully created
vgscan -- WARNING: This program does not do a VGDA backup of your volume group
```

6. Create the physical volumes.

   In the following example, the command creates the physical volumes on `/dev/md0` and `/dev/md1`.

```
# pvcreate /dev/md0
pvcreate -- physical volume "/dev/md0" successfully created
# pvcreate /dev/md1
pvcreate -- physical volume "/dev/md1" successfully created
```

7. Create the volume group.

In the following example, the command creates the vg01 volume group by using the physical volumes `/dev/md0` and `/dev/md1`.

```
# vgcreate vg01 /dev/md0 /dev/md1
vgcreate -- INFO: using default physical extent size 4 MB
vgcreate -- INFO: maximum logical volume size is 255.99 Gigabyte
vgcreate -- doing automatic backup of volume group "vg01"
vgcreate -- volume group "vg01" successfully created and activated
```

8. Deactivate the volume group.

   In the following example, the logical volume subject to the migration belongs to the `vg01` volume group:

```
# vgchange -an vg01
vgchange -- volume group "vg01" successfully deactivated
```

9. Stop the md devices.

   In the following example, the commands stop the md devices `/dev/md0` and `/dev/md1`.

```
# raidstop /dev/md0
# raidstop /dev/md1
```

10. Start the md devices.

    In the following example, the commands start the md devices `/dev/md0` and `/dev/md1`.

```
# raidstart /dev/md0
# raidstart /dev/md1
```

11. Activate the volume group.

    In the following example, the logical volume subject to the change belongs to the `vg01` volume group:

```
# vgchange -ay vg01
vgchange -- volume group "vg01" successfully activated
```

12. Create the logical volume.

    In the following example, the command creates the logical volume (lvol1: 100 MB) by using the `vg01` volume group.

```
# lvcreate -L 100M -n lvol1 vg01
vgchange -- volume group "vg01" successfully activated
```

13. Create a file system.

```
# mkfs -t file-system-type/dev/vg01/lvol1
```

14. Mount the logical volume.

```
# mount /dev/vg01/lvol1 mount-point
```

### 3.7.1.3 When Moving a Logical Volume Created on a SCSI Device to an HDLM Device

This subsection explains the procedure for moving a logical volume created on a SCSI device to an HDLM device. Use the following procedure to create the environment in Figure 3.80.

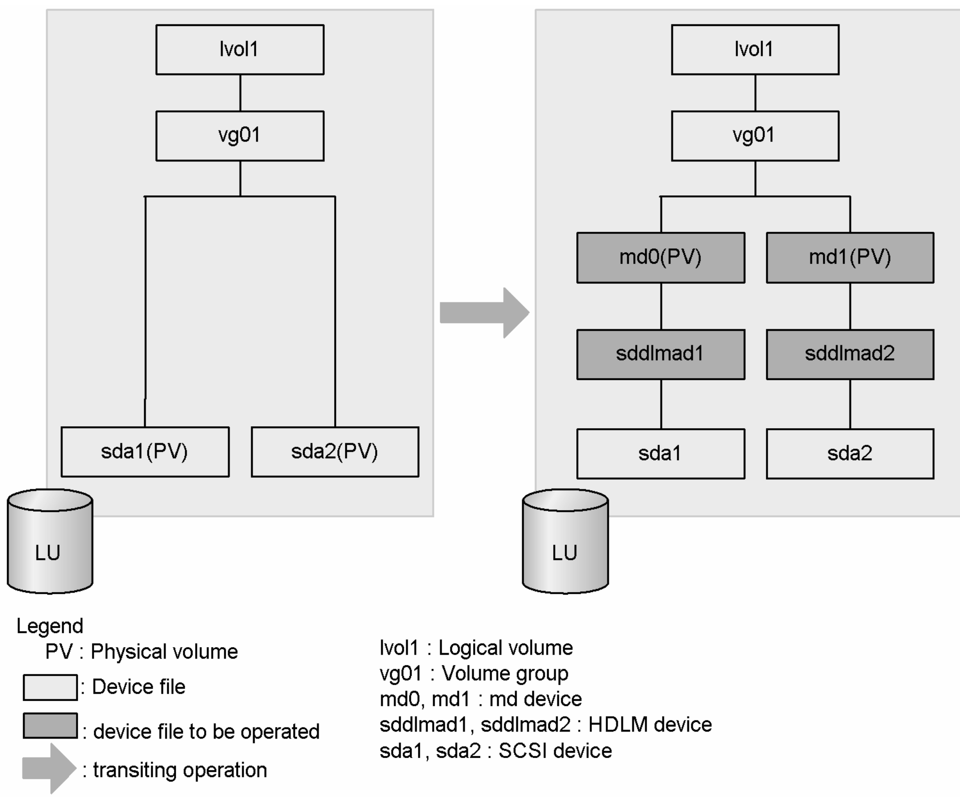


**Figure 3.80 Device Configuration Where the Logical Volume on a SCSI Device Is Migrated to an HDLM Device**

To move the logical volume to an HDLM device:

1. Stop operation of the SCSI and HDLM devices.
2. Back up the data on the logical volume that you want to migrate to an HDLM device.
3. Unmount the logical volume to be migrated.

   In the following example, the logical volume subject to the move has been mounted on `/mnt/lvol1`:

```
# umount /mnt/lvol1
```

4. Deactivate the volume group.

   In the following example, the logical volume subject to the move belongs to the vg01 volume group:

```
# vgchange -an vg01
vgchange -- volume group "vg01" successfully deactivated
```

5. Export the volume group.

   In the following example, the command exports information about the vg01 volume group:

```
# vgexport vg01
vgexport -- volume group "vg01" successfully exported
```

6. Change the system ID for the disk partition from `Linux LVM` (`8e` in hexadecimal representation) to `Linux` (`83` in hexadecimal representation).

   Example of executing the fdisk command (IA32)

```
# fdisk /dev/sddlmad

Command (m for help): p

Disk /dev/sddlmaa: 1075 MB, 1075314688 bytes
255 heads, 63 sectors/track, 130 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes

Device Boot    Start    End    Blocks   Id  System
/dev/sddlmaa1  1        65     522081   8e  Linux LVM
/dev/sddlmaa2  66       130    522112+  8e  Linux LVM

Command (m for help): t
Partition number (1-4): 1
Hex code (type L to list codes): 83
Changed system type of partition 1 to 83 (Linux)

Command (m for help): t
Partition number (1-4): 2
Hex code (type L to list codes): 83
Changed system type of partition 2 to 83 (Linux)

Command (m for help): p

Disk /dev/sddlmad: 1075 MB, 1075314688 bytes
255 heads, 63 sectors/track, 130 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes

Device Boot    Start    End    Blocks   Id  System
/dev/sddlmad1  1        65     522081   83  Linux
/dev/sddlmad2  66       130    522112+  83  Linux

Command (m for help): w
The partition table has been altered!
Calling ioctl() to re-read partition table.
Syncing disks.
```

   Example of executing the parted command (IA32, IPF or EM64T/AMD64)

```
# parted /dev/sddlmad set 1 lvm off
# parted /dev/sddlmad set 2 lvm off

Check partition type with print option.
# parted /dev/sddlmaa print
Disk geometry for /dev/sddlmaa: 0.000-1025.500 megabytes
Disk label type: msdos
Minor    Start     End       Type      Filesystem  Flags
1        0.031     509.875   primary
2        509.875   1019.750  primary
Information: Don't forget to update /etc/fstab, if
necessary.

Update the /proc/partitions using the blockdev command.

# blockdev --rereadpt /dev/sddlmad
```

7. Edit the `/etc/raidtab` file.

   In the following example, `/dev/md0` and `/dev/md1` are defined as md devices:

```
raiddev /dev/md0
raid-level      linear
chunk-size      32
nr-raid-disks   1
persistent-superblock 1
device          /dev/sddlmad1
raid-disk       0

raiddev /dev/md1
raid-level      linear
chunk-size      32
nr-raid-disks   1
persistent-superblock 1
device          /dev/sddlmad2
raid-disk       0
```

**Figure 3.81 Example of Editing the /etc/raidtab File**

8. Create the md devices on the HDLM device.

   In the following example, the command creates `/dev/md0` and `/dev/md1` as md devices:

```
# mkraid -R /dev/md0 /dev/md1
DESTROYING the contents of /dev/md0 in 5 seconds, Ctrl-C if unsure!
handling MD device /dev/md0
analyzing super-block
disk 0: /dev/sddlmad1, 401593kB, raid superblock at 401472kB
DESTROYING the contents of /dev/md1 in 5 seconds, Ctrl-C if unsure!
handling MD device /dev/md1
analyzing super-block
disk 0: /dev/sddlmad2, 401625kB, raid superblock at 401536kB
```

   ***Note:***

   Always specify the `-R` option.

9. Execute the `vgscan` command.

To re-create the volume group, execute the `vgscan` command as shown in the following example:

```
# vgscan
vgscan -- reading all physical volumes (this may take a while...)
vgscan -- found exported volume group "vg01PV_EXP"
vgscan -- "/etc/lvmtab" and "/etc/lvmtab.d" successfully created
vgscan -- WARNING: This program does not do a VGDA backup of your volume group
```

10. Import the volume group.

    In the following example, the command imports the `vg01` volume group:

```
# vgimport vg01 /dev/md0 /dev/md1
vgimport -- doing automatic backup of volume group "vg01"
vgimport -- volume group "vg01" successfully imported and activated
```

11. Deactivate the volume group.

    In the following example, the logical volume subject to the migration belongs to the `vg01` volume group:

```
# vgchange -an vg01
vgchange -- volume group "vg01" successfully deactivated
```

12. Stop the md devices.

    In the following example, the commands stop the md devices `/dev/md0` and `/dev/md1`.

```
# raidstop /dev/md0
# raidstop /dev/md1
```

13. Start the md devices.

    In the following example, the commands start the md devices `/dev/md0` and `/dev/md1`.

```
# raidstart /dev/md0
# raidstart /dev/md1
```

14. Activate the volume group.

    In the following example, the logical volume subject to the migration belongs to the `vg01` volume group:

```
# vgchange -ay vg01
vgchange -- volume group "vg01" successfully activated
```

15. Execute the `pvscan` command to check that the environment has been changed successfully.

    To check that the information about the physical volume for the target logical volume has been changed to the information about the md devices created on the HDLM device, execute the following command:

```
# pvscan
pvscan -- reading all physical volumes (this may take a while...)
pvscan -- WARNING: physical volume "/dev/sddlmad1" belongs to a meta device
pvscan -- WARNING: physical volume "/dev/sddlmaD2" belongs to a meta device
```

```
pvscan -- WARNING: physical volume "/dev/sda1" belongs to a meta device
pvscan -- WARNING: physical volume "/dev/sda2" belongs to a meta device
pvscan -- ACTIVE   PV "/dev/md0"      of VG "vg01" [384 MB / 284 MB free]
pvscan -- ACTIVE   PV "/dev/md1"      of VG "vg01" [384 MB / 384 MB free]
pvscan -- total: 12 [784.19 MB] / in use: 12 [784.19 MB] / in no VG: 0 [0]
```

16. Mount the logical volume that was moved.

    In the following example, the logical volume that was migrated is `/dev/vg01/lvol1` and the command mounts the volume on `/mnt/lvol1`:

```
# mount /dev/vg01/lvol1 /mnt/lvol1
```

17. Move the data that you backed up in step 2 to the mounted logical volume.

#### 3.7.1.4 When Migrating a Logical Volume Created on md Devices That Use a SCSI Device to md Devices That Use an HDLM Device

This subsection explains the procedure for moving a logical volume created on md devices that use a SCSI device, to md devices that use an HDLM device. Use the following procedure to create the environment in Figure 3.82.

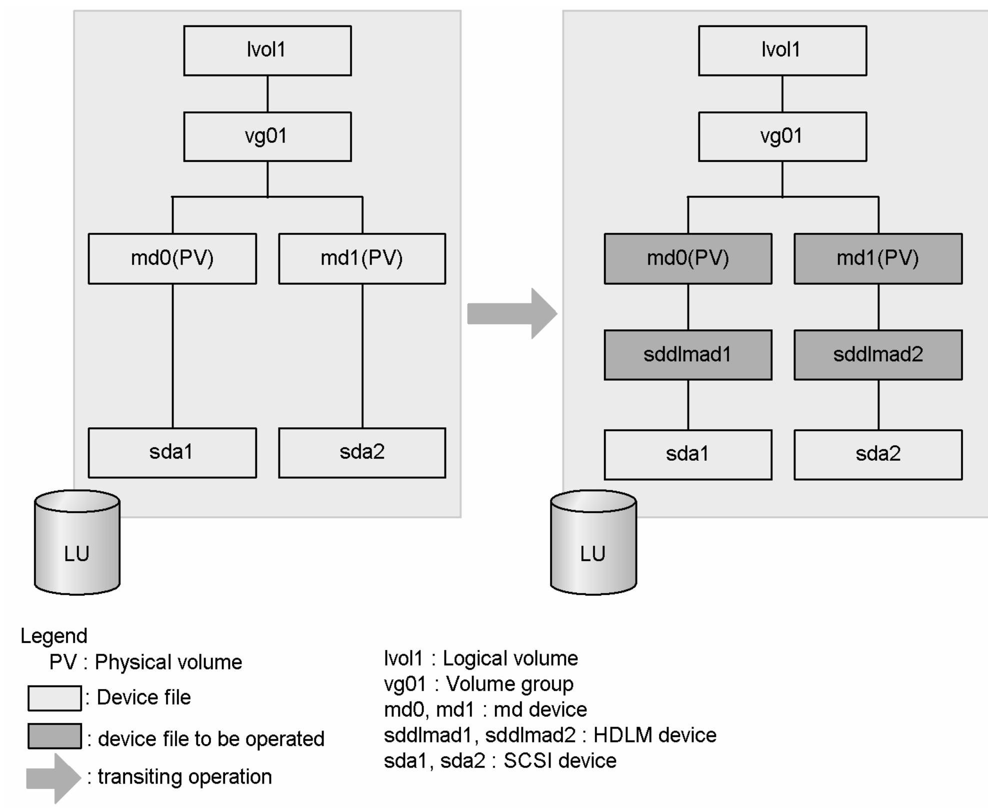


**Figure 3.82 Device Configuration Where the Logical Volume Created on md Devices That Use a SCSI Device Is Moved to md Devices That Use an HDLM Device**

To move the logical volume to md devices that use an HDLM device:

1. Stop operation of the SCSI, md, and HDLM devices.
2. Unmount the logical volume to be migrated.

   In the following example, the logical volume subject to the move has been mounted on `/mnt/lvol1`:

```
# umount /mnt/lvol1
```

3. Deactivate the volume group.

   In the following example, the logical volume subject to the move belongs to the vg01 volume group:

```
# vgchange -an vg01
vgchange -- volume group "vg01" successfully deactivated
```

4. Export the volume group.

   In the following example, the command exports information about the vg01 volume group:

```
# vgexport vg01
vgexport -- volume group "vg01" successfully exported
```

5. Stop the md devices.

   In the following example, the commands stop the md devices `/dev/md0` and `/dev/md1`.

```
# raidstop /dev/md0
# raidstop /dev/md1
```

6. Edit the `/etc/raidtab` file.

   Edit the `/etc/raidtab` file in order to create the md devices on the HDLM device. To create the md devices, modify the entries for the already-registered SCSI device. The following shows an example of editing the `/etc/raidtab` file:

```
raiddev /dev/md0
raid-level      linear
chunk-size      32
nr-raid-disks   1
persistent-superblock 1
device          /dev/sda1
raid-disk       0
raiddev /dev/md1
raid-level      linear
chunk-size      32
nr-raid-disks   1
persistent-superblock 1
device          /dev/sda2
raid-disk       0
```

Editing

```
raiddev /dev/md0
raid-level      linear
chunk-size      32
nr-raid-disks   1
persistent-superblock 1
device          /dev/sddlmad1
raid-disk       0
raiddev /dev/md1
raid-level      linear
chunk-size      32
nr-raid-disks   1
persistent-superblock 1
device          /dev/sddlmad2
raid-disk       0
```

**Figure 3.83 Example of Editing the /etc/raidtab File When the md devices Are Created on an HDLM Device**

7. Create the md devices on the HDLM device.

   In the following example, the command creates `/dev/md0` and `/dev/md1` md devices:

```
# mkraid -R /dev/md0 /dev/md1
DESTROYING the contents of /dev/md0 in 5 seconds, Ctrl-C if unsure!
handling MD device /dev/md0
analyzing super-block
disk 0: /dev/sddlmaa1, 401593kB, raid superblock at 401472kB
DESTROYING the contents of /dev/md1 in 5 seconds, Ctrl-C if unsure!
handling MD device /dev/md1
analyzing super-block
disk 0: /dev/sddlmaa2, 401625kB, raid superblock at 401536kB
```

   ***Note:***

   Always specify the `-R` option.

8. Execute the `vgscan` command.

   To re-create the volume group, execute the `vgscan` command as shown in the following example:

```
# vgscan
vgscan -- reading all physical volumes (this may take a while...)
vgscan -- found exported volume group "vg01PV_EXP"
vgscan -- "/etc/lvmtab" and "/etc/lvmtab.d" successfully created
vgscan -- WARNING: This program does not do a VGDA backup of your volume group
```

9. Import the volume group.

   In the following example, the command imports the `vg01` volume group:

```
# vgimport vg01 /dev/md0 /dev/md1
vgimport -- doing automatic backup of volume group "vg01"
vgimport -- volume group "vg01" successfully imported and activated
```

10. Deactivate the volume group.

    In the following example, the logical volume subject to the move belongs to the `vg01` volume group:

```
# vgchange -an vg01
vgchange -- volume group "vg01" successfully deactivated
```

11. Stop the md devices.

    In the following example, the commands stop the md devices `/dev/md0` and `/dev/md1`:

```
# raidstop /dev/md0
# raidstop /dev/md1
```

12. Start the md devices.

    In the following example, the commands start the md devices `/dev/md0` and `/dev/md1`:

```
# raidstart /dev/md0
# raidstart /dev/md1
```

13. Activate the volume group.

In the following example, the logical volume subject to the move belongs to the `vg01` volume group:

```
# vgchange -ay vg01
vgchange -- volume group "vg01" successfully activated
```

14. Execute the `pvscan` command to check that the environment has been changed successfully.

    To check that the information about the physical volume for the target logical volume has been changed to the information about the md devices created on the HDLM device, execute the following command:

```
# pvscan
pvscan -- reading all physical volumes (this may take a while...)
pvscan -- WARNING: physical volume "/dev/sddlmad1" belongs to a meta device
pvscan -- WARNING: physical volume "/dev/sddlmad2" belongs to a meta device
pvscan -- WARNING: physical volume "/dev/sda1" belongs to a meta device
pvscan -- WARNING: physical volume "/dev/sda2" belongs to a meta device
pvscan -- ACTIVE   PV "/dev/md0"      of VG "vg01" [384 MB / 284 MB free]
pvscan -- ACTIVE   PV "/dev/md1"      of VG "vg01" [384 MB / 384 MB free]
pvscan -- total: 12 [784.19 MB] / in use: 12 [784.19 MB] / in no VG: 0 [0]
```

15. Mount the logical volume that was moved.

    In the following example, the logical volume that was migrated is `/dev/vg01/lvol1` and the command mounts the volume on `/mnt/lvol1`:

```
# mount /dev/vg01/lvol1 /mnt/lvol1
```

### 3.7.2 When using SUSE LINUX Enterprise Server 9, SUSE LINUX Enterprise Server 10, Red Hat Enterprise Linux AS4/ES4

#### 3.7.2.1 Notes on Using Logical Volume Manager (LVM2)

A SCSI device and an HDLM device cannot be used concurrently as an LVM2 physical volume. When LVM2 is used in an environment in which HDLM is installed, create a physical volume, volume group, and logical volume by using an HDLM device. Note that when HDLM is uninstalled, change the settings for using an HDLM device to the settings for using a SCSI device.

#### 3.7.2.2 When Using an HDLM Device to Create a New Logical Volume

This subsection explains the procedure for using an HDLM device to create a new logical volume. Use the following procedure to create the environment shown in Figure 3.84.

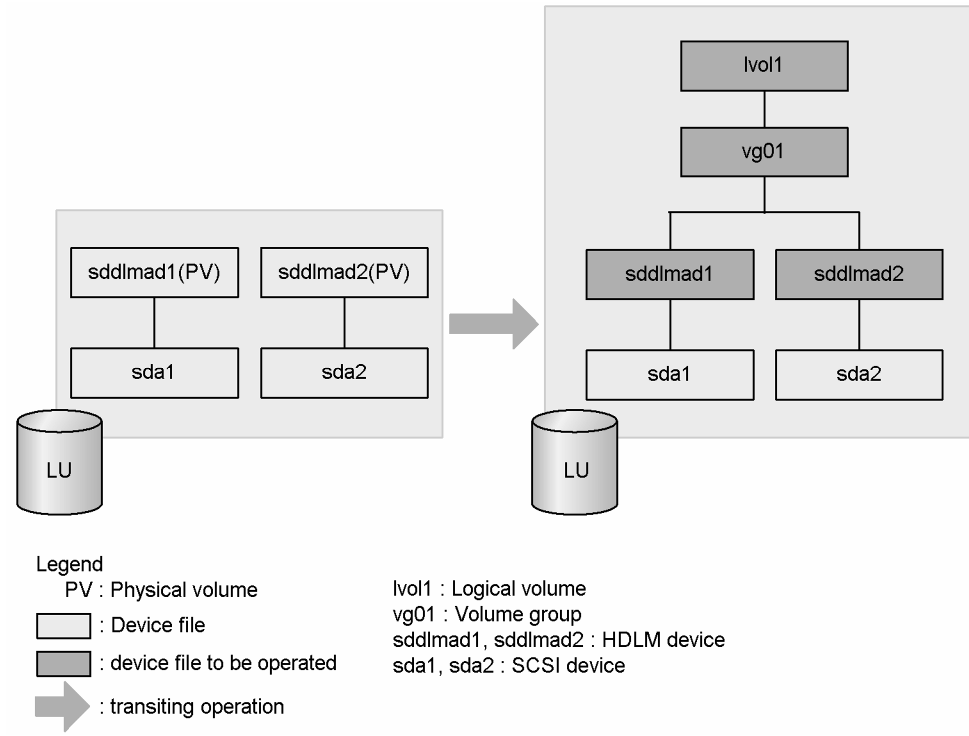


**Figure 3.84 Device configuration When Creating a Logical Volume on an HDLM Device**

To create a new logical volume on an HDLM device:

1. Stop operation of the HDLM device.
2. Create a disk partition on the HDLM device and change the system ID from `Linux` (`83` in hexadecimal representation) to `Linux LVM` (`8e` in hexadecimal representation).

   The following shows an example of how to create two partitions in `/dev/sddlmad` (HDLM device):

   Example of executing the `fdisk` command (IA32):

```
# fdisk /dev/sddlmad

Command (m for help): p

Disk /dev/sddlmad: 255 heads, 63 sectors, 130 cylinders
Units = cylinders of 16065 * 512 bytes

Device Boot    Start    End    Blocks   Id  System
/dev/sddlmad1      1        65     522081   83  Linux
/dev/sddlmad2      66       130    522112+  83  Linux

Command (m for help): t
Partition number (1-4): 1
Hex code (type L to list codes): 8e

Command (m for help): t
Partition number (1-4): 2
Hex code (type L to list codes): 8e

Command (m for help): p

Disk /dev/sddlmad: 255 heads, 63 sectors, 130 cylinders
Units = cylinders of 16065 * 512 bytes

Device Boot    Start    End    Blocks   Id  System
/dev/sddlmad1      1        65     522081   8e  Linux LVM
/dev/sddlmad2      66       130    522112+  8e  Linux LVM

Command (m for help): w
The partition table has been altered!

Calling ioctl() to re-read partition table.
Syncing disks.
```

Example of executing the `parted` command (IA32, IPF or EM64T/AMD64):

```
# parted /dev/sddlmad
GNU Parted 1.6.15
Copyright (C) 1998 - 2004 Free Software Foundation, Inc.
This program is free software, covered by the GNU General Public License.

This program is distributed in the hope that it will be useful, but WITHOUT ANY
WARRANTY; without even the implied warranty of MERCHANTABILITY
or FITNESS FOR A PARTICULAR PURPOSE.  See the GNU General Public License for more
details.

Using /dev/sddlmad
(parted) p
Disk geometry for /dev/sddlmad: 0.000-1025.000 megabytes
Disk label type: msdos
Minor    Start       End     Type      Filesystem  Flags
(parted) mkpart primary 1 500
(parted) mkpart primary 501 1000
(parted) p
Disk geometry for /dev/sddlmad: 0.000-1025.000 megabytes
Disk label type: msdos
Minor    Start       End     Type      Filesystem  Flags
1          1.000    500.000  primary               type=83
2        501.000   1000.000  primary               type=83
(parted) mkfs 1 ext2
(parted) mkfs 2 ext2
(parted) p
Disk geometry for /dev/sddlmad: 0.000-1025.000 megabytes
Disk label type: msdos
Minor    Start       End     Type      Filesystem  Flags
1          1.000    500.000  primary   ext2        type=83
2        501.000   1000.000  primary   ext2        type=83
(parted) set 1 lvm on
(parted) set 2 lvm on
(parted) p
Disk geometry for /dev/sddlmad: 0.000-1025.000 megabytes
Disk label type: msdos
Minor    Start       End     Type      Filesystem  Flags
1          1.000    500.000  primary   ext2        lvm, type=8e
2        501.000   1000.000  primary   ext2        lvm, type=8e
(parted)q

# blockdev --rereadpt /dev/sddlmad
```

When you use the `parted` command to change the partition, to let the system recognize the partitions on the HDLM device, execute the `blockdev` command after the `parted` command.

3. Edit the `/etc/lvm/lvm.conf` file.

   Edit the device section in the file as shown below to disable the SCSI device configuration and enable the HDLM device configuration:

   – Adding and deleting `filter` entry information

     Add the following line and comment out all other `filter` entry information:

     `filter = [ "a|sddlm*|", "r|/dev/sd|" ]`

   – Adding `types` entry information

     Add the following line:

     `types = [ "sddlmfdrv", 16 ]`

     Make sure that `types=[ "fd", 16 ]` was commented out.

   Figure 3.85 shows an example of editing the file when the OS is SUSE LINUX Enterprise Server 9. The shaded section shows the portion to be edited.

```
# This section allows you to configure which block devices should
# be used by the LVM system.
devices {

    # Where do you want your volume groups to appear ?
    dir = "/dev"

    # An array of directories that contain the device nodes you wish
    # to use with LVM2.
    scan = [ "/dev" ]

    # A filter that tells LVM2 to only use a restricted set of devices.
    # The filter consists of an array of regular expressions.  These
    # expressions can be delimited by a character of your choice, and
    # prefixed with either an 'a' (for accept) or 'r' (for reject).
    # The first expression found to match a device name determines if
    # the device will be accepted or rejected (ignored).  Devices that
    # don't match any patterns are accepted.

    # Remember to run vgscan after you change this parameter to ensure
    # that the cache file gets regenerated (see below).

    # By default we accept every block device except udev names:
    # filter = [ "r|/dev/.*/by-path/.*|", "r|/dev/.*/by-id/.*|", "a/.*/" ]
    filter = [ "a|sddlm*|", "r|/dev/sd|" ]

    cache = "/etc/lvm/.cache"

    # You can turn off writing this cache file by setting this to 0.
    write_cache_state = 1

    # Advanced settings.

    # List of pairs of additional acceptable block device types found
    # in /proc/devices with maximum (non-zero) number of partitions.
    # types = [ "fd", 16 ]
    types = [ "sddlmfdrv", 16 ]

    # the block devices it believes are valid.
    # 1 enables; 0 disables.
    sysfs_scan = 1

    # By default, LVM2 will ignore devices used as components of
    # software RAID (md) devices by looking for md superblocks.
    # 1 enables; 0 disables.
    md_component_detection=0
}
```

**Figure 3.85 Example of Editing the /etc/lvm/lvm.conf File**

4. Create the physical volumes.

   The following example shows how to define `/dev/sddlmad1` and `/dev/sddlmad2` as physical volumes:

```
# pvcreate /dev/sddlmad1
  Physical volume "/dev/sddlmad1" successfully created
# pvcreate /dev/sddlmad2
  Physical volume "/dev/sddlmad2" successfully created
```

5. Create a volume group.

   In the following example, the command creates the `vg01` volume group by using the physical volumes `/dev/sddlmad1` and `/dev/sddlmad2`:

```
# vgcreate vg01 /dev/sddlmad1 /dev/sddlmad2
  Volume group "vg01" successfully created
```

6. Create the logical volume.

In the following example, the command creates the logical volume (`lvol1`: 100 MB) by using the `vg01` volume group:

```
# lvcreate -L 100M -n lvol1 vg01
  Logical volume "lvol1" created
```

7. Create a file system.

   The following example shows how to use the `mke2fs` command to create a file system on the `lvol1` logical volume:

```
# mke2fs /dev/vg01/lvol1
Filesystem label=
OS type: Linux
Block size=1024 (log=0)
Fragment size=1024 (log=0)
25688 inodes, 102400 blocks
5120 blocks (5.00%) reserved for the super user
First data block=1
13 block groups
8192 blocks per group, 8192 fragments per group
1976 inodes per group
Superblock backups stored on blocks:
        8193, 24577, 40961, 57345, 73729

Writing inode tables: done
Writing superblocks and filesystem accounting information: done

This filesystem will be automatically checked every 38 mounts or 180 days, whichever comes
first.  Use tune2fs -c or -i to override.
```

8. Create the directory to which the logical volume is to be mounted.

   The following example shows how to create the `/mnt/lvol1` directory:

```
# mkdir /mnt/lvol1
```

9. Mount the logical volume.

   The following example shows how to mount the logical volume to the `/mnt/lvol1` directory:

```
# mount /dev/vg01/lvol1 /mnt/lvol1
```

#### 3.7.2.3 When Moving a Logical Volume Created on a SCSI Device in a Single-Path Environment to an HDLM Device

This subsection explains the procedure for moving a logical volume created on a SCSI device in a single-path environment to an HDLM device in a multi-path environment. Use the following procedure to create the environment in Figure 3.86.

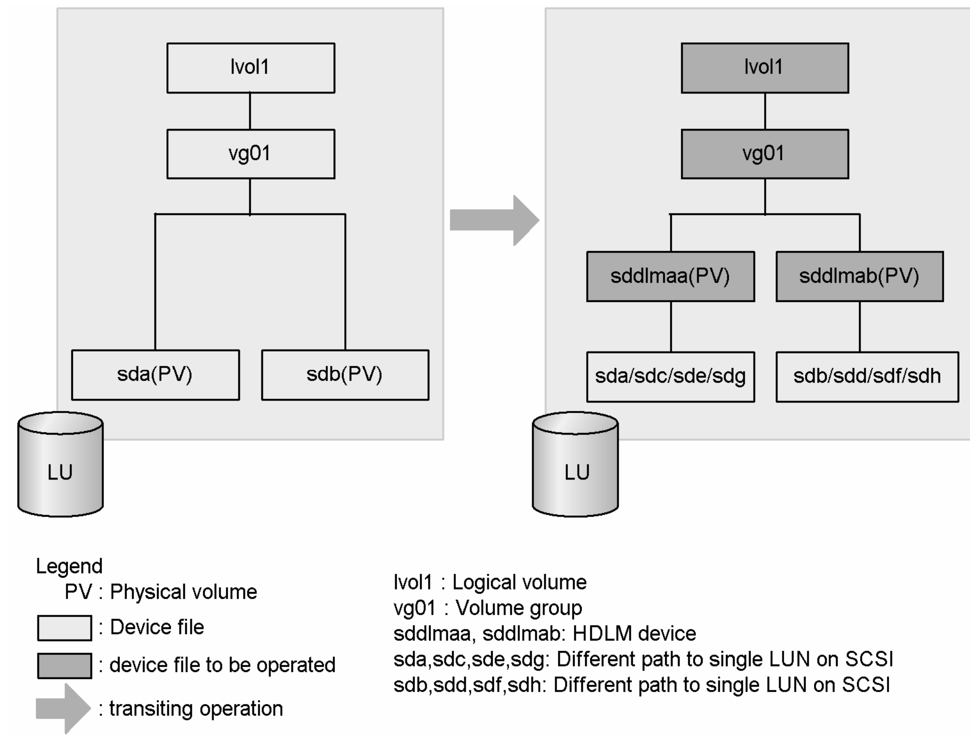


**Figure 3.86 Device Configuration When a Logical Volume on a SCSI Device Is Moved to an HDLM Device**

In Figure 3.86, `sd`*n* indicates the different LUs. `sddlmad1` and `sddlmad2` indicate the HDLM devices corresponding to each LU.

To move the logical volume to an HDLM device:

1. Make sure that HDLM has been installed.

   Execute the `rpm` command to make sure that HDLM has been installed.

   ```
   # rpm -q HDLM
     HDLM-x.xx.x.xxx-x
   ```

2. Make sure that the path is a single path.

   Execute the HDLM-configuration definition utility (`dlmcfgmgr`) with the `-v` parameter specified to make sure that the path to the LU is a single path.

   ```
   # dlmcfgmgr -v
   HDevNam         Management  Device    Host Channel  Target  Lun
   /dev/sddlmaa   configured  /dev/sda   2       0       0      0
   /dev/sddlmab   configured  /dev/sdb   2       0       0      1
   KAPL10302-I /sbin/dlmcfgmgr completed normally.
   ```

3. Stop operation of the SCSI and HDLM devices.

4. Unmount the logical volume to be moved.

   In the following example, the logical volume subject to the move has been mounted on `/mnt/lvol1`:

```
# umount /mnt/lvol1
```

5. Deactivate the volume group.

   In the following example, the logical volume to be moved belongs to the `vg01` volume group:

```
# vgchange -an vg01
  0 logical volume(s) in volume group "vg01" now active
```

6. Export the volume group.

   In the following example, the command exports information about the `vg01` volume group:

```
# vgexport vg01
  Volume group "vg01" successfully exported
```

7. Edit the `/etc/lvm/lvm.conf` file.

   Edit the device section in the file as shown below to disable the SCSI device configuration and enable the HDLM device configuration:

   – Adding and deleting `filter` entry information

     Add the following line and comment out all other `filter` entry information:

     `filter = [ "a|sddlm*|", "r|/dev/sd|" ]`

   – Adding `types` entry information

     Add the following line:

     `types = [ "sddlmfdrv", 16 ]`

     Make sure that `types=[ "fd", 16 ]` was commented out.

   Figure 3.87 shows an example of editing the file when the OS is SUSE LINUX Enterprise Server 9. The shaded section shows the portion to be edited.

```
# This section allows you to configure which block devices should
# be used by the LVM system.
devices {

    # Where do you want your volume groups to appear ?
    dir = "/dev"

    # An array of directories that contain the device nodes you wish
    # to use with LVM2.
    scan = [ "/dev" ]

    # A filter that tells LVM2 to only use a restricted set of devices.
    # The filter consists of an array of regular expressions.  These
    # expressions can be delimited by a character of your choice, and
    # prefixed with either an 'a' (for accept) or 'r' (for reject).
    # The first expression found to match a device name determines if
    # the device will be accepted or rejected (ignored).  Devices that
    # don't match any patterns are accepted.

    # Remember to run vgscan after you change this parameter to ensure
    # that the cache file gets regenerated (see below).

    # By default we accept every block device except udev names:
    # filter = [ "r|/dev/.*/by-path/.*|", "r|/dev/.*/by-id/.*|", "a/.*/" ]
    filter = [ "a|sddlm*|", "r|/dev/sd|" ]

    cache = "/etc/lvm/.cache"

    # You can turn off writing this cache file by setting this to 0.
    write_cache_state = 1

    # Advanced settings.

    # List of pairs of additional acceptable block device types found
    # in /proc/devices with maximum (non-zero) number of partitions.
    # types = [ "fd", 16 ]
    types = [ "sddlmfdrv", 16 ]

    # the block devices it believes are valid.
    # 1 enables; 0 disables.
    sysfs_scan = 1

    # By default, LVM2 will ignore devices used as components of
    # software RAID (md) devices by looking for md superblocks.
    # 1 enables; 0 disables.
    md_component_detection=0
}
```

**Figure 3.87 Example of Editing the /etc/lvm/lvm.conf File**

8. Execute the VG scan.

   To re-create the volume group on the HDLM device, execute the `vgscan` command as shown in the following example:

```
# vgscan
  Reading all physical volumes.  This may take a while...
  Found exported volume group "vg01" using metadata type lvm2
```

9. Import the volume group.

   In the following example, the command imports the `vg01` volume group:

```
# vgimport vg01
  Volume group "vg01" successfully imported
```

10. Execute the `pvscan` command to check that the environment has been changed successfully.

Execute the following command to make sure that the information about the physical volume for the target logical volume has been changed to the information about the HDLM device:

```
# pvscan
  PV /dev/sddlmad1   VG vg01   lvm2 [468.00 MB / 368.00 MB free]
  PV /dev/sddlmad2   VG vg01   lvm2 [548.00 MB / 548.00 MB free]
  Total: 2 [1016.00 MB] / in use: 2 [1016.00 MB] / in no VG: 0 [0   ]
```

11. Activate the volume group.

    In the following example, the logical volume subject to the move belongs to the `vg01` volume group:

```
# vgchange -ay vg01
  1 logical volume(s) in volume group "vg01" now active
```

12. Add a path to the LU.

    Add a path to an existing LU by carrying out the procedure described in 4.4.3.3 in 4.4.3.

13. Mount the logical volume that was moved.

    In the following example, the logical volume that was moved is `/dev/vg01/lvol1` and the command mounts the volume on `/mnt/lvol1`:

```
# mount /dev/vg01/lvol1 /mnt/lvol1
```

## 3.8 Settings for LifeKeeper

This section describes the procedure required to use HDLM together with LifeKeeper. This section also describes how to cancel settings for HDLM and LifeKeeper in the operating environment when you stop using HDLM together with LifeKeeper.

For details about how to operate LifeKeeper, see the LifeKeeper documentation.

### 3.8.1 Notes on Using LifeKeeper

- Although the KAPL05023-E message is output to syslog on a regular basis, operations are not affected.
- The KAPL05023-E message might appear when a resource is added, LifeKeeper is stopped, or a resource is changed. However, operations are not affected.
- If a path is in the `Offline(E)` or `Offline(C)` status when LifeKeeper starts or when a failover makes the standby host active, this path cannot be placed in the `Online` status by an HDLM command or automatic failback.

  In this case, either of the following messages is output to the console window or system log (`/var/log/messages`): If the path could not be placed in the `Online` status by using the appropriate command, the `KAPL01036-E` message is output. If the path could not be placed in the `Online` status by using automatic failback, the `KAPL05702-W` message is output.

  For details about how to place a path in the `Online` status, see section 4.4.4.11 in 4.4.4 or 4.4.4.12 in 4.4.4.
- If a disk partition is created or changed while some of the paths for an LU are disconnected, the configuration of the disk partition might not be recognized even after the disconnected paths are recovered. You must first recover the disconnected paths, and then create or change a disk partition.

  If you created or changed the disk partition while the paths were disconnected, see section 4.4.3.1 in 4.4.3.
- If some of the paths connected to the same LU are repeatedly disconnected and then reconnected at short intervals or if an intermittent error occurs, workload on the entire server temporarily increases due to a problem with the OS. When LifeKeeper is in use, the increased workload may cause a timeout during heartbeat processing with another node, resulting in a failover on the node. Before you disconnect a path, use an HDLM command to place the path in the `Offline` status.
- Use the same version of HDLM on all nodes.

### 3.8.2 When Changing an Environment That Uses a SCSI Device to an Environment That Uses HDLM

To install HDLM and link it to LifeKeeper in an environment where a SCSI device is used and LifeKeeper is running:

1. On all nodes, stop accesses to the SCSI device that is registered as a resource in LifeKeeper.
2. Delete all SCSI device resource information from the active node.

   For example, delete the SCSI device from the FileSystem resources for LifeKeeper.
3. Stop the cluster services of LifeKeeper on all nodes.
4. Install HDLM on all nodes.

   For details about how to install HDLM, see section ¤.
5. On all nodes, change the SCSI device used by LifeKeeper to an HDLM device.

   When the SCSI device is used as a FileSystem resource:

   Change the contents of `/etc/fstab` file to the HDLM device.

   When the SCSI device is used as a raw device resource:

   Change the contents of `/etc/sysconfig/rawdevices` file to the HDLM device.
6. Execute the HDLM `dlmlfk` utility for supporting LifeKeeper on all nodes as shown below:

```
# /opt/DynamicLinkManager/bin/dlmlfk -c
```

7. Start the cluster service of LifeKeeper on all nodes.
8. From the active node, add the HDLM device to the LifeKeeper resources.

   For example, add the HDLM device to the FileSystem resources for LifeKeeper.

### 3.8.3 When Adding LifeKeeper to an Environment That Uses HDLM

To install LifeKeeper and link it to HDLM in an environment that already uses HDLM:

1. Install LifeKeeper on all the applicable nodes.
2. Check the status of the services of LifeKeeper on all nodes, and then stop all the running services.
3. Execute the HDLM `dlmlfk` utility for supporting LifeKeeper on all nodes as shown below:

```
# /opt/DynamicLinkManager/bin/dlmlfk -c
```

4. On all nodes, stop I/O for the HDLM device to be used as a LifeKeeper resource.
5. Execute the following command to make sure that all paths to the HDLM device to be used as a LifeKeeper resource are placed online.

   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           9830           0

PathID  PathName                          DskName                          iLU     ChaPort
  Status    Type IO-Count  IO-Errors  DNum HDevName
000000  0002.0000.0000000000000000.000B   HITACHI .DF600F         .0115   0011    0A
  Online    Non           35          0   0 sddlmaa
000001  0002.0000.0000000000000000.000C   HITACHI .DF600F         .0115   0012    0A
  Online    Own          344          0   0 sddlmab
000002  0002.0000.0000000000000000.000D   HITACHI .DF600F         .0115   001     0A
  Online    Non          132          0   0 sddlmac
000003  0002.0000.0000000000000000.000E   HITACHI .DF600F         .0115   0014    0A
  Online    Own         3923          0   0 sddlmad
000004  0002.0000.0000000000000000.0006   HITACHI .DF600F         .0115   0006    0A
  Online    Own         2953          0   0 sddlmba
000005  0002.0000.0000000000000000.0007   HITACHI .DF600F         .0115   0007    0A
  Online    Non          307          0   0 sddlmbb
000006  0002.0000.0000000000000000.0008   HITACHI .DF600F         .0115   0008    0A
  Online    Own         2111          0   0 sddlmbc
000007  0002.0000.0000000000000000.0009   HITACHI .DF600F         .0115   0009    0A
  Online    Non           25          0   0 sddlmbd
KAPL01001-I The HDLM command completed normally. Operation name =view, completion time =
yyyy/mm/dd hh:mm:ss
#
```

6. Start the cluster service of LifeKeeper on all nodes.
7. From the active node, add the HDLM device to the LifeKeeper resources.

   For example, add the HDLM device to the FileSystem resources for LifeKeeper.

### 3.8.4 When Adding a Node to an Environment That Uses LifeKeeper and HDLM

To add a node to an environment that already uses HDLM and LifeKeeper:

1. On all nodes, stop accesses to the HDLM device included in the LifeKeeper management target.
2. Start the node you want to add.
3. If the HDLM device is defined as a LifeKeeper resource, delete the settings of all HDLM devices from the active node.

   For example, delete the HDLM devices from the FileSystem resources for LifeKeeper.
4. Stop the cluster service of LifeKeeper on the node you want to add.
5. Install HDLM on this node.

   For details about how to install HDLM, see section ¤.
6. Execute the `dlmlfk` utility for supporting LifeKeeper as shown below to change the SCSI device used by LifeKeeper on the added node to the HDLM device.

   ```
   # /opt/DynamicLinkManager/bin/dlmlfk -c
   ```
7. Start the cluster service of LifeKeeper on the added node.
8. From the active node, add the HDLM device to the LifeKeeper resources.

   For example, add the HDLM device to the FileSystem resources for LifeKeeper.

## 3.9 Settings for the Red Hat Cluster Manager

The user must execute the following procedure when HDLM is installed in an environment that uses RHCM.

### 3.9.1 Notes on Using Red Hat Cluster Manager

The HDLM device used by the cluster service and the HDLM device used by Quorum must be the same.

### 3.9.2 Specifying Settings for Red Hat Cluster Manager

To specify settings for Red Hat Cluster Manager:

1. Stop the service of the RHCM cluster.

   Stop the RHCM service with the following command if the service is running:

   ```
   # service clumanager stop
   ```

2. Change the SCSI device being used by RHCM to an HDLM device.

   Edit as shown below if the lower SCSI device (`/dev/sda`) of the HDLM device (`/dev/sddlmaa`) is being used as `/dev/raw/raw1`.

   ```
   # vi /etc/sysconfig/rawdevices
   # /dev/raw/raw1 /dev/sda
   /dev/raw/raw1 /dev/sddlmaa
   ```

   Comment out `/dev/sda` and add `/dev/sddlmaa`.

3. Restart the `rawdevices` service to apply the changes made in Step 2.

   Execute the following command to restart the service:

   ```
   # service rawdevices restart
   ```

4. Start the cluster service of RHCM.

   Executing the following command to start the service:

   ```
   # service clumanager start
   ```

## 3.10 Settings for VERITAS Cluster Server

The following is the procedure that must be executed by the user when HDLM is installed in an environment that uses VCS. For details on how to operate the VCS GUI, see the VCS documentation.

1. Stop the cluster service of VCS.
2. Change the SCSI device used by VCS to an HDLM device.

   Change the `Block Devices` parameter of the `Mount` resource from the SCSI device to an HDLM device.
3. Start the cluster service of VCS.

## 3.11 Checking the Path Configuration

HDLM functionality such as load balancing and failover is 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 `dlnkmgr` 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 section 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 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.

– Make sure that there are different host port numbers and bus numbers for each physical HBA port.

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 to see whether the HDLM device corresponding to the SCSI device has been created.

## 3.12 Setting Up HDLM

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.

### 3.12.1 Checking the Current Settings

For setting the HDLM functionality, this chapter describes how to check the settings before the change by using the `dlnkmgr` command's `set` operation.

Check the current settings by executing the following command to set the HDLM functionality by using the `dlnkmgr` command's `set` operation.

```
# /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
Reservation Status            :
Intermittent Error Monitor    : off
KAPL01001-I The HDLM command completed normally. Operation name = view, completion time =
yyyy/mm/dd hh:mm:ss
#
```

### 3.12.2 Setting up HDLM Functionality

The recommended and default values for the HDLM functions are shown below in Table 3.14.

**Table 3.14 The Recommended and Default Values of Each Function**

| Function | Default value | Recommended value |
|---|---|---|
| Load-balancing | on<br>Algorithm is round robin | on<br>The recommended algorithm depends on the operating environment. |
| Path health checking | on<br>30-minute check interval | on<br>The recommended checking interval depends on the operating environment. |
| Automatic failback | off | off |
| Intermittent Error Monitor | off | off |
| Reservation level | (Not specifiable) | (Not specifiable) |

| Function | Default value | Recommended value |
|---|---|---|
| Remove LU | (Not specifiable) | (Not specifiable) |
| Logging level | 3: Collect all information for errors at the Information level or higher | 3: Collect all information for errors at the Information level 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 |

#### 3.12.2.1 Setting for Load Balancing

You can select whether to enable load balancing.

The following is an example of using a command to set load balancing.

```
# /opt/DynamicLinkManager/bin/dlnkmgr set -lb on -lbtype rr
```

To enable load balancing, specify on. Specify `off` otherwise. When you specify `on`, after the `-lbtype` option specify `rr` for round robin or `exrr` for extended round robin. 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 load balancing without specifying an algorithm, load balancing will be executed according to the setting stored in the system.

#### 3.12.2.2 Setting for Path Health Checking

You can select whether to enable path health checking.

The following is an example of using a command to enable path health checking.

```
# /opt/DynamicLinkManager/bin/dlnkmgr set -pchk on -intvl 10
```

To enable path health checking, specify `on`. Specify `off` otherwise. When you specify `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.

#### 3.12.2.3 Setting for Automatic Failback

You can select whether to enable automatic failback.

When intermittent error monitoring is enabled and the number of error occurrences 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 message (KAPL01080-W) 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 is an example of using a command to enable automatic failback:

```
# /opt/DynamicLinkManager/bin/dlnkmgr set -afb on -intvl 10
```

To enable automatic failback, specify `on`. Specify `off` otherwise. When you specify `on`, you can use the `-intvl` parameter to specify the *checking* interval can be specified with the `-intvl` parameter when `on` is specified. The value that was specified previously will be applied if the *checking* interval is not specified. For example, specify auto failback as `off` after specifying the *checking* interval to five minutes and executing. Then, when executing after specifying auto failback as `on` without specifying the *checking* interval, the five minutes that were specified previously will be applied again.

#### 3.12.2.4 Setting for Intermittent Error Monitoring

Intermittent error monitoring is specifiable 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 number of times, make sure that the condition shown in 3.12.2.3 is satisfied.

To determine whether a path is ineligible for automatic failback, you can use the results of the `dlnkmgr` command's `view` operation.

The following is an example of using a command to enable intermittent error monitoring:

```
# /opt/DynamicLinkManager/bin/dlnkmgr set -iem on -intvl 20 -iemnum 2
```

To enable intermittent error monitoring, specify `on`. To disable intermittent error monitoring, specify `off`. When you specify `on`, 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). In the `-intvl` parameter, specify the monitoring interval for an intermittent error. In the `-iemnum` parameter, specify the number of times that the error is to occur. When these parameters are omitted and 3 or more errors occur within 30 minutes, the system assumes that an intermittent error is occurring.

### 3.12.2.5 Setting for 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 3.15 lists and describes values for the error log collection level setting.

**Table 3.15 Values for the Error Log Collection Level Setting**

| Value | Description |
|---|---|
| 0 | No error logs are collected. |
| 1 | All information for errors at the Error level or higher is collected. |
| 2 | All information for errors at the Warning level or higher is collected. |
| 3 | All information for errors at the Information level or higher is collected. |
| 4 | All information for errors at the Information level or higher (including maintenance information) is collected. |

If an error occurs, you may have to set 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 is an example of using a command to set the error log collection level.

```
# /opt/DynamicLinkManager/bin/dlnkmgr set -ellv 2
```

Specify the error log collection level by a number.

### 3.12.2.6 Setting for the Trace Level

The trace output level can be set.

A trace file in which a trace level can be set is `hdlmtrn.log` (*n* indicates a file number from `1` to `64`).

Table 3.16 lists and describes values for the trace level setting.

**Table 3.16 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 may have to set 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 of using a command to set the trace level.

```
# /opt/DynamicLinkManager/bin/dlnkmgr set -systflv 1
```

Specify the trace level by a number.

### 3.12.2.7 Setting for the Size of Error Log Files

The error log file size (the HDLM manager log (`dlmmgrn.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.

When an error log file reaches the specified size, the information in the old error log file is replaced with new information, beginning with the oldest file. By specifying both the log file size and the number of log files, you can collect up to 32000000 kilobytes (approximately 30 GB) of error logs in total.

The following is an example of using a 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.

#### 3.12.2.8 Setting for 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 30 GB) of error logs in total.

The following is an example of using a command to set the number of error log files.

```
# /opt/DynamicLinkManager/bin/dlnkmgr set -elfn 5
```

Specify the number of error log files by a number.

#### 3.12.2.9 Setting for the Size of Trace Files

The trace file size can be set.

Trace files for which a trace file size can be set 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 (KAPL01097-W) 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 is an example of using a command to set the trace file size:

```
# /opt/DynamicLinkManager/bin/dlnkmgr set -systfs 2000
```

Specify the size of the trace file in kilobytes.

#### 3.12.2.10 Setting for the Number of Trace Files

You can set the number of the trace files.

Trace files for which the number of files can be set 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 message (KAPL01097-W) 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 is an example of using a command to set the number of trace files:

```
# /opt/DynamicLinkManager/bin/dlnkmgr set -systfn 10
```

Specify the number of trace files by a number.

### 3.12.3 Checking the New Settings

This chapter describes steps involved in 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 is 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                   : 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)
Reservation Status           :
Intermittent Error Monitor   : on(2/20)
KAPL01001-I The HDLM command completed normally. Operation name = view, completion time =
yyyy/mm/dd hh:mm:ss
#
```

## 3.13 Setting up Integrated Traces

When HDLM is used, `dlnkmgr` commands logs are output to the following *integrated trace information files* of Hitachi Network Objectplaza Trace Library (HNTRLib2): `/var/opt/hitachi/HNTRLib2/spool/hntr2`*n*`.log` (*n* indicates a file number). Note that HNTRLib, which is an older version of HNTRLib2, is not usable.

If there is a significant amount of integrated trace information output, the information might end up deleted in a short 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, increasing the integrated trace file size and buffer size. ***Note***: specifying excessively large values places a heavy load on the system. When determining these values, consider these operation tradeoffs.

Table 3.17 shows the default and recommended values for the integrated trace file settings.

**Table 3.17 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 cycle | 10 (seconds) | 5 (seconds) |
| | Buffer size per monitoring interval | 64 (kilobytes) | 256 (kilobytes) |
| Number of messages to be output per monitoring interval | Monitoring cycle | 0 (seconds) | 0 (seconds) |
| | Number of messages to be output | 0 | 0 |

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.

### 3.13.1 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

### 3.13.2 Displaying the Hitachi Network Objectplaza Trace Library Setup Menu

To display the Hitachi Network Objectplaza Trace Library setup menu:

1. Log on as a root user.
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.

### 3.13.3 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 will appear. 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)**.

   The specifiable range is from 8 to 8192 KB, with a default of 256. Set this to a value larger than that set in step 2 of 3.13.5. We recommend setting a value of 4096 when 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.

### 3.13.4 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 will appear. 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)**.

   You can specify a value from 1 to 16. The default is 4. 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.

### 3.13.5 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 will appear. 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 the desired size in **New Size (KB)**.

   Set a new buffer size to fit the monitoring interval set in **5: Interval Timer**. The specifiable range is from 8 to 2048 KB, with a default of 64. Set this to a value smaller than that set in step 2 of 3.13.3. We recommend setting a value of 256 when collecting an integrated trace.

   To leave the buffer size as is, leave **New Size (KB)** blank, type `!` and press the **Enter** key. You will be returned 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 will appear. 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 the desired interval in **New Span (sec)**.

   The specifiable range is from 1 to and 300 seconds, with a default of 10. We recommend setting a value of 5 when collecting an integrated trace.

   To leave the monitoring interval as is, leave **New Span (sec)** blank, enter `!` and press the **Enter** key. 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.

### 3.13.6 Adjusting the Number of Messages to Be Output Per Monitoring Interval

To adjust the number of messages to be output per 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 amount of messages output to the integrated trace file will appear. 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)**.

   The specifiable range is from 0 to 3600 seconds, with a default of 0. We recommend setting a value of 0.

   To leave the monitoring interval as is, leave **New Span (sec)** blank, type `!` and press the **Enter** key. You will be returned to the Hitachi Network Objectplaza Trace Library setup menu.

   ***Note:*** 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 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** will appear. 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 specifiable range is from 0 to 500 messages, with a default of 0. If you want to increase the number of messages which are output to the integrated trace file as much as possible, we recommend setting a value of 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.

   To leave the maximum number of messages output as is, leave **New Max (sec)** blank, enter `!` and press the **Enter** key. 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.

### 3.13.7 Finishing the Hitachi Network Objectplaza Trace Library Settings

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 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, enter `Yes`, otherwise, enter `No`.

### 3.13.8 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
   HDLM
   #
   ```

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.

## 3.14 Creating a Character-Type Device File for an HDLM Device

This section describes how to create a character-type device file for an HDLM device.

HDLM creates only a block-type device under `/dev`. A character-type device file is created from the HDLM device using the Linux `raw` command. The following is an example of executing the `raw` command for an HDLM device.

```
# raw /dev/raw/raw1 /dev/sddlmaa
/dev/raw/raw1:  bound to major 253, minor 0
```

**Figure 3.88 Executing the raw Command**

In this example, the command creates `/dev/raw/raw1` as a character-type device file for `/dev/sddlmaa` (major number of `253`, minor number of `0`).

## 3.15 Creating File Systems for HDLM (When Volume Management Software Is Not Used)

This section describes how to build a file system in an HDLM-managed device without using volume management software.

### 3.15.1 Mounting a File System

To mount an HDLM-managed device as a file system by specifying the logical device file name for an HDLM device:

1. Create a file system by using an HDLM device.

   If the file system already exists, skip to step 2.

   ***Note:*** the file systems supported by HDLM are ext2, ext3, ReiserFS, and VERITAS File System (VxFS). For details about creating a file system, see section 3.15.2.

2. Create a directory on which to mount the file system.

```
# mkdir /mnt/hdlm
```

   In this example, the command creates `/mnt/hdlm` to be used for the mount point.

3. Mount the file system.

   Execute a command such as the following:

```
# mount /dev/sddlmaa /mnt/hdlm
```

   In this example, the command mounts `/dev/sddlmaa` (the logical device file name for an HDLM device) `/mnt/hdlm`.

### 3.15.2 Creating a File System

The file systems supported by HDLM are ext2, ext3, ReiserFS, and VxFS (VERITAS File System). The following example shows how to create an ext2 file system in `/dev/sddlmaa1`.

```
# mkfs -t ext2 /dev/sddlmaa1
mke2fs 1.26 (3-Feb-2002)
warning: 184 blocks unused.

Filesystem label=
OS type: Linux
Block size=1024 (log=0)
Fragment size=1024 (log=0)
100744 inodes, 401409 blocks
20079 blocks (5.00%) reserved for the super user
First data block=1
49 block groups
8192 blocks per group, 8192 fragments per group
2056 inodes per group
Superblock backups stored on blocks:
        8193, 24577, 40961, 57345, 73729, 204801, 221185

Writing inode tables: done
Writing superblocks and filesystem accounting information: done

This filesystem will be automatically checked every 39 mounts or
180 days, whichever comes first.  Use tune2fs -c or -i to override.
```

**Figure 3.89 Example of Creating an ext2 File System**

## 3.16 Settings for Automatic Mounting

If you specify the name of a logical device file of an HDLM device at host startup, an HDLM-managed device is automatically mounted in the same way as if you had specified the name of a logical device file of a SCSI device. To mount an HDLM-managed device automatically, you need to edit the `/etc/fstab` file. In this file, you can specify the name of the logical device file of the SCSI device and the mount point of its disk. After this file has been created, the specified HDLM-managed device is mounted automatically when Linux is started. The Linux functionality that adds `LABEL=` to a SCSI device is not supported in HDLM. Do not use this functionality.

The following subsections describe the tasks that are required for using HDLM.

### 3.16.1 Setting the HDLM Device for the First Time

To mount an HDLM-managed device automatically by setting a new HDLM device:

1. Identify the name of the logical device file of the HDLM device to be mounted automatically.

   Execute the `view` operation of the HDLM command, and identify the applicable name of the logical device file of the HDLM and SCSI devices. For details on the `view` operation, see section 6.7. The following is an actual example.

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -drv
PathID HDevName            Device          LDEV
000000 sddlmaa            /dev/sda         9970/9980.15001.05B7
000001 sddlmaa            /dev/sdq         9970/9980.15001.05B7
000002 sddlmab            /dev/sdr         9970/9980.15001.05B0
000003 sddlmac            /dev/sds         9970/9980.15001.05B1
000004 sddlmad            /dev/sdt         9970/9980.15001.05B2


000025 sddlmab            /dev/sdj         9970/9980.15001.05B0
000026 sddlmac            /dev/sdk         9970/9980.15001.05B1
000027 sddlmad            /dev/sdl         9970/9980.15001.05B2
000028 sddlmae            /dev/sdm         9970/9980.15001.05B3
000029 sddlmaf            /dev/sdn         9970/9980.15001.05B4
000030 sddlmag            /dev/sdo         9970/9980.15001.05B5
000031 sddlmah            /dev/sdp         9970/9980.15001.05B6
KAPL01001-I The HDLM command completed normally. Operation name =
view, completion time = yyyy/mm/dd hh:mm:ss
```

**Figure 3.90 Example of Executing the dlnkmgr Command's View Operation (With -drv Specified)**

   In the example, the name of the logical device file of the HDLM device that corresponds to `/dev/sda` can be identified as `sddlmaa`.

2. Add the name of the logical device file of the HDLM device that was confirmed in step 1 to the `/etc/fstab` file.

   The following is an example of editing the `/etc/fstab` file:

```
/dev/sddlmaa     /mnt/sda        ext2    defaults       0 0
```

This example adds a line for `/dev/sddlmaa` (name of the logical device file of the HDLM device).

3. Restart the host or manually mount the HDLM device.

   HDLM-managed devices are mounted by using the logical device file name of the HDLM device.

### 3.16.2 Migrating from an Environment Where a SCSI Device Is Already Set

To move an already set SCSI device to an HDLM device:

1. Identify the name of the logical device file of the HDLM device to be mounted automatically.

   Execute the `dlnkmgr` command's `view` operation, and identify the applicable name of the logical device file of the HDLM and SCSI devices. For details on the `view` operation, see section 6.7. The following is an actual example.

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -drv
PathID HDevName            Device         LDEV
000000 sddlmaa            /dev/sda        9970/9980.15001.05B7
000001 sddlmaa            /dev/sdq        9970/9980.15001.05B7
000002 sddlmab            /dev/sdr        9970/9980.15001.05B0
000003 sddlmac            /dev/sds        9970/9980.15001.05B1
000004 sddlmad            /dev/sdt        9970/9980.15001.05B2


000025 sddlmab            /dev/sdj        9970/9980.15001.05B0
000026 sddlmac            /dev/sdk        9970/9980.15001.05B1
000027 sddlmad            /dev/sdl        9970/9980.15001.05B2
000028 sddlmae            /dev/sdm        9970/9980.15001.05B3
000029 sddlmaf            /dev/sdn        9970/9980.15001.05B4
000030 sddlmag            /dev/sdo        9970/9980.15001.05B5
000031 sddlmah            /dev/sdp        9970/9980.15001.05B6
KAPL01001-I The HDLM command completed normally. Operation name =
view, completion time = yyyy/mm/dd hh:mm:ss
```

**Figure 3.91 Example of Executing the dlnkmgr Command's View Operation (With -drv Specified)**

   In the example, the name of the logical device file of the HDLM device that corresponds to `/dev/sda` can be identified as `sddlmaa`.

2. In the `/etc/fstab` file, change the logical device file name for a SCSI device to that for the HDLM device that you confirmed in step 1.

   The following is an example of editing the `/etc/fstab` file:

```
# /dev/sda          /mnt/sda        ext2    defaults       0 0
/dev/sddlmaa     /mnt/sda        ext2    defaults       0 0
```

   This example comments out the line for `/dev/sda` (logical device file name of the SCSI device) and adds a line for `/dev/sddlmaa` (logical device file name of the HDLM device).

3. Restart the host or manually mount the HDLM device.

   The HDLM-managed device is mounted using the logical device file name of the HDLM device.

## 3.17 Canceling the Settings for HDLM

Return the HDLM environment to the way it was before HDLM was installed by following the procedure flow shown in Figure 3.92.

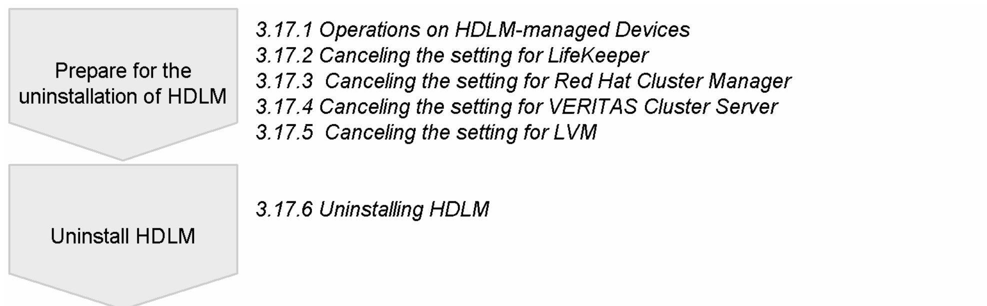


**Figure 3.92 Overview of HDLM Uninstallation**

### 3.17.1 Operations on HDLM-Managed Devices

The following procedure must be performed before uninstalling HDLM:

1. Back up all HDLM-managed devices as required (for example, using a tape backup).
2. Collect all information on the correspondence between logical device file names for the HDLM devices and LU numbers of the LUs in the storage subsystem.

   This step enables you to keep the correspondence between LUs and the logical device file names for SCSI devices once uninstallation is complete.
3. Cancel HDLM device registration in an application.

   An application uses an HDLM device when accessing an LU, so after HDLM is uninstalled, the application cannot use the same HDLM device to access that LU. If an HDLM-managed device is registered in an application (including cluster software, but excluding volume management software), cancel registration of the HDLM device.

   For details how to cancel the LifeKeeper settings, see section 3.17.2. For details how to cancel the RHCM settings, see section 3.17.3. For details on how to cancel the VCS settings, see section 3.17.4.
4. Unmount the HDLM device.

   If HDLM-managed devices have been mounted by specifying HDLM devices, unmount them. Also, if the disks are set to be mounted automatically when the host starts, delete this setting in the `/etc/fstab` file.

   For details on how to cancel the setting, see section 3.5.1.1 in 3.5.1.

   However, perform the procedure by releasing an HDLM disk that is already defined and adding a SCSI device.
5. Cancel HDLM device registration in volume management software.

   If HDLM devices are registered in volume management software, cancel the registration.

For details on how to cancel the LVM settings, see section 3.17.5.

### 3.17.2 Canceling the Settings for LifeKeeper

This section describes how to cancel the HDLM or LifeKeeper settings in an environment where LifeKeeper and HDLM are linked.

#### 3.17.2.1 When Canceling HDLM Settings from an Environment That Is Using HDLM

To cancel HDLM settings from an environment where LifeKeeper and HDLM are linked:

1. On all nodes, stop accesses to the HDLM device that is registered as a resource in LifeKeeper.
2. If the HDLM device is defined as a LifeKeeper resource, delete the settings of all HDLM devices from the active node.

   For example, delete the HDLM devices from the FileSystem resources for LifeKeeper.
3. Stop the cluster service of LifeKeeper on all nodes.
4. On all nodes, change the HDLM device used by LifeKeeper as a resource to a SCSI device.

   When the HDLM device is used as a FileSystem resource:

   Change the contents of `/etc/fstab` file to the SCSI device.

   When the HDLM device is used as a raw device resource:

   Change the contents of `/etc/sysconfig/rawdevices` file to the SCSI device.
5. Execute the HDLM `dlmlfk` utility for supporting LifeKeeper on all nodes as shown below:

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

6. If necessary, uninstall HDLM from all nodes.

   You can use LifeKeeper even if you do not uninstall HDLM.

   For details about how to uninstall HDLM, see section 3.17.6.
7. On all nodes, delete SCSI devices for which the same LU is specified, other than the one specified for `/etc/fstab` file in step 4.

   LifeKeeper does not support a multi-path configuration. To use LifeKeeper in an environment that does not use HDLM, change to a single-path configuration.
8. Start the cluster service of LifeKeeper on all nodes.
9. From the active node, add the SCSI device to the LifeKeeper resources.

   For example, add the SCSI device to the FileSystem resources for LifeKeeper.

#### 3.17.2.2 When Canceling LifeKeeper Settings from an Environment That Is Using LifeKeeper

To cancel LifeKeeper settings from an environment where LifeKeeper and HDLM are linked:

1. On all nodes, stop accesses to the HDLM device that is registered as a resource in LifeKeeper.

2. Delete the settings of all HDLM devices from the active node.

   For example, delete the HDLM devices from the FileSystem resources for LifeKeeper.

3. Stop the cluster service of LifeKeeper on all nodes.

4. Execute the HDLM `dlmlfk` utility for supporting LifeKeeper on all nodes.

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

5. Uninstall LifeKeeper from all nodes.

### 3.17.3 Canceling the Settings for Red Hat Cluster Manager

This section describes how to cancel the HDLM setting in an environment where RHCM is used. If the environment is one where an HDLM device is already being used in RHCM, the following procedure must be executed in order to change to an environment that uses a character-type device of a SCSI device.

1. Stop the cluster service of RHCM.

   If the RHCM service is running, stop the service by executing the following command:

```
# service clumanager stop
```

2. Replace the HDLM device used by RHCM with a SCSI device.

   Edit as shown below to use the SCSI device `/dev/sda`, which is the lower SCSI device of the HDLM device`/dev/sddlmaa`, as `/dev/raw/raw1`.

```
# vi /etc/sysconfig/rawdevices
# /dev/raw/raw1 /dev/sddlmaa
/dev/raw/raw1 /dev/sda
```

   Comment out `/dev/sddlmaa` and add `/dev/sda`.

3. Restart the `rawdevices` service to apply the changes made in Step 2.

   Execute the following command to restart the service:

```
# service rawdevices restart
```

### 3.17.4 Canceling the Settings for VERITAS Cluster Server

For details on how to perform operations from the VCS GUI, see the VCS documentation.

To cancel the HDLM settings in an environment where VCS is used:

1. Stop the cluster service of VCS.

2. Change the HDLM device used by VCS to a SCSI device.

   Change the `Block Devices` parameter in `Mount` resource from an HDLM device to SCSI device.

### 3.17.5 Canceling the Settings for LVM

This section explains how to change from an environment where HDLM devices are used as LVM physical volumes, to an environment where SCSI devices are used as LVM physical volumes. This section explains how to cancel the settings for LVM in an environment using Red Hat Enterprise Linux AS 3 or Red Hat Enterprise Linux ES 3, and for LVM2 in an environment using SUSE LINUX Enterprise Server 9, SUSE LINUX Enterprise Server 10, Red Hat Enterprise Linux AS 4, or Red Hat Enterprise Linux ES 4.

Canceling the settings for LVM

– When moving a logical volume created on an HDLM device to a SCSI device
– When moving a logical volume created on md devices that use an HDLM device, to md devices that use a SCSI device

Canceling the settings for LVM2

– When moving a logical volume created on an HDLM device in a multi-path environment to a SCSI device in a single-path environment.

#### 3.17.5.1 When using Red Hat Enterprise Linux AS 3 or Red Hat Enterprise Linux ES 3

When Migrating a Logical Volume Created on an HDLM Device to a SCSI Device

This subsection explains the procedure for migrating a logical volume created on an HDLM device to a SCSI device. Use the following procedure to create the environment in Figure 3.93.

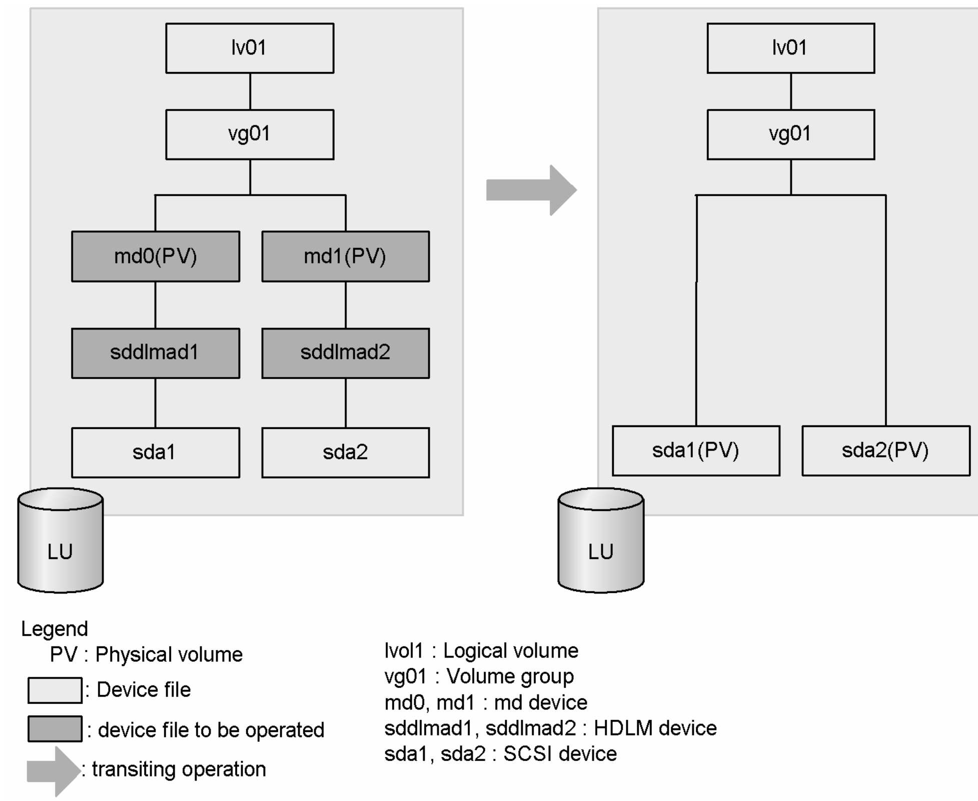


**Figure 3.93 Device Configuration Where a Logical Volume on an HDLM Device Is Moved to a SCSI Device**

To migrate the logical volume to a SCSI device:

1. Stop operation of the SCSI and HDLM devices.
2. Back up the data on the logical volume that you want to migrate to a SCSI device.
3. Unmount the logical volume to be migrated.

   In the following example, the logical volume subject to the migration has been mounted on `/mnt/lvol1`:

   ```
   # umount /mnt/lvol1
   ```

4. Deactivate the volume group.

   In the following example, the logical volume subject to the migration belongs to the vg01 volume group:

   ```
   # vgchange -an vg01
   vgchange -- volume group "vg01" successfully deactivated
   ```

5. Export the volume group.

   In the following example, the command exports information about the vg01 volume group:

   ```
   # vgexport vg01
   vgexport -- volume group "vg01" successfully exported
   ```

6. Stop the md devices.

In the following example, the commands stop `/dev/md0` and `/dev/md1` (md devices created on the HDLM device):

```
# raidstop /dev/md0
# raidstop /dev/md1
```

7. Edit the `/etc/raidtab` file.

   To create md devices on the HDLM device, delete the entries `raiddev /dev/md0` and `raiddev /dev/md1` that were added to the `/etc/raidtab` file. The following is an example of the entries to be deleted.

```
raiddev /dev/md0
raid-level      linear
chunk-size      32
nr-raid-disks   1
persistent-superblock 1
device          /dev/sddlmad1
raid-disk       0

raiddev /dev/md1
raid-level      linear
chunk-size      32
nr-raid-disks   1
persistent-superblock 1
device          /dev/sddlmad2
raid-disk       0
```

8. Change the system ID for the disk partition from `Linux` (`83` in hexadecimal representation) to `Linux LVM` (`8e` in hexadecimal representation).

   Example of executing the fdisk command (IA32)

```
# fdisk /dev/sda

Command (m for help): p

Disk /dev/sda: 255 heads, 63 sectors, 130 cylinders
Units = cylinders of 16065 * 512 bytes

Device Boot    Start    End    Blocks   Id  System
/dev/sda1      1        65     522081   83  Linux
/dev/sda2      66       130    522112+  83  Linux

Command (m for help): t
artition number (1-4): 1
Hex code (type L to list codes): 8e

Command (m for help): t
Partition number (1-4): 2
Hex code (type L to list codes): 8e

Command (m for help): p

Disk /dev/sda: 255 heads, 63 sectors, 130 cylinders
Units = cylinders of 16065 * 512 bytes

Device Boot    Start    End    Blocks   Id  System
/dev/sda1      1        65     522081   8e  Linux LVM
/dev/sda2      66       130    522112+  8e  Linux LVM

Command (m for help): w
The partition table has been altered!

Calling ioctl() to re-read partition table.
Syncing disks.
```

Example of executing the parted command (IA32, IPF or EM64T/AMD64)

```
# parted /dev/sda set 1 lvm on
# parted /dev/sda set 2 lvm on

Check partition type with print option.
# parted /dev/sda print
Disk geometry for /dev/sda: 0.000-10241.000 megabytes
Disk label type: msdos
Minor   Start     End        Type      Filesystem  Flags
1       0.031     5098.754   primary               lvm, type=8e
2       5098.755  10236.730  primary               lvm, type=8e
Information: Don't forget to update /etc/fstab, if necessary.

Update the /proc/partitions by running the blockdev command.
```

9. Execute the `vgscan` command.

   To re-create the volume group, execute the `vgscan` command as shown in the following example:

```
# vgscan
vgscan -- reading all physical volumes (this may take a while...)
vgscan -- found exported volume group "vg01PV_EXP"
vgscan -- "/etc/lvmtab" and "/etc/lvmtab.d" successfully created
vgscan -- WARNING: This program does not do a VGDA backup of your volume group
```

10. Import the volume group.

    In the following example, the command imports the `vg01` volume group:

```
# vgimport vg01 /dev/sda1 /dev/sda2
vgimport -- doing automatic backup of volume group "vg01"
vgimport -- volume group "vg01" successfully imported and activated
```

11. Execute the `pvscan` command to check that the environment has been changed successfully.

    Make sure that the target SCSI device is in the `ACTIVE` status. An example of command execution is as follows:

    To check that the information about the physical volume for the target logical volume has been changed to the information about the SCSI device, execute the following command:

```
# pvscan
pvscan -- reading all physical volumes (this may take a while...)
pvscan -- inactive PV "/dev/sddlmad1" of VG "vg01" [384 MB / 284 MB free]
pvscan -- inactive PV "/dev/sddlmad2" of VG "vg01" [384 MB / 384 MB free]
pvscan -- ACTIVE   PV "/dev/sda1"     of VG "vg01" [384 MB / 284 MB free]
pvscan -- ACTIVE   PV "/dev/sda2"     of VG "vg01" [384 MB / 384 MB free]
pvscan -- total: 10 [3.83 GB] / in use: 10 [3.83 GB] / in no VG: 0 [0]
```

    ***Note:***

    If you uninstall HDLM when multiple paths are set for LVM, only one path is indicated as `ACTIVE`; other paths are indicated as `inactive`. This occurs because LVM does not support a multi-path environment, but there is no problem. We recommend that you use LVM in a single-path environment.

12. Mount the logical volume that was migrated.

    In the following example, the logical volume that was migrated is `/dev/vg01/lvol1` and the command mounts the volume on `/mnt/lvol1`:

```
# mount /dev/vg01/lvol1 /mnt/lvol1
```

13. Move the data that you backed up in step 2 to the mounted logical volume.

When Migrating a Logical Volume Created on md Devices That Use an HDLM Device, to md Devices That Use a SCSI Device

This subsection explains the procedure for migrating a logical volume created on md devices that use an HDLM device, to md devices that use a SCSI device. Use the following procedure to create the environment in Figure 3.94.

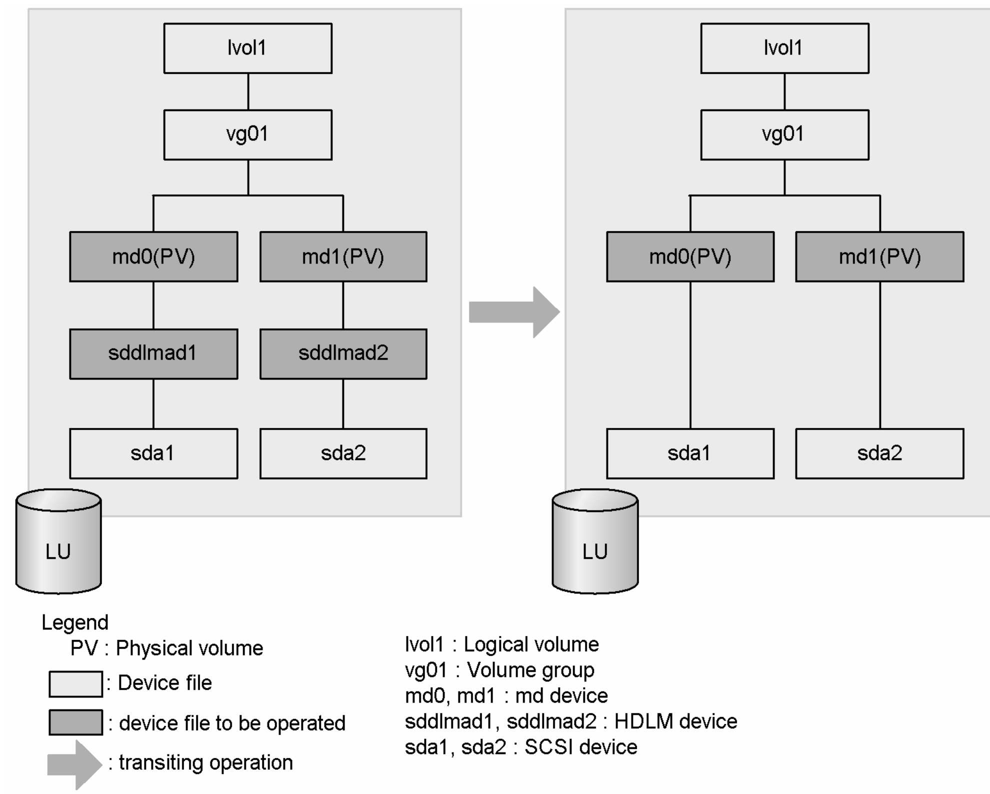


**Figure 3.94 Device Configuration Where the Logical Volume Created on md Devices that Use an HDLM Device to md Devices that Use a SCSI Device**

To migrate the logical volume to md devices that use a SCSI device:

1. Stop operation of the SCSI, md, and HDLM devices.
2. Unmount the logical volume to be migrated.

   In the following example, the logical volume subject to the migration has been mounted on `/mnt/lvol1`:

   ```
   # umount /mnt/lvol1
   ```

3. Deactivate the volume group.

   In the following example, the logical volume subject to the migration belongs to the vg01 volume group:

   ```
   # vgchange -an vg01
   vgchange -- volume group "vg01" successfully deactivated
   ```

4. Export the volume group.

   In the following example, the command exports information about the vg01 volume group:

   ```
   # vgexport vg01
   vgexport -- volume group "vg01" successfully exported
   ```

5. Stop the md devices.

In the following example, the commands stop the md devices `/dev/md0` and `/dev/md1`:

```
# raidstop /dev/md0
# raidstop /dev/md1
```

6. Edit the `/etc/raidtab` file.

   Edit the `/etc/raidtab` file in order to create md devices on the SCSI device. To create the md devices, modify the entries of the already-registered HDLM device. The following shows an example of editing the `/etc/raidtab` file:

```
raiddev /dev/md0
raid-level      linear
chunk-size      32
nr-raid-disks   1
persistent-superblock 1
device          /dev/sddlmad1
raid-disk       0
raiddev /dev/md1
raid-level      linear
chunk-size      32
nr-raid-disks   1
persistent-superblock 1
device          /dev/sddlmad2
raid-disk       0
```

Editing

```
raiddev /dev/md0
raid-level      linear
chunk-size      32
nr-raid-disks   1
persistent-superblock 1
device          /dev/sda1
raid-disk       0
raiddev /dev/md1
raid-level      linear
chunk-size      32
nr-raid-disks   1
persistent-superblock 1
device          /dev/sda2
raid-disk       0
```

**Figure 3.95 Example of Editing the /etc/raidtab File When the md Devices Are Created on the SCSI Devices**

7. Create the md devices on the SCSI device.

   In the following example, the command creates `/dev/md0` and `/dev/md1` as md devices:

```
# mkraid -R /dev/md0 /dev/md1
DESTROYING the contents of /dev/md0 in 5 seconds, Ctrl-C if unsure!
handling MD device /dev/md0
analyzing super-block
disk 0: /dev/sddlmad1, 401593kB, raid superblock at 401472kB
DESTROYING the contents of /dev/md1 in 5 seconds, Ctrl-C if unsure!
handling MD device /dev/md1
analyzing super-block
disk 0: /dev/sddlmad2, 401625kB, raid superblock at 401536kB
```

   ***Note:***

   Always specify the `-R` option.

8. Execute the `vgscan` command.

   To re-create the volume group, execute the `vgscan` command as shown in the following example:

```
# vgscan
vgscan -- reading all physical volumes (this may take a while...)
vgscan -- found exported volume group "vg01PV_EXP"
vgscan -- "/etc/lvmtab" and "/etc/lvmtab.d" successfully created
vgscan -- WARNING: This program does not do a VGDA backup of your volume group
```

9. Import the volume group.

In the following example, the command imports the `vg01` volume group:

```
# vgimport vg01 /dev/md0 /dev/md1
vgimport -- doing automatic backup of volume group "vg01"
vgimport -- volume group "vg01" successfully imported and activated
```

10. Deactivate the volume group.

    In the following example, the logical volume subject to the migration belongs to the `vg01` volume group:

```
# vgchange -an vg01
vgchange -- volume group "vg01" successfully deactivated
```

11. Stop the md devices.

    In the following example, the commands stop the md devices `/dev/md0` and `/dev/md1`:

```
# raidstop /dev/md0
# raidstop /dev/md1
```

12. Start the md devices.

    In the following example, the commands start the md devices `/dev/md0` and `/dev/md1`:

```
# raidstart /dev/md0
# raidstart /dev/md1
```

13. Activate the volume group.

    In the following example, the logical volume subject to the migration belongs to the `vg01` volume group:

```
# vgchange -ay vg01
vgchange -- volume group "vg01" successfully activated
```

14. Execute the `pvscan` command to check that the environment has been changed successfully.

    To check that the information about the physical volume for the target logical volume has been changed to the information about the md devices created on the SCSI device, execute the following command:

```
# pvscan
pvscan -- reading all physical volumes (this may take a while...)
pvscan -- WARNING: physical volume "/dev/sddlmad1" belongs to a meta device
pvscan -- WARNING: physical volume "/dev/sddlmad2" belongs to a meta device
pvscan -- WARNING: physical volume "/dev/sda1" belongs to a meta device
pvscan -- WARNING: physical volume "/dev/sda2" belongs to a meta device
pvscan -- ACTIVE   PV "/dev/md0"      of VG "vg01" [384 MB / 284 MB free]
pvscan -- ACTIVE   PV "/dev/md1"      of VG "vg01" [384 MB / 384 MB free]
pvscan -- total: 12 [784.19 MB] / in use: 12 [784.19 MB] / in no VG: 0 [0]
```

15. Mount the logical volume that was migrated.

    In the following example, the logical volume that was migrated is `/dev/vg01/lvol1` and the command mounts the volume on `/mnt/lvol1`:

```
# mount /dev/vg01/lvol1 /mnt/lvol1
```

### 3.17.5.2 When using SUSE LINUX Enterprise Server 9, SUSE LINUX Enterprise Server 10, Red Hat Enterprise Linux AS 4, or Red Hat Enterprise Linux ES 4

This subsection explains the procedure for moving a logical volume created on an HDLM device in a multi-path environment to a SCSI device in a single-path environment.

To create the environment shown in Figure 3.96, carry out the procedure below. To uninstall HDLM, carry out the steps before the uninstallation.

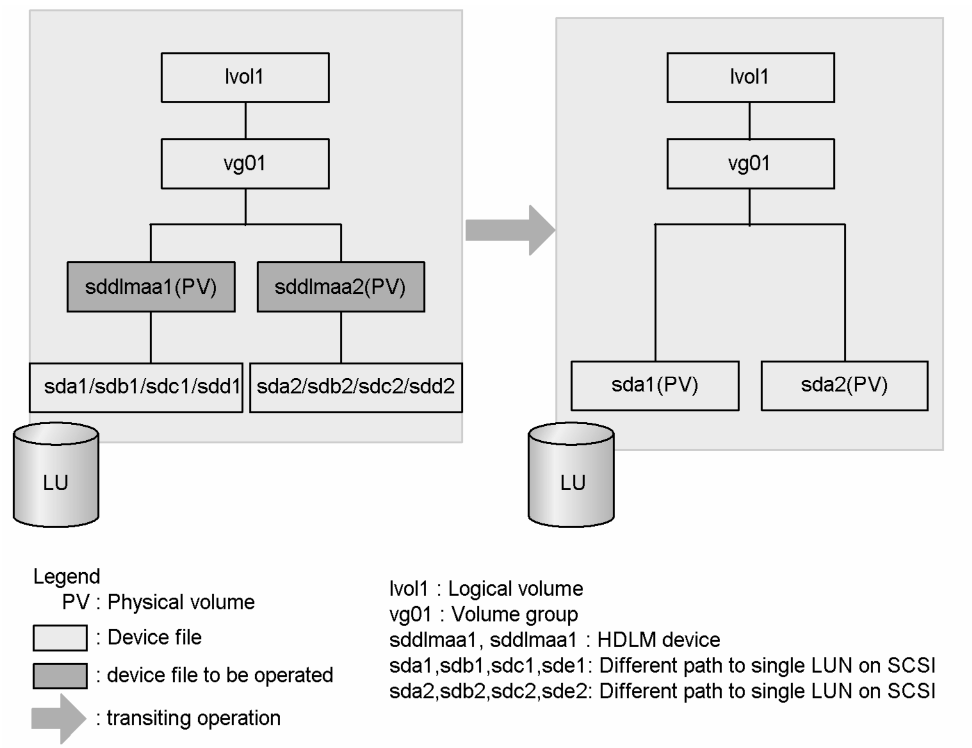


**Figure 3.96 Device Configuration When a Logical Volume on an HDLM Device Is Moved to a SCSI Device**

In Figure 3.96, `sda1`, `sdb1`, `sdc1`, and `sde1` indicate the same LU. Likewise, `sda2`, `sdb2`, `sdc2`, and `sde2` indicate the same LU. `sddlmaa1` and `sddlmaa2` indicate the HDLM devices corresponding to each LU.

To move the logical volume to a SCSI device:

1. Stop operation of the SCSI and HDLM devices.
2. Unmount the logical volume to be moved.

   In the following example, the logical volume subject to the move has been mounted on `/mnt/lvol1`:

```
# umount /mnt/lvol1
```

3. Deactivate the volume group.

   In the following example, the logical volume to be moved belongs to the `vg01` volume group:

```
# vgchange -an vg01
0 logical volume(s) in volume group "vg01" now active
```

4. Export the volume group.

   In the following example, the command exports information about the `vg01` volume group:

```
# vgexport vg01
Volume group "vg01" successfully exported
```

5. If the LVM environment uses an HDLM device in a multi-path configuration, configure the path to an existing LU as a single path.

   Change the path to an existing LU as a single-path configuration by carrying out the procedure described in 4.4.3.5 in 4.4.3.

   ***Note:***

   When you restart the host, if the `/etc/fstab` file contains a description about mounting the target logical volume, comment out the description. Also, stop the applications that access the target logical volume.

6. Edit the `/etc/lvm/lvm.conf` file.

   Edit the device section in the file as shown below to enable the SCSI device configuration and disable the HDLM device configuration:

   – Adding and deleting the `filter` entry information

     Add the following line:

     `filter = [ "r|sddlm*|", "a|/dev/sd|" ]`

     Delete or comment out the following line:

     `filter = [ "a|sddlm*|", "r|/dev/sd|" ]`

   – Adding and deleting `types` entry information

     Delete or comment out the following line:

     `types = [ "sddlmfdrv", 16 ]`

   To uninstall HDLM, perform the following operations:

   Delete `filter = [ "a|sddlm*|", "r|/dev/sd|" ]` and `types = [ "sddlmfdrv", 16 ]`.

   Delete the comment out character (#) from the `filter` entry to return the definition file to the state it was in before the HDLM device configuration was enabled.

   Figure 3.97 shows an example of editing the file by adding and commenting out lines. In this example, SUSE LINUX Enterprise Server 9 is used as the OS. The shaded parts show the locations that were edited.

```
# This section allows you to configure which block devices should
# be used by the LVM system.
devices {

    # Where do you want your volume groups to appear ?
    dir = "/dev"

    # An array of directories that contain the device nodes you wish
    # to use with LVM2.
    scan = [ "/dev" ]

    # A filter that tells LVM2 to only use a restricted set of devices.
    # The filter consists of an array of regular expressions.  These
    # expressions can be delimited by a character of your choice, and
    # prefixed with either an 'a' (for accept) or 'r' (for reject).
    # The first expression found to match a device name determines if
    # the device will be accepted or rejected (ignored).  Devices that
    # don't match any patterns are accepted.

    # Remember to run vgscan after you change this parameter to ensure
    # that the cache file gets regenerated (see below).

    # By default we accept every block device except udev names:
    # filter = [ "r|/dev/.*/by-path/.*|", "r|/dev/.*/by-id/.*|", "a/.*/" ]
    filter = [ "r|sddlm*|", "a|/dev/sd|" ]
    # filter = [ "a|sddlm*|", "r|/dev/sd|" ]

    cache = "/etc/lvm/.cache"

    # You can turn off writing this cache file by setting this to 0.
    write_cache_state = 1

    # Advanced settings.

    # List of pairs of additional acceptable block device types found
    # in /proc/devices with maximum (non-zero) number of partitions.
    # types = [ "fd", 16 ]
    # types = [ "sddlmfdrv", 16 ]

    # the block devices it believes are valid.
    # 1 enables; 0 disables.
    sysfs_scan = 1

    # By default, LVM2 will ignore devices used as components of
    # software RAID (md) devices by looking for md superblocks.
    # 1 enables; 0 disables.
    md_component_detection=0
}
```

**Figure 3.97 Example of Editing the /etc/lvm/lvm.conf File**

7. Execute the VG scan.

   To re-create the volume group on the HDLM device, execute the `vgscan` command as shown in the following example:

```
# vgscan
Reading all physical volumes.  This may take a while...
Found exported volume group "vg01" using metadata type lvm2
```

8. Import the volume group.

   In the following example, the command imports the `vg01` volume group:

```
# vgimport vg01
Volume group "vg01" successfully imported
```

9. Execute the `pvscan` command to check that the environment has been changed successfully.

   The following shows an example of executing the `pvscan` command:

```
# pvscan
PV /dev/sda1   VG vg01   lvm2 [468.00 MB / 368.00 MB free]
PV /dev/sda2   VG vg01   lvm2 [548.00 MB / 548.00 MB free]
Total: 2 [1016.00 MB] / in use: 2 [1016.00 MB] / in no VG: 0 [0   ]
```

10. Activate the volume group.

    In the following example, the logical volume subject to the move belongs to the vg01 volume group:

```
# vgchange -ay vg01
  1 logical volume(s) in volume group "vg01" now active
```

11. If the /etc/fstab file was edited (a section was commented out) in step 5, return the commented out section to its original state.
12. Mount the logical volume that was moved.

    In the following example, the logical volume that was moved is /dev/vg01/lvol1 and the command mounts the volume on /mnt/lvol1:

```
# mount /dev/vg01/lvol1 /mnt/lvol1
```

### 3.17.6 Uninstalling HDLM

This subsection describes how to uninstall HDLM.

#### 3.17.6.1 Uninstalling HDLM

***Note:***

- If the system initiates path health checking or automatic failback while HDLM is being uninstalled, an error message (KAPL04023) may be output to syslog. However, HDLM operations are not affected.
- If you uninstall HDLM, the error logs (/var/tmp/hdlminstlog/installhdlm[01-10].log) generated during installation of HDLM are not uninstalled. Delete the /var/tmp/hdlminstlog/installhdlm[01-10].log files as needed.

  For details on the installhdlm[01-10].log files, see section 3.4.2.
- When using release 10.1.0.3.0 or later of Oracle RAC 10g, the value set in the MISSCOUNT parameter, for which the threshold of the I/O timeout of the voting disk is set, must be returned to the value that was set before HDLM was installed. For details on how to set the value of the MISSCOUNT parameter, contact the company with which you have a support service contract for Oracle.
- When using release 10.2.0.2.0 or later of Oracle RAC 10g, the value set in the DISKTIMEOUT parameter, must be returned to the value that was set before HDLM was installed. For details on how to set the value of the DISKTIMEOUT parameter, contact the company with which you have a support service contract for Oracle.

– 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 on to Linux as a root user.
2. Stop all processes and services that are using HDLM-managed paths.

   Stop all DBMS, various application processes, and services that are using a path managed by HDLM.
3. Unmount all HDLM devices.

   Execute the following command:

```
# umount mount-point
```

4. Change the setting of the environment variable PATH in the root environment setup file.

   Remove `/opt/DynamicLinkManager/bin` from the environment variable `PATH` in the environment setup file for the shell used by the root user.
5. Execute the uninstall command.

   Execute the following command to uninstall HDLM:

```
# rpm -e HDLM
```

6. Make sure that the uninstall processing has finished.

   Execute the following command to make sure that the uninstall processing has finished. The following shows the information that is output.

```
# rpm -qi HDLM
package HDLM is not installed
```

**Figure 3.98 Example of Checking HDLM Uninstallation**

7. Start the processes and services stopped in step 2.

   The cluster software and volume management software will start up after the host starts.

### 3.17.6.2 Uninstalling HDLM When an HDLM Device Is Used as a Boot Disk

This subsection explains how to uninstall HDLM from a multi-path boot disk environment that uses an HDLM device, and to change the environment to a boot disk environment that uses a SCSI device. (Note that if settings are incorrect, the OS might not start.)

The following procedure describes how to set up a boot disk environment that uses LVM2. This procedure assumes that grub is used as the boot loader.

To uninstall HDLM in such a configuration:

1. Log in to Linux as a user with root privileges.
2. See the `/etc/fstab` file to check the correspondence between the root directory and the HDLM device.

   Make sure that the definitions in the `/etc/fstab` file contain the HDLM device. Make sure that the root directory is associated with the HDLM device. Also make sure that the mount points for HDLM management-target devices other than `/boot` directory and `/boot/efi` directory are defined so that the HDLM device will be mounted.

   – When the boot loader is LILO or GRUB

     Figure 3.99 shows an example of `/etc/fstab` file.

| | | | | |
|---|---|---|---|---|
| /dev/sddlmaa2 | / | ext2 | defaults | 1 1 |
| /dev/sda1 | /boot | ext2 | defaults | 1 2 |
| /dev/sddlmaa4 | /tmp | ext2 | defaults | 1 2 |
| /dev/sddlmaa5 | /var | ext2 | defaults | 1 2 |
| /dev/sddlmaa6 | /usr | ext2 | defaults | 1 2 |
| none | /dev/pts | devpts | gid=5,mode=620 | 0 0 |
| none | /proc | proc | defaults | 0 0 |
| none | /dev/shm | tmpfs | defaults | 0 0 |
| /dev/sddlmaa3 | swap | swap | defaults | 0 0 |
| /dev/cdrom | /mnt/cdrom/ | udf,iso9660 | noauto,owner,kudzu,ro | 0 0 |
| /dev/fd0 | mnt/floppy | auto | noauto,owner,kudzu | 0 0 |

**Figure 3.99 Display Example of /etc/fstab file When the Boot Loader Is LILO or GRUB**

   – When the boot loader is ELILO running on an IPF host

     Figure 3.100 shows an example of `/etc/fstab` file.

| | | | | |
|---|---|---|---|---|
| /dev/sddlmaa2 | / | ext3 | defaults | 1 1 |
| /dev/sda1 | /boot/efi | vfat | defaults | 0 0 |
| /dev/sddlmaa4 | /tmp | ext3 | defaults | 1 2 |
| /dev/sddlmaa5 | /var | ext3 | defaults | 1 2 |
| /dev/sddlmaa6 | /usr | ext3 | defaults | 1 2 |
| none | /dev/pts | devpts | gid=5,mode=620 | 0 0 |
| none | /proc | proc | defaults | 0 0 |
| none | /dev/shm | tmpfs | defaults | 0 0 |
| /dev/sddlmaa3 | swap | swap | defaults | 0 0 |
| /dev/cdrom | /mnt/cdrom | udf ,iso9660 | noauto,owner,kudzu,ro | 0 0 |

**Figure 3.100 Display Example of /etc/fstab file When the Boot Loader Is ELILO**

   – When using LVM2 Figure 3.101 shows an example of `/etc/fstab` file.

```
/dev/VolGroup00/LogVol00   /          ext3    defaults            1 1
#LABEL=/boot               /boot      ext3    defaults            1 2
/dev/sda1                  /boot      ext3    defaults            1 2
none                       /dev/pts   devpts  gid=5,mode=620      0 0
none                       /dev/shm   tmpfs   defaults            0 0
none                       /proc      proc    defaults            0 0
none                       /sys       sysfs   defaults            0 0
/dev/VolGroup00/LogVol02   /tmp       ext3    defaults            1 2
/dev/VolGroup00/LogVol03   /var       ext3    defaults            1 2
/dev/VolGroup00/LogVol01   swap       swap    defaults            0 0
```

**Figure 3.101 Display Example of /etc/fstab file When using LVM2**

3. Execute the HDLM-configuration definition utility (`dlmcfgmgr`) with the `-v` parameter specified to check the correspondence between the HDLM device and the SCSI device.

   If the SCSI device name has been changed by using the `udev` function, specify the `-udev` parameter also.

   Figure 3.102 shows an example of executing the `dlmcfgmgr` utility with the `-v` parameter specified. Figure 3.103 shows an example of executing the dlmcfgmgr utility with the `-v` and `-udev` parameters specified.

```
# dlmcfgmgr -v
HDevName       Management   Device         Host   Channel Target  Lun
/dev/sddlmaa   configured   /dev/sda          0    0          0     0
                            /dev/sdb          0    0          1     0
KAPL10302-I /sbin/dlmcfgmgr completed normally.
```

**Figure 3.102 Example of Executing the dlmcfgmgr Utility with the -v Parameter Specified**

```
# dlmcfgmgr -v -udev
HDevName       Management   Device         Host   Channel Target  Lun    Udev
/dev/sddlmaa   configured   /dev/sda          0    0          0     0    /dev/aaaaaaaa
KAPL10302-I /sbin/dlmcfgmgr completed normally.
```

**Figure 3.103 Example of Executing the dlmcfgmgr Utility with the -v and -udev Parameters Specified**

When LVM2 is used, go to step 8.

4. When LVM2 is not used, edit the `/etc/fstab` file.

   Change the definition so that the SCSI device, rather than the HDLM device, will be mounted on the root directory. Also, change the definition so that the swap partition uses the SCSI device.

   – When the boot loader is LILO or GRUB

     Figure 3.104 shows an example of editing the `/etc/fstab` file.

```
#/dev/sddlmaa2   /             ext2         defaults                  1 1
/dev/sda2        /             ext2         defaults                  1 1
/dev/sda1        /boot         ext2         defaults                  1 2
#/dev/sddlmaa4   /tmp          ext2         defaults                  1 2
/dev/sda4        /tmp          ext2         defaults                  1 2
#/dev/sddlmaa5   /var          ext2         defaults                  1 2
/dev/sda5        /var          ext2         defaults                  1 2
#/dev/sddlmaa6   /usr          ext2         defaults                  1 2
/dev/sda6        /usr          ext2         defaults                  1 2
none             /dev/pts      devpts       gid=5,mode=620            0 0
none             /proc         proc         defaults                  0 0
none             /dev/shm      tmpfs        defaults                  0 0
#/dev/sddlmaa3   swap          swap         defaults                  0 0
/dev/sda3        swap          swap         defaults                  0 0
/dev/cdrom       /mnt/cdrom/   udf ,iso9660 noauto,owner,kudzu,ro     0 0
/dev/fd0         mnt/floppy    auto         noauto,owner,kudzu        0 0
```

**Figure 3.104 Example of Editing /etc/fstab file When the Boot Loader Is LILO or GRUB**

– When the boot loader is ELILO running on an IPF host

  Figure 3.105 shows an example of editing the `/etc/fstab` file.

```
#/dev/sddlmaa2   /             ext3         defaults                  1 1
/dev/sda2        /             ext3         defaults                  1 1
/dev/sda1        /boot/efi     vfat         defaults                  0 0
#/dev/sddlmaa4   /tmp          ext3         defaults                  1 2
/dev/sda4        /tmp          ext3         defaults                  1 2
#/dev/sddlmaa5   /var          ext3         defaults                  1 2
/dev/sda5        /var          ext3         defaults                  1 2
#/dev/sddlmaa6   /usr          ext3         defaults                  1 2
/dev/sda6        /usr          ext3         defaults                  1 2
none             /dev/pts      devpts       gid=5,mode=620            0 0
none             /proc         proc         defaults                  0 0
none             /dev/shm      tmpfs        defaults                  0 0
#/dev/sddlmaa3   swap          swap         defaults                  0 0
/dev/sda3        swap          swap         defaults                  0 0
/dev/cdrom       /mnt/cdrom    udf ,iso9660 noauto,owner,kudzu,ro     0 0
```

**Figure 3.105 Example of Editing /etc/fstab file When the Boot Loader Is ELILO**

5. Copy the registration of the HDLM device specification.
6. Comment out the existing HDLM device specification by placing a hash mark (`#`) at the beginning of each line.
7. Convert the copied registration from an HDLM device specification to a SCSI device specification according to the correspondence between the HDLM device and the SCSI device that you checked in step 3.

   Go to step 9.
8. When LVM2 is used, edit the `/etc/lvm/lvm.conf` file so that the LVM2 recognizes the SCSI device, not the HDLM device.

   Figure 3.106 show examples of editing the `/etc/lvm/lvm.conf` file.

```
linux: /etc/lvm # vi lvm.conf
# This section allows you to configure which block devices should
# be used by the LVM system.
devices {

                    :
                    :

    # filter = [ "a|sddlm*|", "r|/dev/sd|" ]
    filter = [ "r|sddlm*|", "a|/dev/sd|" ]

                    :
                    :

     types = [ "fd", 16 ]
     #types = [ "sddlmfdrv", 16 ]

                    :
                    :

     md_component_detection=0
}
```

**Figure 3.106 Example of Editing /etc/fstab file**

Comment out the existing `filter` and `types` lines, and add the shaded lines in the figure.

In addition, set `md_component_detection` to its original value before installation.

9. Edit the boot loader configuration file for SCSI startup.

   Edit the boot loader configuration file as shown in Figure 3.107, Figure 3.108, Figure 3.109, and Figure 3.110.

   The detailed procedure is explained in steps 10 to 13.

   Separate examples are provided for when the boot loader is LILO, when the boot loader is GRUB, when the boot loader is ELILO running on an IPF host, and when LVM2 is used.

   – When the boot loader is LILO

     Figure 3.107 shows an example of editing the configuration file.

     After modifying the definition, execute `/sbin/lilo` command to activate the new setting. You can delete the configuration for specifying the HDLM device without causing any problems.

```
boot=/dev/sda
map=/boot/map
install=/boot/boot.b
prompt
timeout=50
linear
#default=HDLM_32.0.1.ELsmp
default=linux

image=/boot/vmlinuz-2.4.21-32.0.1.ELsmp
    label=HDLM_32.0.1.ELsmp
    initrd=/boot/initrd-hdlm-2.4.21-32.0.1.ELsmp.gz
    read-only
    append="option ramdisk_size=8517 option"

image=/boot/vmlinuz-2.4.21-32.0.1.ELsmp
    label=linux
    initrd=/boot/initrd-2.4.21-32.0.1.ELsmp
    read-only
    append="option"
    root=/dev/sda2
```

Legend:
  option:  option that depends on the user environment is specified

**Figure 3.107 Example of Editing /etc/lilo.conf file**

- When the boot loader is GRUB

  Figure 3.108 shows an example of editing the configuration file. Make sure that the value specified in `default` is correct. You can delete the configuration for specifying the HDLM device without causing any problems.

```
#default=0
default=1
timeout=10
splashimage=(hd0,0)/grub/splash.xpm.gz

title Red Hat Enterprise Linux AS 3.0 (HDLM-32.0.1.ELsmp)
   root (hd0,0)
   kernel /vmlinuz-2.4.21-32.0.1.ELsmp ro option ramdisk_size=8517 option
   initrd /initrd-hdlm-2.4.21-32.0.1.ELsmp.gz

title Red Hat Enterprise Linux AS 3.0 (2.4.21-32.0.1.ELsmp)
   root (hd0,0)
   kernel /vmlinuz-2.4.21-32.0.1.ELsmp ro root=/dev/sda2 option
   initrd /initrd- 2.4.21-32.0.1.ELsmp
```

Legend:
  option: option that depends on the user environment is specified

**Figure 3.108 Example of Editing /etc/grub.conf file**

- When the boot loader is ELILO running on an IPF host

  Figure 3.109 shows an example of editing the configuration file. You can delete the configuration for specifying the HDLM device without causing any problems.

```
prompt
timeout=50
#default= HDLM_2.4.21-27.EL
default= 2.4.21-27.EL-sd

image=vmlinuz-2.4.21-27.EL
    label=HDLM_2.4.21-27.EL
    initrd=initrd-hdlm-2.4.21-27.EL.img
    read-only
    append="option ramdisk_size=16000 option"

image=vmlinuz-2.4.21-27.EL
    label=2.4.21-27.EL-sd
    initrd=initrd-2.4.21-27.EL.img
    read-only
    append="option"
    root=/dev/sda2


Legend:
  option:  option that depends on the user environment is specified
```

**Figure 3.109 Example of Editing /etc/elilo.conf file**

– When using LVM2

Figure 3.110 shows an example of editing the configuration file. You can delete the configuration for specifying the HDLM device without causing any problems.

```
#default=0
default=1
timeout=5
splashimage=(hd1,0)/grub/splash.xpm.gz
Hiddenmenu

title Red Hat Enterprise Linux AS (HDLM 2.6.9-11.EL)
   root (hd1,0)
   kernel /vmlinuz-2.6.9-11.EL ro rhgb quiet
   initrd /initrd-hdlm-2.6.9-11.EL.img

title Red Hat Enterprise Linux AS (2.6.9-11.EL)
   root (hd1,0)
   kernel /vmlinuz-2.6.9-11.EL ro root=/dev/VolGroup00/LogVol00 rhgb quiet
#    initrd /initrd-hdlm-2.6.9-11.EL.img
   initrd /initrd-2.6.9-11.EL.img
```

**Figure 3.110 Example of Editing /etc/grub.conf file**

10. Copy the configuration that was used for startup from the HDLM device.

11. Assign a name to the copied configuration.

    You can assign any name that indicates that the configuration is used for specifying the SCSI device.

    Change `label` if the boot loader is LILO, or ELILO running on an IPF host. If LILO is the boot loader, you must then execute `/sbin/lilo` command to activate the new setting.

    Change `title` if you will use GRUB as the boot loader.

12. In `root`, specify the SCSI device you checked in step 3.

When LVM2 is used, specify `root` as the name of the device mounted to `/` in the `/etc/fstab` file.

13. In `initrd`, specify the initial RAM disk image file for the SCSI device. Delete `ramdisk_size` that was used for startup from the HDLM device if Red Hat Enterprise Linux AS3/ES3 is being used.

    When editing the file, be careful of the following:

    - When an option that depends on the user environment is specified:

      Be careful not to delete the option.

    - When the boot loader is LILO, or ELILO running on an IPF host, and an option that depends on the user environment is not specified:

      Comment out or delete `append`.

    Create the initial RAM disk image file for the SCSI device according to the `mkinitrd` command documentation for the OS.

14. Shut down the host.

    Execute the following command to shut down the host:

```
# shutdown -h now
```

15. Change the configuration from a multi-path configuration to a single-path configuration.

16. Start the host

    When LVM2 is used, go to step 18.

17. Make sure that the SCSI device is used by the file system and the swap partition.

    Make sure of the following three things: that the root directory is associated with the SCSI device, that the SCSI device is mounted, and that the SCSI device is assigned to the swap partition.

    - When the boot loader is LILO or GRUB

      Figure 3.111 shows an example of executing the `mount` command.

```
# mount
/dev/sda2 on / type ext2 (rw)
none on /proc type proc (rw)
none on /dev/pts type devpts (rw,gid=5,mode=620)
usbdevfs on /proc/bus/usb type usbdevfs (rw)
/dev/sda1 on /boot type ext2 (rw)
/dev/sda4 on /tmp type ext2 (rw)
/dev/sda5 on /var type ext2 (rw)
/dev/sda6 on /usr type ext2 (rw)
none on /dev/shm type tmpfs (rw)
```

**Figure 3.111 Example of Executing the mount Command When the Boot Loader Is LILO or GRUB**

- When the boot loader is ELILO running on an IPF host

  Figure 3.112 shows an example of executing the `mount` command.

```
# mount
/dev/sda2 on / type ext3 (rw)
none on /proc type proc (rw)
none on /dev/pts type devpts (rw,gid=5,mode=620)
usbdevfs on /proc/bus/usb type usbdevfs (rw)
/dev/sda1 on /boot/efi type vfat (rw)
/dev/sda4 on /tmp type ext3 (rw)
/dev/sda5 on /var type ext3 (rw)
/dev/sda6 on /usr type ext3 (rw)
none on /dev/shm type tmpfs (rw)
```

**Figure 3.112 Example of Executing the mount Command When the Boot Loader Is ELILO**

Go to step 19.

18. Check the relationship between the volume group and HDLM devices.

    Execute the following command to make sure that the physical volumes that make up the volume group are not HDLM deivces.

19. Delete the initial RAM disk image file that was created with the `dlmmkinitrd` utility for supporting a boot disk.

    To do this, execute the following command:

    – When an IA32 host is used:

    ```
    # rm /boot/initrd-2.4.21-27.EL.img
    ```

    – When an IPF host is used:

    ```
    # rm /boot/efi/efi/redhat/initrd-hdlm-2.4.21-27.EL.img
    ```

20. Stop HDLM that was used in a location other than the root directory.

    For details on how to stop HDLM, see steps 2 to 4 in 3.17.1 and 3.17.6.1.

21. Uninstall HDLM.

    To uninstall HDLM, execute the following command:

    ```
    # rpm -e HDLM
    ```

22. Make sure that the uninstallation has been completed.

    You must execute the following command to make sure that the uninstallation has been completed. The following information is output.

```
# rpm -qi HDLM
package HDLM is not installed
```

**Figure 3.113 Example of Checking HDLM Uninstallation**

# 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

Note the following when using HDLM:

When your OS is Red Hat Enterprise Linux AS3/ES3, see section 4.1.1 and 4.1.2.

When your OS is SUSE LINUX Enterprise Server 9 or Red Hat Enterprise Linux AS4/ES4, see section 4.1.1, 4.1.2 and 4.1.3.

When your OS is SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10, see section 4.1.1, 4.1.3, and 4.1.4.

### 4.1.1 Notes Common to OSs

The following notes are common to OSs:

- When a path error exists, the period of time required to return a response to the application that made an I/O request to HDLM depends on the timeout values specified for the switch and HBA.

  Therefore, if an error occurs in all the paths for a single LU, this response time will be the total period of time required for the corresponding switches and HBAs to detect a timeout for the I/O requested for each path.

  For details on the timeout values for switches and HBAs, see the appropriate manual for the device. Figure 4.1 gives an overview of the period required to return a response: from the time an application requests an I/O, to the time the application receives the response.

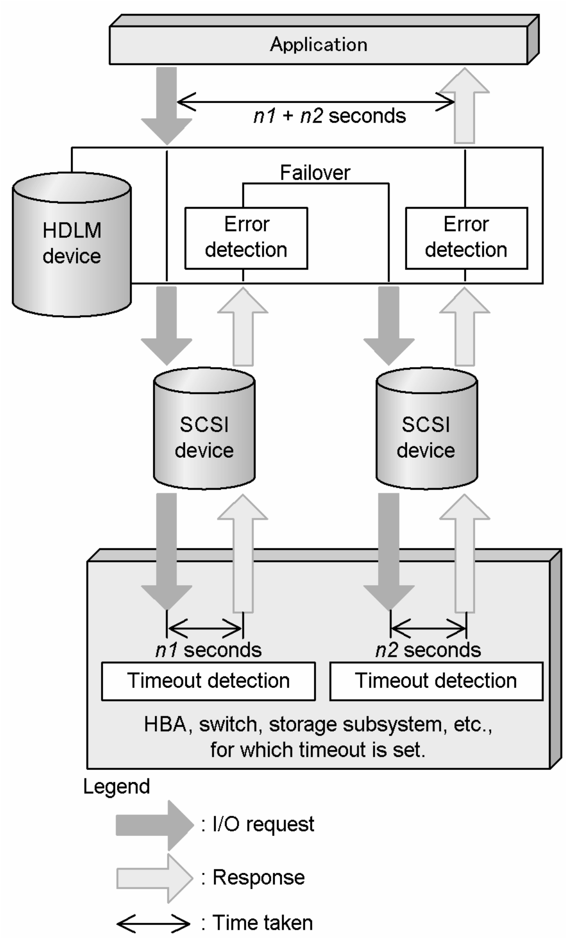


**Figure 4.1 Overview of the Period Required to Respond to an Application's I/O Request**

As shown in the preceding diagram, when an HDLM device has two paths (SCSI devices), the maximum period of time required to respond to the application's I/O request is n1 + n2 seconds; where n1 indicates the timeout value specified for the path that uses SCSI device A, and n2 indicates the timeout value specified for the path that uses SCSI device B.

- When using HDLM in a cluster environment, a node in the cluster may fail over before the path completes failover. To avoid this, when you set up the failover timeout value for the node by using cluster software, make sure that you specify a period longer than the response time that is calculated as described in the preceding note.
- If a path error occurs while creating a file system, or formatting, or executing `fsck`, the operation may not finish. In such a case, perform the operation again after restoring the path error.
- The name of the HDLM device file will not be displayed on the Hardware Browser of Red Hat Enterprise Linux.
- The Linux functionality that adds `LABEL=` to a SCSI device is not supported in HDLM. Do not use this functionality. With HDLM, access can constantly be made to the same LU if the name of the HDLM device file is the same.

- If multiple `shred` commands are executed simultaneously for an HDLM device, and if the load balance algorithm being used is the extended round robin algorithm, then the CPU usage may become very high temporarily. Please ensure that multiple `shred` commands do not execute simultaneously.

- If the Oracle command `oracleasm` is executed for HDLM devices, the KAPL05023-E message might be output to syslog.

  If the following message indicating the successful completion of the `oracleasm` command is output after this message, there is no problem with HDLM operation.

```
oracleasm:  succeeded
```

- If you use HDLM with Oracle RAC 10g and use ASMLib of Oracle RAC 10g, you must change the settings for the ASMLib configuration file.

  After installing ASMLib, do the following for each node:

  a) Execute the following command to create the ASMLib configuration file:

  ```
  # /etc/init.d/oracleasm configure
  ```

  b) Open the ASMLib configuration file (`/etc/sysconfig/oracleasm`).

  c) Change the `ORACLEASM_SCANORDER` line as follows:

  ```
  ORACLEASM_SCANORDER ="sddlm"
  ```

  If an ASM disk defined in an HDLM device already exists, perform this procedure, and then restart all the nodes.

  This step must be done to enable the settings.

- If all the following conditions are satisfied, even a path in the normal status might become Offline(E) or Online(E):
  - A storage subsystem that has a SATA drive is being used.
  - There is a heavy I/O load on the storage subsystem.
  - Path health checking is set to ON.

  If this problem happens, reduce the I/O load, and then return the path status to Online by performing an online operation.

### 4.1.2 Notes When Using Red Hat Enterprise Linux AS3/ES3 and Red Hat Enterprise Linux AS4/ES4

Note the following when your OS is Red Hat Enterprise Linux AS4/ES4:

- Before you update the kernel package, uninstall HDLM. After you have updated the kernel package, re-install HDLM.

  For details on HDLM installation and uninstallation, see section 3.5 and 3.17.6.

### 4.1.3 Notes When Using SUSE LINUX Enterprise Server 9, SUSE LINUX Enterprise Server 10 and Red Hat Enterprise Linux AS4/ES4

Note the following when your OS is SUSE LINUX Enterprise Server 9, SUSE LINUX Enterprise Server 10 or Red Hat Enterprise Linux AS4/ES4:

- SCSI devices corresponding to a HDLM device can be removed dynamically even while the HDLM device is in use. Before doing so, perform the following operations:

  a) Unmount any devices mounted on the HDLM device.

  Execute the `dlmcfgmgr -o` command to exclude the HDLM device from management.

- You cannot create md devices on the HDLM device.

### 4.1.4 Notes When Using SUSE LINUX Enterprise Server 9 and SUSE LINUX Enterprise Server 10

Note the following when your OS is SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10:

- If you execute the `sync` command on the file system that is mounted by Ext3/reiserfs, the following message is output to syslog, but there is no problem with HDLM operations.

```
JBD: barrier-based sync failed on sddlmaa1 - disabling barriers
```

- Before you update the OS service pack, uninstall HDLM. After you have updated the OS service pack, re-install HDLM.

  For details on HDLM installation and uninstallation, see section 3.5 and 3.17.6.

- If the number of I/O operations that have occurred in the system exceeds the execution performance of the system (in an environment in which all of the conditions below exist), the KAPL05008-E message and the call trace of a kernel might be output.

  – Your OS is SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10.

  – TagmaStore WMS exists in HDLM management-target devices.

  – In the utility for setting HDLM driver option (`dlmsetopt`), the number of retries for allocating memory is set to 0 through 2,147,483,646.

  In this case, the system might become unstable, and an error might occur in the I/O that is being executed. In the `dlmsetopt` utility, set a larger value or -1 for the number of retries for allocating memory, and then restart the system. When a call trace is output even after the retry count is set to -1, restart the system again. Also, consider adding memory and stopping unnecessary processes.

## 4.2 HDLM Operations Using Commands

This section explains how to use the HDLM commands to utilize HDLM functions. For details on commands, 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 an HDLM 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 sddlmaa
Paths:000004 OnlinePaths:000004
PathStatus   IO-Count   IO-Errors
Online        0               0
PathID   PathName                            DskName                    iLU    ChaPort Status  Type    IO-Count  IO-Errors  DNum     HDevName
000000   0000.0000.0000000000000000.0000   HITACHI .DF600F        .0115 0000  0A      Online  Own            0          0        0  sddlmaa
000020   0000.0000.0000000000000001.0000   HITACHI .DF600F        .0115 0000  1A      Online  Non            0          0        0  sddlmaa
000040   0001.0000.0000000000000000.0000   HITACHI .DF600F        .0115 0000  0A      Online  Own            0          0        0  sddlmaa
000041   0001.0000.0000000000000001.0000   HITACHI .DF600F        .0115 0000  1A      Online  Non            0          0        0  sddlmaa
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 section 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 example shows how to execute 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 section 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 is an example of executing the command:

```
# /opt/DynamicLinkManager/bin/dlnkmgr online -hba 0010.0000
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 is 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 is 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 section 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 is an example of executing the command:

```
# /opt/DynamicLinkManager/bin/dlnkmgr offline -hba 0010.0000
```

```
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 is 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 an HDLM command.

To display LU information, execute the `dlnkmgr` command's `view` operation with the `-lu` parameter specified. The following is an example of executing the command:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -lu
Product       : 9500V
SerialNumber  : 0115
LUs           : 5
```

```
iLU  HDevName Device   PathID Status
0000 sddlmac  /dev/sds 000002 Online
              /dev/sdd 000007 Online
              /dev/sdi 000012 Online
              /dev/sdn 000017 Online
0001 sddlmad  /dev/sdt 000003 Online
              /dev/sde 000008 Online
              /dev/sdj 000013 Online
              /dev/sdo 000018 Online
0002 sddlmae  /dev/sdu 000004 Online
              /dev/sdf 000009 Online
              /dev/sdk 000014 Online
              /dev/sdp 000019 Online
0003 sddlmaa  /dev/sdq 000000 Online
              /dev/sdv 000005 Online
              /dev/sdg 000010 Online
              /dev/sdl 000015 Online
0004 sddlmab  /dev/sdr 000001 Online
              /dev/sdw 000006 Online
              /dev/sdh 000011 Online
              /dev/sdm 000016 Online
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 section 6.7.

### 4.2.5 Displaying Corresponding Information About an HDLM Device, SCSI Device, and LDEV

This subsection explains how to display corresponding information about an HDLM device, SCSI device, and LDEV by using an HDLM command.

You can display corresponding information about an HDLM device, SCSI device, and LDEV by executing the `dlnkmgr` command's `view` operation with the `-drv` parameter specified. For details on the `view` operation, see section 6.7. The execution result of this operation is displayed on a single line for each path.

The following is an example of executing the `dlnkmgr` command's `view` operation:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -drv
PathID HDevName    Device          LDEV
000000 sddlmaa     /dev/sdag       9970/9980.15001.05B7
000001 sddlmaa     /dev/sdq        9970/9980.15001.05B7
000002 sddlmab     /dev/sdr        9970/9980.15001.05B0
000003 sddlmac     /dev/sds        9970/9980.15001.05B1
000004 sddlmad     /dev/sdt        9970/9980.15001.05B2
000005 sddlmae     /dev/sdu        9970/9980.15001.05B3
000006 sddlmaf     /dev/sdv        9970/9980.15001.05B4
000007 sddlmag     /dev/sdw        9970/9980.15001.05B5
             :
000028 sddlmae     /dev/sdm        9970/9980.15001.05B3
000029 sddlmaf     /dev/sdn        9970/9980.15001.05B4
000030 sddlmag     /dev/sdo        9970/9980.15001.05B5
000031 sddlmah     /dev/sdp        9970/9980.15001.05B6
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 section 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 is 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 is 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 is 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 is 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
Reservation Status           :
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 section 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

For details on how to set each function, see section 6.6.

For example, to set up the log 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 is 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, see section 4.2.7.1.

### 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 is 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 section 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 the license, you can only execute one command at a time. If you attempt to execute more than one dlnkmgr command containing the set operation with the -lic parameter, 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 is 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 is 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)          : 1000
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
Reservation Status           :
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 1000
HDLM Driver Ver      WakeupTime
Alive       xx-xx    yyyy/mm/dd hh:mm:ss
LicenseType 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 is 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
Reservation Status           :
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 1000
HDLM Driver Ver      WakeupTime
Alive       xx-xx    yyyy/mm/dd hh:mm:ss
LicenseType 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 is 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 HDLM to recover from the error.

### 4.3.1 Starting the HDLM Manager

To start the HDLM manager, log in to Linux as a user with root privileges and then execute the following command.

```
# /etc/init.d/DLMManager start
Starting DLMManager:                              [  OK  ]
```

Lower case characters can be used for the command name (`dlmmanager`).

```
# /etc/init.d/dlmmanager start
Starting DLMManager:                              [  OK  ]
```

The startup script that was set up during HDLM installation runs, and starts the HDLM manager.

***Notes:***

- When you start the HDLM manager immediately after stopping it, confirm that the HDLM manager has stopped (`Dead` is displayed when executing `/opt/DynamicLinkManager/bin/dlnkmgr view -sys -msrv`), and then start it.
- Even if the HDLM manager is running, `Dead` might be displayed when you execute `/opt/DynamicLinkManager/bin/dlnkmgr view -sys -msrv`. If this happens, see section 5.4.3.

Important

When stopping the HDLM manager, even if a message indicating that the HDLM manager has successfully stopped, it takes a few seconds more to stop processes. Therefore, if you start the HDLM manager immediately after stopping it, the HDLM manager is regarded as being started, and a message indicating that the HDLM manager has successfully started is output. However, starting of the HDLM manager is cancelled, and it will remain inactivate.

Use one of the following procedures to confirm that the HDLM manager is active:

Use the `dlnkmgr` command's `view` operation:

Execute the following command:

```
# /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.

Use a script for confirmation:

Execute the following command to confirm that the HDLM manager is active:

```
# /etc/init.d/DLMManager status
dlmmgr (pid 14889 14886 14884 14880) is running...
```

***Note:***

If an error occurs, the execution results of the `dlnkmgr` command's view operation might be different from those of the `DLMManager` status command. In this case, see the execution results of the `dlnkmgr` command.

## 4.3.2 Stopping the HDLM Manager

To stop the HDLM manager, log in to Linux as a user with root privileges, and then execute the following command:

```
# /etc/init.d/DLMManager stop
Stopping DLMManager:                      [  OK  ]
```

Lower case characters can be used for the command name (`dlmmanager`).

```
# /etc/init.d/dlmmanager stop
Stopping DLMManager:                      [  OK  ]
```

The stop script that was set up during HDLM installation runs, and stops the HDLM manager.

Use one of the following procedures to confirm that the HDLM manager has stopped:

Use the `dlnkmgr` command's `view` operation:

Execute the following command.

```
# /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.

Use a script for confirmation:

Execute the following command to confirm that the HDLM manager is inactive:

```
# /etc/init.d/DLMManager status
dlmmgr is stopped.
```

***Note:***

If an error occurs, the execution results of the `dlnkmgr` command's view operation might be different from those of the `DLMManager` status command. In this case, see the execution results of the `dlnkmgr` command.

## 4.4 Reconfiguring the HDLM Operating Environment

This chapter describes the operation required for changing the configuration of an HDLM operating environment. To add a new LU or delete an existing LU, to add or reduce a path to or from an existing LU, or to include or exclude an HDLM device to or from the management target, you need to change the configuration of the path for HDLM device.

To do this, use the HDLM-configuration definition utility (`dlmcfgmgr`) or the management program of the storage subsystem.

### 4.4.1 Replacing a Fiber Cable

If there are multiple online paths for a device in an LU, you can replace a desired fiber cable while running your applications by placing offline only the path that goes through the fiber cable to be replaced, and using other paths to continue accesses.

The following steps show an example procedure for replacing a fiber cable. Note that this procedure only applies to replacing a fiber cable.

This example assumes the system configuration shown below:

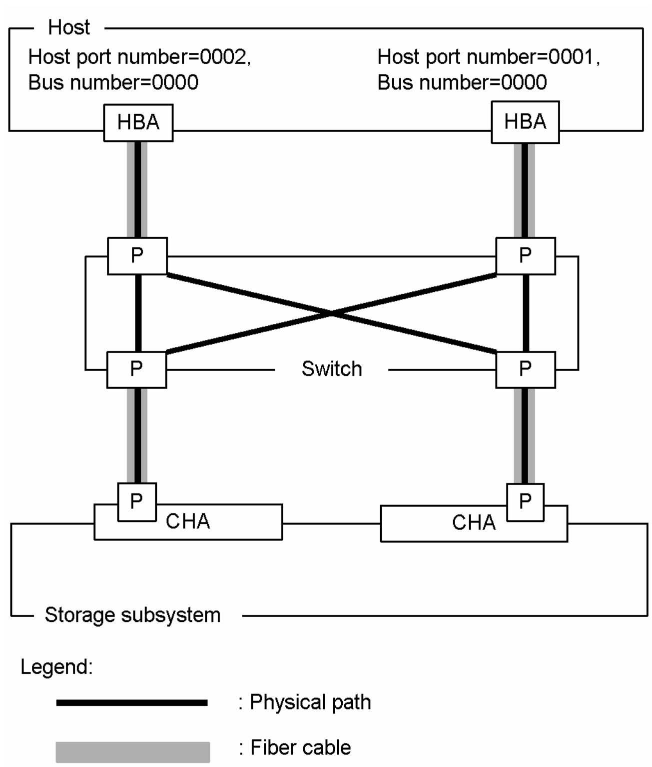


**Figure 4.2 System Configuration for Replacing a Fiber Cable (Explained in the Following Steps)**

To replace a fiber cable:

1. Place in Offline (C) status the path that goes through the fiber cable to be replaced (path that goes through the host bus adapter to which the fiber cable is connected).

   For example, to place in Offline (C) status the path that goes through the host bus adapter with host port number `0001` and bus number `0000`, execute the following command:

   ```
   # /opt/DynamicLinkManager/bin/dlnkmgr offline -hba 0001.0000
   ```

   While the path that goes through the host bus adapter with host port number `0001` is placed Offline, the status of the path that goes through the host bus adapter with host port number `0002` is Online.

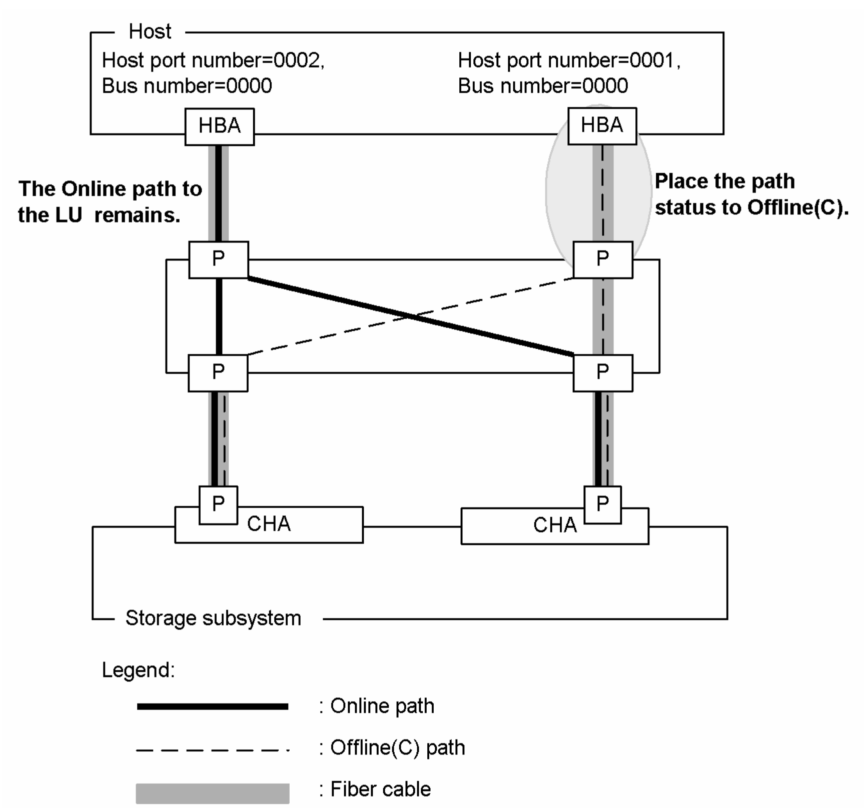


**Figure 4.3 When the Status of the Path That Goes Through the Fiber Cable to Be Replaced Is Offline(C)**

2. Replace the fiber cable for the path that has been placed in Offline(C) status.

   In this example, replace the fiber cable connected to the host bus adapter with host port number `0001`.

3. Place in Online status the path that goes through the replaced fiber cable (that is, the path that goes through the host bus adapter connected to the fiber cable).

   For example, to place in Online status the path that goes through the host bus adapter with host port number `0001` and bus number `0000`, execute the following command:

   ```
   # /opt/DynamicLinkManager/bin/dlnkmgr online -hba 0001.0000
   ```

4. Check the path information.

For example, execute the following command:

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

For details about the path information, see section 6.7.

### 4.4.2 Replacing the Fibre Channel Switch

If there are multiple online paths for a device in an LU, you can replace a desired Fibre Channel switch while running your applications by placing offline only the path that goes through the Fibre Channel switch to be replaced, and using other paths to continue accesses.

The following steps show an example procedure for replacing a Fibre Channel switch. Note that this procedure only applies for replacing a Fibre Channel switch.

This example assumes the system configuration as below:

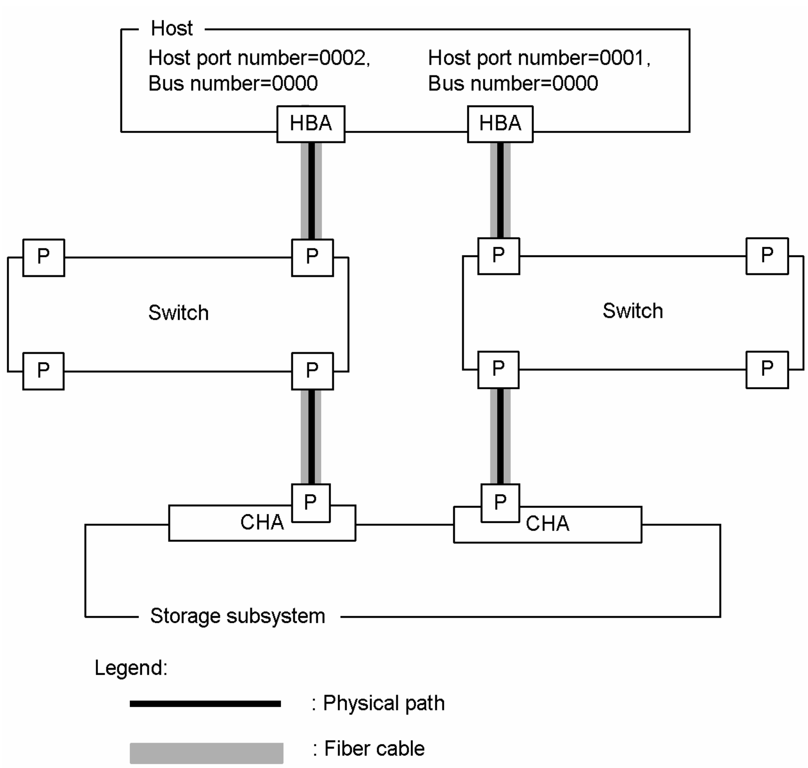


**Figure 4.4 System Configuration for Replacing a Fibre Channel Switch (Explained in the Following Steps)**

To replace the Fibre Channel switch:

1. Place in Offline (C) status the path that goes through the Fibre Channel switch to be replaced (path that goes through the host bus adapter to which the Fibre Channel switch is connected).

   For example, to place in Offline (C) status the path that goes through the host bus adapter with host port number `0001` and bus number `0000`, execute the following command:

```
# /opt/DynamicLinkManager/bin/dlnkmgr offline -hba 0001.0000
```

While the path that goes through the host bus adapter with host port number `0001` is placed Offline, the status of the path that goes through the host bus adapter with host port number `0002` is Online.

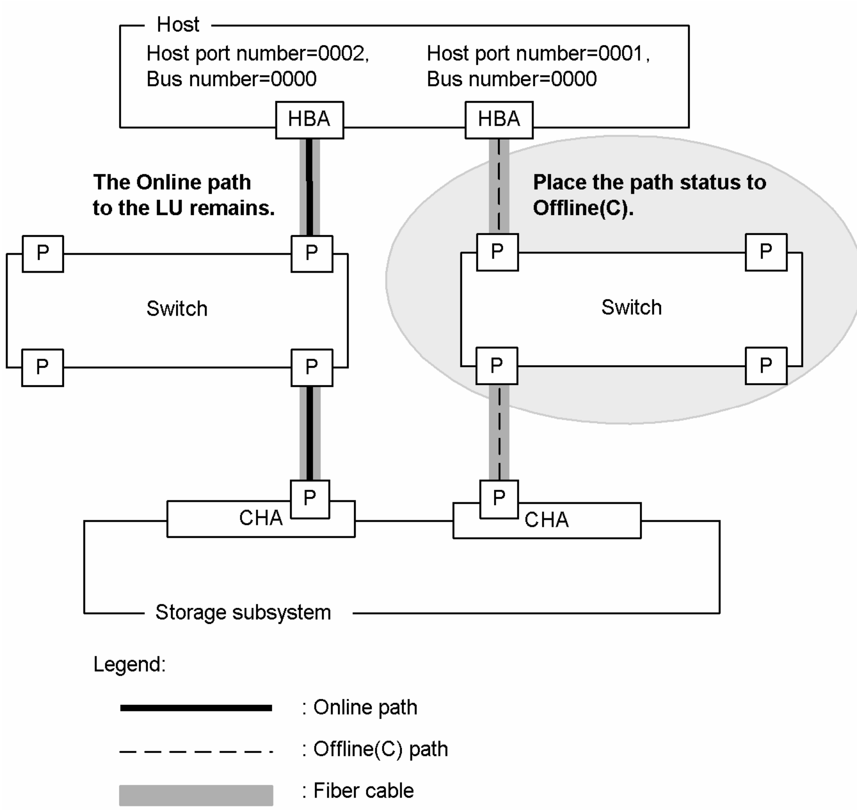


**Figure 4.5 When the Status of the Path That Goes Through the Fibre Channel Switch to Be Replaced Is Offline(C)**

2. Replace the Fibre Channel switch for the path that has been placed in Offline(C) status.

   In this example, replace the Fibre Channel switch connected to the host bus adapter with host port number `0001`.

3. Place in Online status the path that goes through the replaced Fibre Channel switch.

   For example, to place in Online status the path that goes through the host bus adapter with host port number `0001` and bus number `0000`, execute the following command:

```
# /opt/DynamicLinkManager/bin/dlnkmgr online -hba 0001.0000
```

4. Check the path information.

   For example, execute the following command:

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

For details about the path information, see section 6.7.

### 4.4.3 Changing the HDLM Device Configuration

This section describes the actions and procedures to change the HDLM device configuration.

To set up or change (add, change or delete) an LU in a storage subsystem, use the management program of the storage subsystem, rather than the HDLM functionality. For more information on configuring or changing an LU, see the manual provided with the storage subsystem.

All the actions performed via the HDLM-configuration definition utility (`dlmcfgmgr`) are logged in the `/var/opt/DynamicLinkManager/log/dlmcfgmgr1.log` file. By looking in this file, you can examine the change(s) made to the configuration, or the reason why the execution of the `dlmcfgmgr` utility failed. For more information on the `dlmcfgmgr` utility log, see also 2.11.1. For more information on the `dlmcfgmgr` utility, see section 7.3.

#### 4.4.3.1 Notes on Changing the HDLM Device Configuration

Note the following when changing the HDLM device configuration:

- When an HDLM device is excluded from HDLM management, the `AutoPATH_ID`s assigned to the paths for that device will be released. Because of this, after paths for a HDLM device are released, the `AutoPATH_ID`'s of the paths may be displayed in a non-consecutive sequence.
- When an HDLM device is included in HDLM management, HDLM assigns an `AutoPATH_ID` in sequential order (from the lowest available number) to the path of that HDLM device.
- When an HDLM device that is not managed by HDLM is included in HDLM management, HDLM places online all the normal paths to that HDLM device regardless of the previous status (such as `Online(E)` or `Offline(C)`) of the path.
- Do not execute the HDLM-configuration definition utility (`dlmcfgmgr`) with the `-o`, `-i`, or `-u` parameter while the LU is in use (for example, when performing I/O or a mount).
- If a disk partition is created or changed while some of the paths for an LU are disconnected, the disk partition may not be recognized even after the disconnected paths have been recovered. After connecting the disconnected paths, execute the `dlmcfgmgr` utility or restart the host.
- If you are using a QLogic HBA driver, and you add an LU, add a path, or restore the path of an HDLM device that was disconnected at host boot time without restarting the host, perform the following procedure every time you add a SCSI device:

  For SUSE LINUX Enterprise Server 9, Red Hat Enterprise Linux AS4 or Red Hat Enterprise Linux ES4

  a) Execute the following command:

  ```
  # echo "scsi-qlascan" > /proc/scsi/driver-name/host port number
  ```

  `driver-name : qla2xxx` (Name of the driver being used)

  Check that the contents shown below have been output to the `/var/log/messages` file, and then execute the add command for the SCSI device.

```
kernel: dpc(host port number): qla2x00: RESCAN .
kernel: dpc(host port number): qla2x00: RESCAN... done.
```

For SUSE LINUX Enterprise Server 10

a) Execute the following command:

```
# echo "1" > /sys/class/fc_host/hostABC/issue_lip
# echo "- - -" > /sys/class/scsi_host/hostABC/scan
```

*ABC* : host port number

Check that the contents shown below have been output to the `/var/log/messages` file.

```
kernel: SCSI device recognized-SCSI-device-name: ...
kernel: recognized-SCSI-device-name: ...
...
```

#### 4.4.3.2 Adding a New LU

To add a new LU (HDLM device) to the host on which the HDLM is installed:

1. Use the management program of the storage subsystem to allocate an LU, which is not yet allocated to the relevant host, to a port of the storage subsystem that has a path to the relevant host and configure the path.

   For more information on configuring the path, see the manual provided with the storage subsystem.

2. When you want to add an LU without restarting the host, perform step 3 and step 4. When you want to add an LU while restarting the host, go to step 5.

3. Create a SCSI device for each path configured in step 1.

   Since the host has not been restarted, a SCSI device for the LU that was added in step 1 was not created. Execute the command for adding the device shown in Table 4.1 to manually create a SCSI device.

**Table 4.1 Adding or Deleting a SCSI Device**

| Operation | Command |
|---|---|
| Adding a device | ▪ When using a QLogic HBA driver on SUSE LINUX Enterprise Server 9, Red Hat Enterprise Linux AS4, or Red Hat Enterprise Linux ES4, execute the commands in the following order:<br>`# echo "scsi-qlascan" > /proc/scsi/`*driver-name*`/`*adapter-id*<br>`# cat /proc/scsi/`*driver-name*`/`*adapter-id*<br>View the `Id:Lun` part, which is output under `SCSI LUN Information:`, of the output command result, and then execute the following command based on the target ID of the added device:<br>`# echo "scsi add-single-device `*a b c d*`" > /proc/scsi/scsi`<br>▪ When using a QLogic HBA driver on SUSE LINUX Enterprise Server 10, execute the commands in the following order:<br>`# echo "1" > /sys/class/fc_host/host`*a*`/issue_lip`<br>`# echo "- - -" > /sys/class/scsi_host/host`*a*`/scan`<br>*a : host port number* |

| | |
|---|---|
| | ▪ When using an Emulex HBA driver, execute the command as follows:<br>`# echo "scsi add-single-device` *a b c d*`" > /proc/scsi/scsi` |
| Deleting a device | ▪ For SUSE LINUX Enterprise Server 9, SUSE LINUX Enterprise Server 10, Red Hat Enterprise Linux AS4, or Red Hat Enterprise Linux ES4<br>`# dlmcfgmgr -o logical-device-file-name-of-the-relevant-HDLM-device`<br>`# echo "scsi remove-single-device` *a b c d*`" > /proc/scsi/scsi`<br>`# dlmcfgmgr -i logical-device-file-name-of-the-relevant-HDLM-device`<br>▪ For Red Hat Enterprise Linux AS3 or Red Hat Enterprise Linux ES3<br>`# echo "scsi remove-single-device` *a b c d*`" > /proc/scsi/scsi` |
| | *a*: Host ID (host port number)<br>*b*: Channel number (bus number)<br>*c*: Target ID (target ID)<br>*d*: Lun (host LU number)<br>*driver-name*: `qla2200` or `qla2300` (driver name in use)<br>*adapter-id*: host port number<br>Example 1 (when using a QLogic HBA driver)<br>`# echo "scsi-qlascan" > /proc/scsi/qla2200/1`<br>Example 2 (adding a device)<br>`# echo "scsi add-single-device 1 0 1 1" > /proc/scsi/scsi`<br>Example 3 (deleting a device)<br>`# echo "scsi remove-single-device 1 0 1 1" > /proc/scsi/scsi` |

4. Execute the HDLM-configuration definition utility (`dlmcfgmgr -r`).

   Based on the SCSI device created in step 2, an HDLM device for the path configured in step 1 is added.

   The following is an example of executing the `dlmcfgmgr` utility:

```
# /sbin/dlmcfgmgr -r
```

5. Restart the host on which the HDLM is installed.

   If you performed step 3 and step 4, the host does not have to be restarted. Go to step 6.

   While restarting, the HDLM device mapped with the path configured in the above step 1 is added automatically.

6. Refer to `/var/opt/DynamicLinkManager/log/dlmcfgmgr1.log` or execute the `dlmcfgmgr` utility (`dlmcfgmgr -v`) to check if an LU was created with a new HDLM device name.

   The following is an example of executing the `dlmcfgmgr` utility:

```
# /sbin/dlmcfgmgr -v
```

### 4.4.3.3 Adding a Path to an Existing LU

To add a path to an existing LU (HDLM device):

1. Use the management program of the storage subsystem to configure an additional path from an LU mapped with the relevant HDLM device to the relevant host.

   For more information on configuring an additional path, see the manual provided with the storage subsystem.

2. When you want to add a path to the LU without restarting the host, perform step 3 and step 4. When you want to add a path to the LU while restarting the host, go to step 5.

3. Create a SCSI device for each path configured in step 1.

   Since the host has not been restarted, the SCSI device for the path that was added in step 1 was not created. Execute the command for adding the device shown in Table 4.1 to manually create a SCSI device.

4. Execute the HDLM-configuration definition utility (`dlmcfgmgr -r`).

   Based on the SCSI device created in step 2, the path configured in step 1 is added as a path for the HDLM device.

   The following is an example of executing the `dlmcfgmgr` utility:

```
# /sbin/dlmcfgmgr -r
```

5. Restart the host on which the HDLM is installed.

   If you performed step 3 and step 4, the host does not have to be restarted. Go to step 6.

   While restarting, the additional path configured in the above step 1 is automatically added as a path to the relevant HDLM device.

6. Refer to `/var/opt/DynamicLinkManager/log/dlmcfgmgr1.log` or execute the `dlmcfgmgr` utility (`dlmcfgmgr -v`) to check if the number of paths of the relevant HDLM device has increased.

   The following is an example of executing the `dlmcfgmgr` utility:

```
# /sbin/dlmcfgmgr -v
```

#### 4.4.3.4 Deleting an Existing LU

To delete an HDLM device that corresponds to an existing LU:

1. Stop all accesses to the relevant LU.

2. Use the management program of the storage subsystem to cancel all of the path allocations (that is, to delete the paths) to the relevant LU from ports of the storage subsystem.

   For details on canceling path allocations (deleting paths), see the manual provided with the storage subsystem.

3. If you want to delete an HDLM device that corresponds to the LU without restarting the host, perform steps 4 and 5. If you want to delete the LU and restart the host, go to step 6.

4. Delete the SCSI device for the path whose allocation was cancelled in step 2.

   Since the host has not been restarted, the SCSI device for the path whose allocation was cancelled in step 2 was not deleted. To delete the SCSI device, you need to manually execute the command for deleting the device shown in Table 4.1.

5. Execute the HDLM-configuration definition utility (`dlmcfgmgr -r`).

   Apply, to the HDLM device, the information about the SCSI device that was deleted in step 4.

   The following is an example of executing the `dlmcfgmgr` utility:

```
# /sbin/dlmcfgmgr -r
```

6. Restart the host on which HDLM is installed.

   If you performed steps 4 and 5, you do not need to restart the host. Go to step 7.

7. Execute the `dlmcfgmgr` utility (`dlmcfgmgr -v`) or the `view` operation (`dlnkmgr view -drv`).

   Make sure that all the SCSI device names corresponding to the LU for which path allocations were cancelled in step 2 are represented with a hyphen (-) (optional).

   The following is an example of executing the `dlmcfgmgr` utility with the `-v` option:

```
# /sbin/dlmcfgmgr -v
```

   The following is an example of executing the `view` operation with the `-drv` option:

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

8. Execute the `dlmcfgmgr` utility (`dlmcfgmgr -u` *logical_device_file_name_of_HDLM_device*) on all the HDLM devices corresponding to the LU for which path allocations were cancelled in step 2 to unregister the HDLM devices.

   Path information on the relevant HDLM device is deleted.

   The following is an example of executing the `dlmcfgmgr` utility:

```
# /sbin/dlmcfgmgr -u logical_device_file_name_of_HDLM_device
```

9. Check `/var/opt/DynamicLinkManager/log/dlmcfgmgr1.log` or execute the `dlmcfgmgr` utility (`dlmcfgmgr -v`) to make sure that the HDLM device that corresponds to the deleted LU was deleted.

   The following is an example of executing the `dlmcfgmgr` utility:

```
# /sbin/dlmcfgmgr -v
```

### 4.4.3.5 Deleting a Path to an Existing LU

To delete a path to an existing LU (HDLM device):

1. Stop all accesses to the relevant LU.
2. Use the management program of the storage subsystem to delete a path to the relevant LU.

   For more information on deleting a path, see the manual provided with the storage subsystem.

3. When you want to delete a path to the LU without restarting the host, go to step 4. When you want to delete a path while restarting the host, go to step 5.
4. Delete the SCSI device for the path that was deleted in step 2.

Since the host has not been restarted, the SCSI device for the path that was deleted in step 2 was not deleted. Execute the command for deleting the device shown in Table 4.1 to manually delete the SCSI device.

5. Restart the host on which the HDLM is installed.

   If you performed step 4, the host does not have to be restarted. Go to step 6.

6. Execute the HDLM-configuration definition utility (`dlmcfgmgr -v`) or the `view` operation of the `dlnkmgr` command (`dlnkmgr view -drv`) (optional).

   Check if the SCSI device name of the HDLM device whose path was deleted in the above step 2 is represented with a hyphen (-).

   The following are examples of executing the command or utility:

   When executing the `dlmcfgmgr` utility with the `-v` option:

```
# /sbin/dlmcfgmgr -v
```

   When executing the `view` operation with the `-drv` option:

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

   Perform this operation if necessary. If it is not necessary to check this, go to step 7.

   If the SCSI device name is not represented with a hyphen (-), to change the SCSI device name to a hyphen, wait until the path health checking has been executed or execute the following `dlnkmgr` command with the *AutoPATH_ID* of the HDLM device.

   Example:

```
# /opt/DynamicLinkManager/bin/dlnkmgr online -pathid AutoPATH_ID
```

7. Execute the `dlmcfgmgr` utility (`dlmcfgmgr -u` *logical_device_file_name_of_HDLM_device*) on the relevant HDLM device whose path was deleted in the above step 2 to unregister the HDLM device.

   Path information on the relevant HDLM device is deleted.

   The following is an example of executing the `dlmcfgmgr` utility:

```
# /sbin/dlmcfgmgr -u name-of-logical-device-file-for-HDLM-device
```

8. Refer to the `/var/opt/DynamicLinkManager/log/dlmcfgmgr1.log` or execute the `dlmcfgmgr` (`dlmcfgmgr -v`) to check if the path mapped with the relevant HDLM device is deleted.

   The following is an example of executing the `dlmcfgmgr` utility:

```
# /sbin/dlmcfgmgr -v
```

#### 4.4.3.6 Changing a Device Managed by HDLM into One Not Managed by HDLM:

The following operation can be done without restarting the host.

To change a device managed by HDLM into one not managed by HDLM:

1. Stop all accesses to the relevant LU.

2. Execute the HDLM-configuration definition utility (`dlmcfgmgr -v`) to check if the device you want to exclude from HDLM management is being managed by HDLM.

   The following is an example of executing the `dlmcfgmgr` utility:

```
# /sbin/dlmcfgmgr -v
```

If the Management status of the relevant HDLM device is represented as `configured` in the execution result, this means that the device is managed by HDLM.

3. Execute the `dlmcfgmgr` utility (`dlmcfgmgr -o` *logical_device_file_name_of_HDLM_device*) on the relevant device you want to exclude from HDLM management.

   The specified device is excluded from HDLM management.

   The following is an example of executing the `dlmcfgmgr` utility:

```
# /sbin/dlmcfgmgr -o name-of-logical-device-file-for-HDLM-device
```

4. Refer to `/var/opt/DynamicLinkManager/log/dlmcfgmgr1.log` or execute the `dlmcfgmgr` utility (`dlmcfgmgr -v`) or the `view` operation (`dlnkmgr view -path`) of the `dlnkmgr` command to make sure that the relevant device is no longer managed by HDLM.

   The following is an example of executing the `dlmcfgmgr` utility:

```
# /sbin/dlmcfgmgr -v
```

### 4.4.3.7 Changing a Device Not Managed by HDLM Into One Managed by HDLM

The following operation can be done without restarting the host.

To change a device not managed by HDLM into one managed by HDLM:

1. Stop all accesses to the relevant LU.
2. Execute the HDLM-configuration definition utility (`dlmcfgmgr -v`) to check if the device you want HDLM to manage is not being managed.

   If the Management status of the relevant HDLM device is represented as `unconfigured` in the execution result, the device is not being managed by HDLM. Perform this operation if necessary. If it is not necessary to check this, go to step 3.

3. Execute the `dlmcfgmgr` utility (`dlmcfgmgr -i` *logical_device_file_name_of_HDLM_device*) on the relevant device you want to HDLM to manage.

   The specified device is managed by HDLM.

   The following is an example of executing the `dlmcfgmgr` utility:

```
# /sbin/dlmcfgmgr -i name-of-logical-device-file-for-HDLM-device
```

4. Refer to `/var/opt/DynamicLinkManager/log/dlmcfgmgr1.log` or execute the `dlmcfgmgr` utility (`dlmcfgmgr -v`) or the view operation (`dlnkmgr view -path`) of the HDLM command to make sure that the relevant device is now managed by HDLM.

   The following is an example of executing the `dlmcfgmgr` utility:

```
# /sbin/dlmcfgmgr -v
```

### 4.4.3.8 Restoring the Path of an HDLM Device Started in Disconnected Status (When a Restart Is Required)

This subsection describes how to restore the path to an HDLM device that was in disconnected status (the cable was pulled out or broken) when the host first started. In this method, the host is restarted. Note that this operation is applicable to an HDLM device managed by HDLM only. Also note that the target path of restoration can be a single, multiple, or all paths to the relevant HDLM device.

To restore the path of an HDLM device started in disconnected status (this procedure restarts the host):

1. Execute the HDLM-configuration definition utility (`dlmcfgmgr -v`) or the `view` operation (`dlnkmgr view -drv`) of the `dlnkmgr` command to check if the host was started while the path to the relevant HDLM device was in disconnected status.

   The following are examples of executing the command or utility:

   When executing the `dlmcfgmgr` utility with the `-v` option:

   ```
   # /sbin/dlmcfgmgr -v
   ```

   When executing the `view` operation with the `-drv` option:

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

   If the SCSI device name mapped with the path to the relevant HDLM device is represented with a hyphen (-) in the execution result, this means that the path was in disconnected status.

   Perform this operation if necessary. If it is not necessary to check this, go to step 2.

2. Connect the cable.

   If the cable is broken, replace it.

3. Restart the host on which the HDLM is installed.

   The SCSI device is recreated, the path to the HDLM device, which was determined in step 1 to be in disconnected status when the host was started, is restored automatically, and the status changes to online.

4. Refer to `/var/opt/DynamicLinkManager/log/dlmcfgmgr1.log` or execute the `dlmcfgmgr` utility (`dlmcfgmgr -v`) to check if the path to the relevant HDLM device is restored.

   The following is an example of executing the `dlmcfgmgr` utility:

   ```
   # /sbin/dlmcfgmgr -v
   ```

### 4.4.3.9 Restoring the Path of an HDLM Device Started in Disconnected Status (When a Restart Is Not Required)

This subsection describes how to restore the path to an HDLM device that was in disconnected status (the cable was pulled out or broken) when the host first started. In this method, the host is not restarted. Note that this operation is applicable to the HDLM device managed by HDLM only. Also note that the target path of restoration can be a single, multiple, or all path to the HDLM device.

To restore the path of an HDLM device started in disconnected status (this procedure does not restart the host):

1. Execute the HDLM-configuration definition utility (`dlmcfgmgr -v`) or the `view` operation (`dlnkmgr view -drv`) of the HDLM command to check if the host was started while the path to the relevant HDLM device was in disconnected status.

   The following are examples of executing the command or utility:

   When executing the `dlmcfgmgr` utility with the `-v` option:

   ```
   # /sbin/dlmcfgmgr -v
   ```

   When executing the `view` operation with the `-drv` option:

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

   If the SCSI device name mapped with the path to the relevant HDLM device is represented with a hyphen (-), in the execution result, this means that the path was in disconnected status.

   Perform this operation if necessary. If it is not necessary to check this, go to step 2.

2. Connect the cable.

   If the cable is broken, replace it.

3. Create the SCSI device for the path to be restored.

   Since you restore the path without restarting, the path for the SCSI device which was in disconnected status during bootup will not be created and you must execute the command for adding a device, shown in Table 4.2, to manually create the SCSI device.

**Table 4.2 Add SCSI Devices**

| Purpose | Command |
|---|---|
| Add a device | ▪ When using a QLogic HBA driver on SUSE LINUX Enterprise Server 9, Red Hat Enterprise Linux AS4, or Red Hat Enterprise Linux ES4, execute the commands in the following order:<br>`# echo "scsi-qlascan" > /proc/scsi/`*driver-name*`/`*adapter-id*<br>`# cat /proc/scsi/`*driver-name*`/`*adapter-id*<br>View the `Id:Lun` part, which is output under `SCSI LUN Information:`, of the output command result, and then execute the following command based on the target ID of the added device:<br>`# echo "scsi add-single-device `*a b c d*`" > /proc/scsi/scsi`<br>▪ When using a QLogic HBA driver on SUSE LINUX Enterprise Server 10, execute the commands in the following order:<br>`# echo "1" > /sys/class/fc_host/host`*a*`/issue_lip`<br>`# echo "- - -" > /sys/class/scsi_host/host`*a*`/scan`<br>*a : host port number*<br>▪ When using an Emulex HBA driver, execute the command as follows:<br>`# echo "scsi add-single-device `*a b c d*`" > /proc/scsi/scsi` |
| | *a*: Host ID (Port number of the host)<br>*b*: Channel number (Bus number)<br>*c*: Target ID (Target ID)<br>*d*: Lun (LU number of the host)<br>*driver-name*: `qla2200` or `qla2300` (driver name in use)<br>*adapter-id*: host port number |

```
Example 1 (when using a QLogic HBA driver)
  # echo "scsi-qlascan" > /proc/scsi/qla2200/1
Example 2 (to add a device)
  # echo "scsi add-single-device 1 0 1 1" > /proc/scsi/scsi
```

4. Wait until the automatic failback function starts, or execute the `dlnkmgr` command (`dlnkmgr online -pathid` *AutoPATH_ID*).

   The relevant path is placed online.

   The following is an example of executing the HDLM command:

   ```
   # /opt/DynamicLinkManager/bin/dlnkmgr online -pathid AutoPATH_ID
   ```

5. Refer to `/var/opt/DynamicLinkManager/log/dlmcfgmgr1.log` or execute the `dlmcfgmgr` utility (`dlmcfgmgr -v`) to check if the path to the relevant HDLM device is restored.

   The following is an example of executing the `dlmcfgmgr` utility:

   ```
   # /sbin/dlmcfgmgr -v
   ```

## 4.4.4 Changing the Configuration of an Environment Where HDLM and LifeKeeper Are Linked

This section describes the procedures for changing the configuration of an environment where HDLM and LifeKeeper are linked.

After changing the configuration, you may need to restart the system.

### 4.4.4.1 Adding a New LU

To add a new LU:

1. Use the management program of the storage subsystem to allocate an LU, which has not yet been allocated to the relevant host, to a port of the storage subsystem that has a path to the relevant host, and then configure the path.

   You can add multiple paths as long as the number of paths does not exceed the maximum.

   For more information on configuring paths, see the manual provided with the storage subsystem.

2. Restart the active and standby nodes that share the added LU.

   When the host is started, an HDLM device for the newly allocated LU is automatically added.

3. Refer to `/var/opt/DynamicLinkManager/log/dlmcfgmgr1.log` or execute the `dlmcfgmgr` utility (`dlmcfgmgr -v`) to make sure that the HDLM device was created for the added LU.

   The following is an example of executing the `dlmcfgmgr` utility:

   ```
   # /sbin/dlmcfgmgr -v
   ```

4. Register the created HDLM device as a LifeKeeper resource.

   For details on registration, see the LifeKeeper documentation.

#### 4.4.4.2 Adding a Path to an LU (When a Restart Is Required)

To add a path to an LU with restart:

1. Stop all accesses to the LU to which you want to add a path.
2. Stop LifeKeeper on the active and standby nodes.
3. Increase the number of logical paths to the active and standby nodes for the LU that the relevant HDLM device uses.

   You can add multiple paths.
4. Restart the active and standby nodes.

   When the host is started, `dlmcfgmgr -r` is automatically executed and a path to an existing HDLM device is added.

   Even if the host is restarted, HDLM does not recognize paths to an HDLM device that have been excluded from HDLM management by executing `dlmcfgmgr -o`. To include paths to such HDLM devices in HDLM management again, execute `dlmcfgmgr -i` and then restart the host.
5. Refer to `/var/opt/DynamicLinkManager/log/dlmcfgmgr1.log` or execute the `dlmcfgmgr` utility (`dlmcfgmgr -v`) to make sure that a path to the HDLM device was created for the added LU.

   The following is an example of executing the `dlmcfgmgr` utility:

```
# /sbin/dlmcfgmgr -v
```

#### 4.4.4.3 Adding Paths to An LU (When a Restart Is Not Required)

To add a path to an LU without restarting the host:

1. Stop all accesses to the LU to which you want to add a path.
2. Delete the HDLM device to which you want to add a path from the LifeKeeper resources.
3. Execute the following command to unmount the HDLM device:

```
# umount mount-point-name
```

4. Increase the number of logical paths to the active and standby nodes for the LU that the relevant HDLM device uses.

   You can add multiple paths.
5. Since a SCSI device file for the added path has not been created, create this SCSI device file on the active and standby nodes according to Table 4.1.
6. Execute the `dlmcfgmgr -r` command to register and add the path to the existing HDLM device.
7. Refer to `/var/opt/DynamicLinkManager/log/dlmcfgmgr1.log` or execute the `dlmcfgmgr` utility (`dlmcfgmgr -v`) to make sure that a path to the HDLM device was created for the added LU.

   The following is an example of executing the `dlmcfgmgr` utility:

```
# /sbin/dlmcfgmgr -v
```

8. Execute the following command to mount the HDLM device on its original mount point:

```
# mount mount-point-name
```

9. Register the HDLM device to which you added the path to the LifeKeeper resources.

#### 4.4.4.4 Deleting an LU

To delete an LU:

1. Stop all accesses to the LU from which you want to delete paths.
2. Delete the HDLM device for the relevant LU from the LifeKeeper resources.

   For details on how to delete the HDLM device, see the LifeKeeper documentation.
3. Execute the following command to unmount the HDLM device.

```
# umount mount-point-name
```

   If the `/etc/fstab` file contain no entries, go to step 5.
4. If /etc/fstab file contains entries, delete them.

   Execute the following command to check `/etc/fstab` file:

```
# cat /etc/fstab
LABEL=/                 /                 ext3    defaults   1 1
LABEL=/boot            /boot            ext3    defaults   1 2
...
/dev/HDLM-device-name    mount-point-name    ext2    rw          0 2
```

   If an unnecessary mount point name is defined, use an editor to comment out or delete the line containing the mount point name.
5. Cancel allocation of the relevant LU from the ports of the storage subsystem that have paths (physical paths) to the relevant host.
6. Restart the active and standby nodes.
7. Execute the `dlmcfgmgr -v` command or `dlnkmgr view -drv` command to make sure that the SCSI device name for the relevant HDLM device is represented with a hyphen (-).

   The following is an example of executing the `dlmcfgmgr -v` command:

```
# /sbin/dlmcfgmgr -v
HDevName        Management   Device  Host  Channel  Target   Lun
/dev/sddlmaa    configured    -       0       0        0       0
                                  -       0       0       1       0
                                  -       1       0       0       0
                                  -       1       0       1       0
```

8. Execute the `dlmcfgmgr` utility (`dlmcfgmgr -u` *name-of-logical-device-file-for-HDLM-device*) to unregister (delete) the HDLM device.

```
# /sbin/dlmcfgmgr -u /dev/sddlmaa
```

   Once the command is executed, the target path will not be displayed even if you execute the `dlmcfgmgr -v` command or `dlnkmgr view -drv` command.

9. Refer to `/var/opt/DynamicLinkManager/log/dlmcfgmgr1.log` or execute the `dlmcfgmgr` utility (`dlmcfgmgr -v`) to make sure that the HDLM device was deleted.

   The following is an example of executing the `dlmcfgmgr` utility:

```
# /sbin/dlmcfgmgr -v
```

### 4.4.4.5 Deleting a Path from an LU (When a Restart Is Required)

To delete a path from an LU with restart:

1. Stop all accesses to the LU for which you want to delete a path.
2. Reduce the number of logical paths that the LU uses.
3. Restart the active and standby nodes.
4. Execute the `dlmcfgmgr -v` command or `dlnkmgr view -drv` command to make sure that the SCSI device name for the relevant path is represented with a hyphen (-).

   The following is an example of executing the `dlmcfgmgr -v` command:

```
# /sbin/dlmcfgmgr -v
HDevName        Management   Device  Host  Channel  Target   Lun
/dev/sddlmaa    configured  /dev/sda  0       0       0       0
                                     -       0       0       1       0
                                /dev/sdag 1       0       0       0
                                     -       1       0       1       0
```

5. Execute the `dlmcfgmgr` utility (`dlmcfgmgr -u HDLM` *name-of-logical-device-file-for-HDLM-device*) to unregister (delete) the path.

```
# /sbin/dlmcfgmgr -u /dev/sddlmaa
```

   Once the command is executed, the target path will not be displayed even if you execute the `dlmcfgmgr -v` command or `dlnkmgr view -drv` command.

6. Refer to `/var/opt/DynamicLinkManager/log/dlmcfgmgr1.log` or execute the `dlmcfgmgr` utility (`dlmcfgmgr -v`) to make sure that the path was deleted.

   The following is an example of executing the `dlmcfgmgr -v` command:

```
# /sbin/dlmcfgmgr -v
HDevName        Management   Device  Host  Channel  Target   Lun
/dev/sddlmaa    configured  /dev/sda  0       0       0       0
                                /dev/sdag 1       0       0       0
```

### 4.4.4.6 Deleting a Path from an LU (When a Restart Is Not Required)

To delete a path from an LU without restarting the host:

1. Stop all accesses to the LU for which you want to delete a path.
2. Delete the HDLM device for the relevant LU from the LifeKeeper resources.
3. Unmount the HDLM device.

```
# umount mount-point-name
```

4. Reduce the number of logical paths that the LU uses.

5. Delete the SCSI device file for the path according to the method shown in Table 4.1.

6. Execute the `dlmcfgmgr -r` command to update information about the path in the existing HDLM device.

7. Execute the `dlmcfgmgr -v` command or `dlnkmgr view -drv` command to make sure that the SCSI device name for the relevant path is represented with a hyphen (-).

   The following is an example of executing the `dlmcfgmgr -v` command:

```
# /sbin/dlmcfgmgr -v
HDevName        Management   Device  Host  Channel  Target   Lun
/dev/sddlmaa    configured  /dev/sda  0       0       0       0
                                    -       0       0       1       0
                                /dev/sdag 1       0       0       0
                                    -       1       0       1       0
```

   If the SCSI device name is not represented with a hyphen (-), do either of the following:

   – Wait until the path health check is executed.

   – Execute the following HDLM command with the *AutoPATH_ID* of the HDLM device:

```
dlnkmgr online -pathid AutoPATH_ID
```

8. Execute the `dlmcfgmgr` utility (`dlmcfgmgr -u` *name-of-logical-device-file-for-HDLM-device*) to unregister (delete) the path.

```
# /sbin/dlmcfgmgr -u /dev/sddlmaa
```

   Once the command is executed, the target path will not be displayed even if you execute the `dlmcfgmgr -v` command or `dlnkmgr view -drv` command.

9. Refer to `/var/opt/DynamicLinkManager/log/dlmcfgmgr1.log` or execute the `dlmcfgmgr` utility (`dlmcfgmgr -v`) to make sure that the path was deleted.

   The following is an example of executing the `dlmcfgmgr -v` command:

```
# /sbin/dlmcfgmgr -v
HDevName        Management   Device  Host  Channel  Target   Lun
/dev/sddlmaa    configured  /dev/sda  0       0       0       0
                                /dev/sdag 1       0       0       0
```

10. Mount the HDLM device on its original mount point.

```
# mount mount-point-name
```

11. Register the HDLM device for which you deleted the path in the LifeKeeper resources.

#### 4.4.4.7 Adding, Deleting, and Changing Partition Information

To add, delete, or change path information:

1. Stop all accesses to the LU for which you want to add, delete, or change partition information.

2. Delete the HDLM device for the relevant LU from the LifeKeeper resources.

3. Unmount the HDLM device.

```
# umount mount-point-name
```

4. Add, delete, or change partition information for the LU.
5. Restart the active and standby nodes that share the LU on which processing was performed.
6. Mount the HDLM device on its original mount point.

```
# mount mount-point-name
```

7. Register the HDLM device in the LifeKeeper resources.

#### 4.4.4.8 Changing a Device Managed by HDLM into One Not Managed by HDLM:

To change a device managed by HDLM into one not managed by HDLM:

1. Stop all accesses to the relevant LU.
2. Delete the HDLM device for the relevant LU from the LifeKeeper resources.
3. Execute the following command to unmount the HDLM device.

```
# umount mount-point-name
```

   If the `/etc/fstab` file contain no entries, go to step 5.

4. If /etc/fstab file contains entries, delete them.

   Execute the following command to check `/etc/fstab` file:

```
# cat /etc/fstab
LABEL=/                    /              ext3    defaults        1 1
LABEL=/boot               /boot          ext3    defaults        1 2
...
/dev/HDLM-device-name  mount-point-name  ext2    rw              0 2
```

   If an unnecessary mount point name is defined, use an editor to comment out or delete the line containing the mount point name.

5. Execute the `dlmcfgmgr -v` command to check whether the relevant path is still being managed.

   The following is an example of executing the `dlmcfgmgr -v` command:

```
# /sbin/dlmcfgmgr -v
HDevName        Management   Device  Host  Channel  Target   Lun
/dev/sddlmaa    configured  /dev/sda  0       0       0       0
                                /dev/sdi  0       0       1       0
                                /dev/sdag 1       0       0       0
                                /dev/sdao 1       0       1       0
```

   If `configured` is displayed for `Management`, the path is still being managed.

6. Execute the `dlmcfgmgr` utility (`dlmcfgmgr -o` *name-of-logical-device-file-for-HDLM-device*) to exclude the HDLM device as a management target.

```
# /sbin/dlmcfgmgr -o /dev/sddlmaa
```

7. Refer to `/var/opt/DynamicLinkManager/log/dlmcfgmgr1.log` or execute the `dlmcfgmgr` utility (`dlmcfgmgr -v`) or the `dlnkmgr view -path` command to make sure that the path is not being managed.

The following is an example of executing the `dlmcfgmgr` utility:

```
# /sbin/dlmcfgmgr -v
HDevName        Management   Device  Host  Channel  Target   Lun
/dev/sddlmaa  unconfigured  /dev/sda  0       0       0       0
                                 /dev/sdi  0       0       1       0
                                 /dev/sdag 1       0       0       0
                                 /dev/sdao 1       0       1       0
```

If `unconfigured` is displayed for `Management`, the path has been removed from management.

#### 4.4.4.9 Changing a Device Not Managed by HDLM Into One Managed by HDLM

To change a device not managed by HDLM into one managed by HDLM:

1. Execute the `dlmcfgmgr -v` command to check whether the relevant path is not being managed.

```
# /sbin/dlmcfgmgr -v
HDevName        Management   Device  Host  Channel  Target   Lun
/dev/sddlmaa  unconfigured  /dev/sda  0       0       0       0
                                 /dev/sdi  0       0       1       0
                                 /dev/sdag 1       0       0       0
                                 /dev/sdao 1       0       1       0
```

   If `unconfigured` is displayed for `Management`, the path is not being managed.

2. Execute the `dlmcfgmgr` utility (`dlmcfgmgr -i` *name-of-logical-device-file-for-HDLM-device*) to include the HDLM device in the management targets.

```
# /sbin/dlmcfgmgr -i /dev/sddlmaa
```

3. Refer to `/var/opt/DynamicLinkManager/log/dlmcfgmgr1.log` or execute the `dlmcfgmgr` utility (`dlmcfgmgr -v`) or the `dlnkmgr view -path` command to make sure that the path is being managed.

   The following is an example of executing the `dlmcfgmgr` utility:

```
# /sbin/dlmcfgmgr -v
HDevName        Management   Device  Host  Channel  Target   Lun
/dev/sddlmaa    configured  /dev/sda  0       0       0       0
                                /dev/sdi  0       0       1       0
                                /dev/sdag 1       0       0       0
                                /dev/sdao 1       0       1       0
```

   If `configured` is displayed for `Management`, the path is now being managed.

4. Register the management-target HDLM device as a LifeKeeper resource.

   For details about how to register a resource, see the LifeKeeper documentation.

#### 4.4.4.10 Restoring the Path to an HDLM Device Started in a Disconnected Status

To restore the path to an HDLM device that was started in a disconnected status:

1. Execute the `dlmcfgmgr -v` command or `dlnkmgr view -drv` command to check whether the host was started while the relevant path was disconnected.

If the SCSI device for the path to the relevant HDLM device is represented with a hyphen (-), this means the path was disconnected when the host was started.

The following is an example of executing the `dlmcfgmgr` utility:

```
# /sbin/dlmcfgmgr -v
HDevName        Management   Device  Host  Channel  Target   Lun
/dev/sddlmaa    configured  /dev/sda  0       0       0       0
                                      -     0       0       1       0
                                /dev/sdag 1       0       0       0
                                      -     1       0       1       0
```

2. Connect the disconnected cable.
3. Restart the host.

   The SCSI device is re-created, and the path is recognized automatically when the host is started.
4. Refer to `/var/opt/DynamicLinkManager/log/dlmcfgmgr1.log` or execute the `dlmcfgmgr` utility (`dlmcfgmgr -v`) or the `dlnkmgr view -path` command to make sure that an HDLM device was created for the added LU.

   The following is an example of executing the `dlmcfgmgr` utility:

```
# /sbin/dlmcfgmgr -v
HDevName        Management   Device  Host  Channel  Target   Lun
/dev/sddlmaa    configured  /dev/sda  0       0       0       0
                                /dev/sdi  0       0       1       0
                                /dev/sdag 1       0       0       0
                                /dev/sdo  1       0       1       0
```

#### 4.4.4.11 Restoring the Path of an HDLM Device When a Standby Node Containing a Disconnected or Offline(C) Path Is Changed to Active (When Using a Block-Type Device File)

To restore the path of an HDLM device:

1. Delete the HDLM device connected to the relevant LU from the LifeKeeper resources.
2. Unmount the HDLM device.

```
# umount mount-point-name
```

3. Connect the disconnected cable.
4. If there is a path placed in the `Offline(C)` status, return it to the `Online` status.

```
# /opt/DynamicLinkManager/bin/dlnkmgr online -s
```

5. Mount the HDLM device on its original mount point.

```
# mount mount-point-name
```

6. Register the HDLM device connected in the relevant LU to the LifeKeeper resources.

Once you change the node to the standby status, you can use HDLM automatic failback or a command to put the paths online.

#### 4.4.4.12 Restoring the Path of an HDLM Device When a Standby Node Containing a Disconnected or Offline(C) Path Is Changed to Active (When Using a Character-Type Device File)

1. Delete the HDLM device connected to the relevant LU from the LifeKeeper resources.
2. Connect the disconnected cable.
3. If there is a path placed in the `Offline(C)` status, return it to the `Online` status.
```
# /opt/DynamicLinkManager/bin/dlnkmgr online -s
```
4. Register the HDLM device for the relevant path in the LifeKeeper resources.

   For details on how to register, see the LifeKeeper documentation.

Once you change the node to the standby status, you can use HDLM automatic failback or a command to put the path online.

### 4.4.5 About Creating a New HDLM Device

While restarting the host, HDLM may detect a new path. The HDLM device to be allocated depends on whether the SCSI device for the detected path is already registered or a new SCSI device is registered. This section describes a situation where the HDLM device that is to be allocated to the newly detected path is already registered and a situation where a new HDLM device is created.

If the SCSI device of the new path that HDLM detected is already registered, the HDLM device mapped with the path already registered is allocated. For example, if the path between the host and an LU of the storage subsystem already exists and a new path is added, the HDLM device that already exists is allocated to the path.

If the SCSI device of the new path that HDLM detected is not yet registered, a new HDLM device is allocated to the path. For example, if you define a new path from the host to an LU of the storage subsystem, a new HDLM device is allocated to the path. The lowest available alphabetical letter is allocated as the logical device file name of the HDLM device.

Examples of the logical device file name of the HDLM device when a new HDLM device is allocated are shown in Table 4.3.

**Table 4.3 Example: Allocation of New HDLM Devices**

| Allocation before restarting the host | Allocation after restarting the host |
|---|---|
| none | /dev/sddlmaa# |
| /dev/sddlmaa | /dev/sddlmaa<br>/dev/sddlmab# |
| /dev/sddlmaa<br>/dev/sddlmab | /dev/sddlmaa<br>/dev/sddlmab<br>/dev/sddlmac# |

| Allocation before restarting the host | Allocation after restarting the host |
|---|---|
| /dev/sddlmaa<br>/dev/sddlmac | /dev/sddlmaa<br>/dev/sddlmab#<br>/dev/sddlmac |
| /dev/sddlmab<br>/dev/sddlmac | /dev/sddlmaa#<br>/dev/sddlmab<br>/dev/sddlmac |

#

The newly allocated logical device file name of the HDLM device

If a new LU was recognized when the number of HDLM devices created reached the upper limit, HDLM displays the message KAPL10357-E and does not add an HDLM device. In this case, you must execute the HDLM-configuration definition utility (dlmcfgmgr -u) to delete the HDLM device(s) not in use (or, represented with a hyphen (-)) to release the available name(s) to create an HDLM device file for the new LU.

Even if you add a path to an HDLM device that is not managed by HDLM and restart the host, the HDLM device will not be managed by HDLM.

An AutoPATH_ID is allocated sequentially as the HDLM recognizes a SCSI device during start up of the host. Therefore, an AutoPATH_ID for the same path may change whenever the host starts.

If a different LU is allocated to the same path name, a new HDLM device will be allocated to the newly allocated LU after restarting the host restarts. In such case, the previous HDLM device is unregistered automatically.

# Chapter 5 Troubleshooting

This chapter describes how to check HDLM error information, and then how to take action if an error has occurred in HDLM. The actions you should take are described separately for path errors, for HDLM program errors, and for other types of errors.

## 5.1 Information Collected by the DLMgetras Utility for Collecting HDLM Error Information

Immediately after an error occurs, execute the error information collection utility `DLMgetras`, since restarting the machine may delete information collected by `DLMgetras`.

For details about the `DLMgetras` utility and the error information it collects, see section 7.2.

## 5.2 Checking Error 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 `syslogd settings` 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 section 6.7.

The following is an example of a message:

KAPL08022-E Error in path occurred. 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*

The elements of the message are explained below.

`ErrorCode`

The error number generated when Linux 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 are re-assigned to each path of the LU when you execute the `dlmcfgmgr` utility.

The path ID is also the same as `PathID` displayed by the `dlnkmgr` command's `view` operation. For details on this `view` operation, see section 6.7.

`PathName`

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 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 `PathName` displayed by the command's `view` operation. For details on the path name, see section 6.7.

`DNum`

A Dev number, which is equivalent to a partition number in Linux.

A device number beginning from `0` is assigned to the device in the LU.

In Linux, this value is fixed to `0`.

This is the same as the `DNum` that is displayed by the `view` operation. For details on the `view` operation, see section 6.7.

`HDevName`

The name of the host device.

The logical device file name of the HDLM device to which access is made by way of the path, minus the partition number, is displayed in the format `sddlm`[*a-p*][*a-p*] (for example, `sddlmaa`), where *a-p* indicates a single letter from `a` to `p`.

This is the same as the `HDevName` that is displayed by the `view` operation. For details on the `view` operation, see section 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.

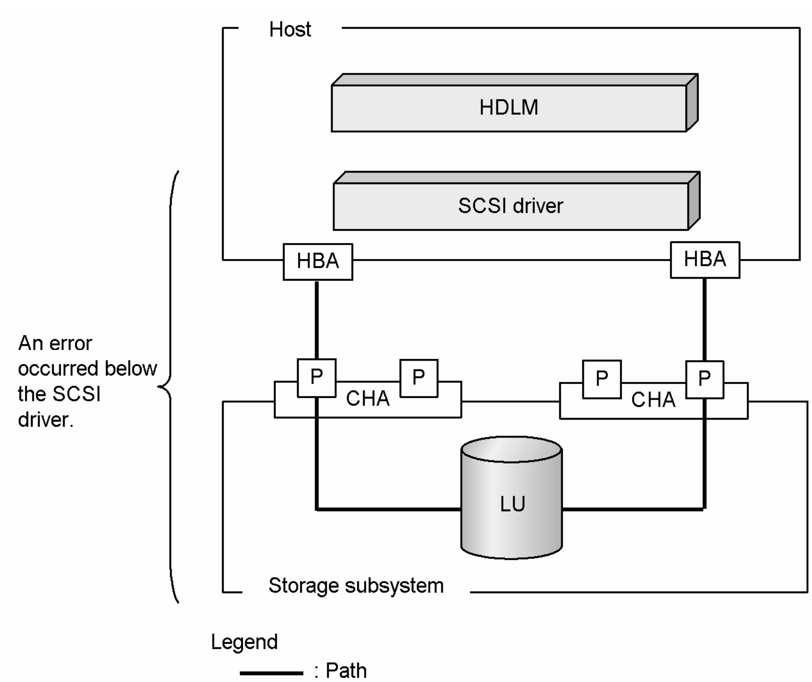


**Figure 5.1 Error Location When the KAPL08022-E Message Is Output**

Figure 5.2 shows the troubleshooting procedure when the KAPL08022-E message is output.

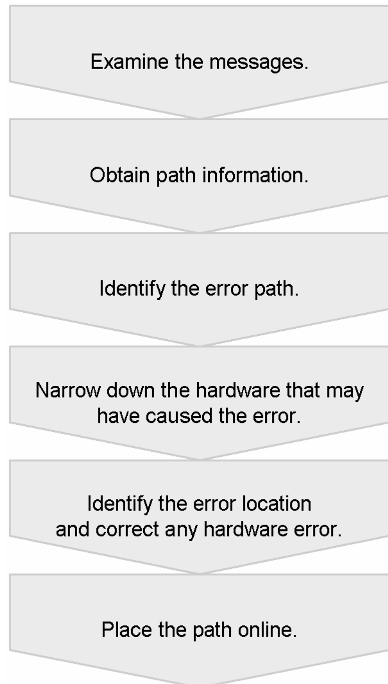


**Figure 5.2 Troubleshooting Procedure When a Path Error Occurs**

The following shows the procedure for using the HDLM command (`dlnkmgr`) 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 that message to check the path in which the error occurs. For details on the message, see section 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 value displayed in the **Subsystem**, **LUN**, and **CHA** 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 **Subsystem**, **LUN**, and **CHA**, use the information provided by the storage-subsystem management program.

### 5.3.5 Identify the Error Location and Correct Any Hardware Error

Use the Linux 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 Path Management window to place online the path that went offline due to the error. For details on the online operation, see section 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. Type n to cancel the processing.

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

## 5.4 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.

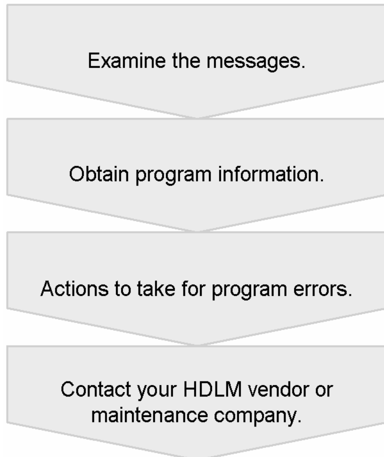


**Figure 5.3 Troubleshooting Procedure When a Program Error Occurs**

The following shows the procedure for using the HDLM command (`dlnkmgr`) to handle a path error.

### 5.4.1 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. Messages with error level `E` (Error) or higher require corrective action.

### 5.4.2 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 to collect the error information. For details about this utility and the information it collects, see section 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.

### 5.4.3 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 section 6.7

For example, execute the following command:

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

If the KAPL01012-E message appears as a result of executing the command:

The KAPL01012-E message is as follows:

```
KAPL01012-E Could not connect the HDLM manager. Operation name = view
```

In this situation, start the HDLM manager. For details about how to do this, see section 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* are character strings.

In this situation, restart the host.

If the same error re-occurs after you take action, go to the following subsection 5.4.4.

### 5.4.4 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 utility for collecting HDLM error information (`DLMgetras`).

## 5.5 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, and then report the collected information to the HDLM vendor or maintenance company if there is a maintenance contract for HDLM. For details about the `DLMgetras` utility and the information it collects, see section 7.2.

# Chapter 6 Command Reference

This chapter describes the HDLM command (`dlnkmgr`) and its operations.

## 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 `dlnkmgr` 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 section 6.2. |
| `help` | Displays the format of the operation used for HDLM. For details, see section 6.3. |
| `offline` | Places offline an online path or paths. For details, see section 6.4. |
| `online` | Places online an offline path or paths. For details, see section 6.5. |
| `set` | Sets the HDLM operating environment. For details, see section 6.6. |
| `view` | Displays HDLM program information, path information, LU information, and corresponding information about each HDLM device, SCSI device, and LDEV. For details, see section 6.7. |

***Notes:***

- 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 quotes (`"`).

## 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 HDLM 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 HDLM 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 one of the following operation names:

- `clear`
- `help`
- `online`
- `offline`
- `set`
- `view`

If you do not specify any operation names, the `help` operation displays the names of all operations available for the HDLM command.

### Examples

Example 1

The following example shows how to display the names of all the operations available in the HDLM command.

```
# /opt/DynamicLinkManager/bin/dlnkmgr help
dlnkmgr:
  Format
    dlnkmgr { clear | help | offline | online | set | view }
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 [-s]
    dlnkmgr online [-path] -cha -pathid AutoPATH_ID [-s]
    dlnkmgr online [-path] [-pathid AutoPATH_ID] [-s]
  Valid value
```

```
    AutoPATH_ID    { 000000 - 999999 }(Decimal)
offline:
  Format
    dlnkmgr offline [-path] -hba HBAPortNumber.BusNumber [-s]
    dlnkmgr offline [-path] -cha -pathid AutoPATH_ID [-s]
    dlnkmgr offline [-path] -pathid AutoPATH_ID [-s]
  Valid value
    AutoPATH_ID    { 000000 - 999999 }(Decimal)
help:
  Format
    dlnkmgr help { clear | offline | online | set | view }
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 by the `help` operation

```
# /opt/DymanicLinkManager/bin/dlnkmgr help help
help:
  Format
    dlnkmgr help { clear | offline | online | set | view }
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 a 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 about the `view` operation, see section 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*

`|-cha -pathid` *AutoPATH_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`, `-cha`, 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 about the `view` operation, see section 6.7.

Example

The following example shows how to place offline all paths connected to the HBA port whose host port number is `0010` 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 offline -hba 0010.0000
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]:
KAPL01061-I 3 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 about the `view` operation, see section 6.7. Leading zeros can be omitted (`000001` and `1` indicate the same *AutoPATH_ID*); however, 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 use 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.

When the confirmation message is displayed, the user enters `y` to continue, or `n` to cancel the operation.

```
# /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]:
KAPL01061-I 2 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 about the `view` operation, see section 6.7. Leading zeros can be omitted (`000001` and `1` indicate the same *AutoPATH_ID*); however, 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 use 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 [-s]
    dlnkmgr offline [-path] -cha -pathid AutoPATH_ID [-s]
    dlnkmgr offline [-path] -pathid AutoPATH_ID [-s]
  Valid value
    AutoPATH_ID     { 000000 - 999999 }(Decimal)
KAPL01001-I The HDLM command completed normally. Operation name = offline, completion
time = yyyy/mm/dd hh:mm:ss
#
```

## Reference

Using 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 about the `view` operation, see section 6.7.

We recommend that you use 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 a 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*

`|-cha -pathid` *AutoPATH_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`, `-cha`, 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 asks whether you would like to continue processing. 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 with 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 about the `view` operation, see section 6.7.

Example

The following example shows how to place online all paths connected to an HBA port whose host port number is `0010` 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 0010.0000
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
#
```

`-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 about the `view` operation, see section 6.7. Leading zeros can be omitted (`000001` and `1` indicate the same *AutoPATH_ID*); however, 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 HDLM-configuration definition utility (`dlmcfgmgr`). Always make sure that you use 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.

When the confirmation message is displayed, the user enters `y` to continue, or `n` to cancel the operation.

```
# /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
#
```

`-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 about the `view` operation, see section 6.7. Leading zeros can be omitted (`000001` and `1` indicate the same *AutoPATH_ID*); however, 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 use 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 [-s]
    dlnkmgr online [-path] -cha -pathid AutoPATH_ID [-s]
    dlnkmgr online [-path] [-pathid AutoPATH_ID] [-s]
  Valid value
    AutoPATH_ID     { 000000 - 999999 }(Decimal)
KAPL01001-I The HDLM command completed normally. Operation name = online, completion
time = yyyy/mm/dd hh:mm:ss
#
```

#### Reference

Using 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 about the `view` operation, see section 6.7.

We recommend that you use 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 defaults and recommended values for each setting. If you change the value of the `set` operation, the new value takes effect immediately.

**Table 6.2 Default values and recommended values**

| Item name | Default value | Recommended setting |
|---|---|---|
| Load balancing | on | on |

| Item name | Default value | Recommended setting |
|---|---|---|
| | Algorithm is round robin | The recommended algorithm depends on the operating environment. |
| Error log collection level | 3: Collect error information for the Information level and higher. | 3: Collect error information for the Information level and 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<br>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 the 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 HDLM manager log (`dlmmgr[1-16].log`) contains log files in which an error log collection level can be set.

Table 6.3 shows the values of the error log collection level. If an error occurs, you may have to set the error log collection level to 1 or higher to collect log information.

**Table 6.3 Values of the Error Log Collection Levels**

| Value | Description |
|---|---|
| 0 | Collects no error log. |
| 1 | Collects error information for the Error or higher level. |

| Value | Description |
|---|---|
| 2 | Collects error information for the Warning or higher level. |
| 3 | Collects error information for the Information or higher level. |
| 4 | Collects error information for the Information (including maintenance information) or higher level. |

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 2,000,000 (in kilobytes) for the size of the error log files. Log files (`dlmmgr`*n*`.log` (*n* indicates a file number from `1` to `16`)). The specified files size is applied to HDLM manager logs.

By specifying both the error log file size and the number of error log files, you can collect up to 32,000,000 kilobytes (approximately 30 GB) of error logs in total.

When the size of all log files in a log file group reaches a specified value, new log data overwrites existing log data (the oldest log file is overwritten first).

`-elfn` *number-of-error-log-files*

Specify the number of error log files (`dlmmgr`*n*`.log` (*n* indicates a file number from `1` to `16`)). Specify a value from `2` to `16`.

By specifying both the error log file size and the number of error log files, you can collect up to 32,000,000 kilobytes (approximately 30 GB) of error logs in total.

`-systflv` *trace-level*

Specify the trace output level.

Log files for which the trace level can be set are `hdlmtr`*n*`.log` (where n is a value 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. |

| Value | Description |
|---|---|
| 4 | Outputs all information. |

The larger the trace level value, the larger the amount of log information that is output. As the amount of log information that is output increases, the amount of time before existing information is overwritten becomes shorter.

`-systfs` *trace-file-size*

Specify a value from 100 to 16,000 for the size of the trace log file (in kilobytes).

By specifying both the trace file size and the number of trace files, you can collect up to 1,024,000 kilobytes of trace files in total.

If you specify a value smaller than the set value, the `KAPL01097-W` message is displayed to confirm execution, and the trace file is temporarily deleted.

Trace files for which the trace file size can be set are `hdlmtr`*n*`.log` (where n is a value from 1 to 64). The length of a trace file is fixed. Therefore, 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. 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 a value from 2 to 64 for the number of trace log files.

By specifying both the trace file size and the number of trace files, you can collect up to 1,024,000 kilobytes of trace data in total.

If you specify a value smaller than the set value, the `KAPL01097-W` message is displayed to confirm execution, and the trace file is temporarily deleted.

Trace files for which the number of files can be set are `hdlmtr`*n*`.log` (where n is a value 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 Thunder 9200, 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 used 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 change 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.

See Table 6.5 for the relationship between automatic failback and intermittent error monitoring.

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 used 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 automatic failback is set to `on`.

The default is `off`. 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 or the number of times that an error is to occur, 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 or the number of times that an error is to occur, the new settings take effect after the next time automatic failback succeeds. Because the errors and elapsed time are not counted or measured while intermittent errors are not monitored, those values do not change.

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 in order for the system to determine that an intermittent error has occurred, the number of errors and the time that has passed since intermittent error monitoring starts are reset to 0. Then, the changed setting takes effect and intermittent error monitoring starts.

Outside the duration of intermittent error monitoring, if you change the number of times that the error is to occur in order for the system to determine that an intermittent error has occurred, from the next time automatic failback finishes normally, the changed values takes effect. 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 that has passed 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 functionality 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 (AFB) and Intermittent Error Monitoring (IEM) and the Executable Operations**

| Setting | | Available 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.<br>▪ The error count, elapsed monitoring time, and information about paths not subject to automatic failback are cleared. |
| | | 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.<br>▪ 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<br>▪ When a path is not being monitored (outside the |

| Setting | | Available operation | Result of operation |
|---|---|---|---|
| | | | 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.<br>▪ The error count, elapsed monitoring time, and information about paths not subject to automatic failback are cleared. |
| | `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

Since 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 stores an HDLM license key.

If a license key file is provided:

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 a license key is provided:

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 on an ongoing basis. |
| Temporary license key# | A 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 | An 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

```
# /opt/DynamicLinkManager/bin/dlnkmgr set -help
set:
  Format
    dlnkmgr set { -lb on [ -lbtype { rr | exrr } ]
                | -lb off
                | -ellv ElogLevel
                | -elfs ElogFileSize
                | -elfn Number-Of-ElogFiles
                | -systflv TraceLevel
                | -systfs TraceFileSize
                | -systfn Number-Of-TraceFiles
                | -pchk on [ -intvl Interval-Time ]
                | -pchk off
                | -afb  on [ -intvl Interval-Time ]
                | -afb  off
                | -iem  on
                        [ -intvl Error-Monitor-Interval ]
                        [ -iemnum Number-Of-Times ]
                | -iem  off
                | -lic
                }
                [-s]
  Valid value
    ElogLevel               { 0 | 1 | 2 | 3 | 4 } (Default Value 3)
    ElogFileSize            { 100 - 2000000 }(KB) (Default Value 9900)
    Number-Of-ElogFiles     { 2 - 16 }(Files)     (Default Value 2)
    TraceLevel              { 0 | 1 | 2 | 3 | 4 } (Default Value 0)
    TraceFileSize           { 100 - 16000 }(KB)   (Default Value 1000)
    Number-Of-TraceFiles    { 2 - 64 }(Files)     (Default Value 4)
    Interval-Time           { 1 - 1440 }(Minute)  (Default Value 30)
      (pchk)
    Interval-Time           { 1 - 1440 }(Minute)  (Default Value 1)
      (afb)
    Error-Monitor-Interval  { 1 - 1440 }(Minute)  (Default Value 30)
    Number-Of-Times         { 1 - 99 }(Times)     (Default Value 3)
KAPL01001-I The HDLM command completed normally. Operation name = set, completion time =
yyyy/mm/dd hh:mm:ss
#
```

## 6.7 view (Displays Information)

The `dlnkmgr` command's `view` operation displays HDLM program information, path information, LU information, and corresponding information about an HDLM device, SCSI 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 path information (by abbreviating the list items)**

```
/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 selecting items to be displayed)

```
/opt/DynamicLinkManager/bin/dlnkmgr view -lu -item

    [ [slpr] [pn] [cp] [clpr] [type] [ic] [ie] [dnu] [iep]

    [hctl]|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 corresponding information about an HDLM device, SCSI device, and LDEV

```
/opt/DynamicLinkManager/bin/dlnkmgr view -drv [-t]
```

### 6.7.1.5 To display the format of the view operation

```
/opt/DynamicLinkManager/bin/dlnkmgr view -help
```

## 6.7.2 Parameters (To display program information)

This section describes the parameters for the view operation, in the following order:

(1) To display program information

(2) To display path information

(3) To display LU information

(4) To display corresponding information about an HDLM device, SCSI device, and LDEV

(5) To display the format of the view operation

### 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 shows the parameters you can specify and the displayed information.

`-t`

Does not display the title for each information item.

**Table 6.7 Function settings**

| Parameter and program information | Item | Description |
|---|---|---|
| `-sfunc`<br>Information about the HDLM function settings | `HDLM Version` | HDLM version number. |
| | `Service Pack Version` | HDLM SP version number. This item is blank if a SP is not installed. |
| | `Load Balance` | Setting for load balancing<br>▪ Setting status<br>`on`: Enabled<br>`off`: Disabled<br>▪ Algorithm<br>When the setting status of load balancing is `on`, the type of algorithm used for load balancing is displayed after the `on`.<br>`rr`: Round robin<br>`extended rr`: Extended round robin |
| | `Support Cluster` | Blank# |
| | `Elog Level` | Error logging level<br>▪ `0`: Collects no error information.<br>▪ `1`: Collects error information at the Error level or higher.<br>▪ `2`: Collects error information at the Warning level or higher.<br>▪ `3`: Collects error information at the Information level or higher.<br>▪ `4`: Collects error information at the Information or higher level (including maintenance information). |
| | `Elog File Size (KB)` | Size of the error log file in kilobytes |
| | `Number Of Elog Files` | Number of error log files |

| Parameter and program information | Item | Description |
|---|---|---|
| | Trace Level | Trace output level<br>▪ 0: Does not output any trace.<br>▪ 1: Only outputs error information.<br>▪ 2: Outputs a summary of program operation.<br>▪ 3: Outputs details of program operation.<br>▪ 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<br>on: Enabled<br>off: Disabled<br>▪ Checking interval<br>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<br>on: Enabled<br>off: Disabled<br>▪ Checking interval<br>When the setting of the automatic failback is on, the checking interval of automatic failback is displayed within the parentheses, ( ), after on. The time is in minutes. |
| | Reservation Status | Blank |
| | Intermittent Error Monitor | ▪ Setting for intermittent error monitoring<br>on: Enabled<br>off: Disabled<br>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.<br>▪ Intermittent error monitoring interval and number of times that the error is to occur<br>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<br>Information about the HDLM manager | HDLM Manager | Status of the HDLM manager<br>Alive: Normal<br>Dead: Stopped |
| | Ver | Version number of the HDLM manager |
| | WakeupTime | Startup time of the HDLM manager |

| Parameter and program information | Item | Description |
|---|---|---|
| -adrv<br>Information about the HDLM alert driver | HDLM Alert Driver | Status of the HDLM alert driver<br>Alive: Normal<br>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. Number of error log files. |
| -pdrv<br>Information about the HDLM driver | HDLM Driver | Status of the HDLM driver<br>Alive: Normal<br>Dead: Stopped |
| | Ver | Version number of the HDLM driver |
| | WakeupTime | Startup time of the HDLM driver |
| -lic<br>Information about the HDLM license | License Type | License type<br>▪ Permanent: permanent license<br>▪ Temporary: temporary license<br>▪ Emergency: emergency license |
| | Expiration | License expiration<br>When using a permanent license: -<br>When using a temporary license or emergency license: The license expiration period is displayed in the 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.<br>▪ When there are 100 days left until the license period (2006/08/21) expires<br>2006/08/21(100days after) |

\#

When you use cluster software, the settings of the cluster function and the kinds of cluster servers are not displayed. However, the cluster 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 Files Size(KB)          : 1000
Number Of Trace Files         : 4
```

```
Path Health Checking           : on(30)
Auto Failback                  : off
Reservation Status             :
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 1000
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, path information and displayed items.

#### 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 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 subsequent sub-parameters (following -path) are:

-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 partition number), /dev/sddlm[*a-p*][*a-p*], for the HDLM device. ([*a-p*]: a single letter from a to p)

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 whose name is `sddlmaa`.

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -path -hdev sddlmaa
Paths:000004 OnlinePaths:000004
PathStatus   IO-Count   IO-Errors
Online         0                0

PathID   PathName                            DskName                    iLU   ChaPort Status   Type    IO-Count  IO-Errors  DNum     HDevName
000000   0000.0000.0000000000000000.0000   HITACHI .DF600F       .0115 0000  0A      Online   Own            0          0        0  sddlmaa
000020   0000.0000.0000000000000001.0000   HITACHI .DF600F       .0115 0000  1A      Online   Non            0          0        0  sddlmaa
000040   0001.0000.0000000000000000.0000   HITACHI .DF600F       .0115 0000  0A      Online   Own            0          0        0  sddlmaa
000041   0001.0000.0000000000000001.0000   HITACHI .DF600F       .0115 0000  1A      Online   Non            0          0        0  sddlmaa

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.

When the value of the `-item` parameter is omitted, only the `PathID` and the `Status` fields are displayed.

Table 6.8 shows the correspondence between the items that can be displayed by the `dlnkmgr view -path -item` command and the parameter values for displaying each item.

**Table 6.8 Correspondence between the items displayed by the dlnkmgr view -path -item command and the values of the -item parameter**

| Items displayed by the dlnkmgr view -path -item command | Values specified after the -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 subsequent sub-parameters (following `-path -item`) are:

`-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 partition number), `/dev/sddlm`[*a-p*][*a-p*], for the HDLM device. ([*a-p*]: a single letter from `a` to `p`)

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.

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -path -item ic -srt lu -stname
Paths:000010  OnlinePaths:000010
PathStatus IO-Count   IO-Errors
Online     500        0

PathID DskName                              iLU              Status     IO-Count
000003 HITACHI .9500V          .0123        0180             Online
400
000009 HITACHI .9500V          .0123        0180             Online
420
000004 HITACHI .9500V          .0123        0181             Online
410
000010 HITACHI .9500V          .0123        0181             Online
399
000005 HITACHI .9500V          .0123        0182             Online
405
000011 HITACHI .9500V          .0123        0182             Online
405
000000 HITACHI .USP_V          .0014050     000050           Online
1005
000006 HITACHI .USP_V          .0014050     000050           Online
897
000001 HITACHI .USP_V          .0014050     000051           Online
0
000007 HITACHI .USP_V          .0014050     000051           Online
0
000002 HITACHI .USP_V          .0014050     000052           Online
0
000008 HITACHI .USP_V          .0014050     000052           Online
0
KAPL01001-I The HDLM command completed normally. Operation name = view, completion
time = yyyy/mm/dd hh:mm:ss
#
```

### To display path information, by abbreviating the list items

`-path -c`

When you specify the `-path` parameter together with the `-c` parameter, the command selects certain items to be displayed, and shortens the contents of each item so that the information about each path fits into a single line on the screen.

The items that are displayed are `PathID`, `DskName`, `iLU`, `CP`, `Status`, and `Type`.

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 ellipses (...).

The subsequent sub-parameters (following `-path -c`) are:

`-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 abbreviate the display of information about the paths, ordered by iLU.

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -path -c -srt lu
Paths:000012 OnlinePaths:000012
PathStatus   IO-Count   IO-Errors
Online       25          0

PathID DskName                        iLU              CP Status     Type
000000 HITACHI .DF600F    .0115            0000             1A Online     Non
000001 HITACHI .DF600F    .0115            0001             1A Online     Own
000002 HITACHI .DF600F    .0115            0002             1A Online     Non
000003 HITACHI .DF600F    .0115            0003             1A Online     Own
000004 HITACHI .DF600F    .0115            0004             1A Online     Non
000005 HITACHI .DF600F    .0115            0005             1A Online     Own
000006 HITACHI .DF600F    .0115            0000             0A Online     Own
000007 HITACHI .DF600F    .0115            0001             0A Online     Non
000008 HITACHI .DF600F    .0115            0002             0A Online     Own
000009 HITACHI .DF600F    .0115            0003             0A Online     Non
000010 HITACHI .DF600F    .0115            0004             0A Online     Own
000011 HITACHI .DF600F    .0115            0005             0A Online     Non
KAPL01001-I The HDLM command completed normally. Operation name = view, completion
time = yyyy/mm/dd hh:mm:ss
#
```

## Items of 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 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 are offline. In this case, you should check the offline 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:<br>▪ Online: All paths are available.<br>▪ Reduced: Some paths are not available.<br>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.<br>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 HDLM-configuration definition utility (dlmcfgmgr). |
| PathName# | - | 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 physical path that will be affected by the change. Path name consists of the following four elements, separated by periods:<br>▪ Host port number (hexadecimal number)<br>▪ Bus number (hexadecimal number)<br>▪ Target ID (hexadecimal number)<br>▪ Host LU number (hexadecimal number)<br>For details about each element of the path name and its representation in Linux, see Table 6.10. |
| DskName# | DskName | Storage subsystem name, which identifies the storage subsystem that is accessed by a path.<br>A storage subsystem name consists of the following three elements, separated by periods:<br>▪ Vendor ID: The name of the storage subsystem vendor (for example, HITACHI).<br>▪ Product ID: Indicates the storage subsystem product ID, emulation type, or model ID (for example, OPEN-3).<br>For more details, see Table 6.11.<br>▪ Serial number: The serial number of the storage subsystem (for example, 15001).<br>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.<br>This number combined with the storage subsystem name (shown in DskName) |

| Displayed Item | | Description |
|---|---|---|
| No summary displayed | Summary displayed | |
| | | identifies the LU that is accessed by a path.<br>▪ For 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.<br>▪ For the 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.<br>▪ 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. |
| `ChaPort`# | `CP` | Port number of the channel adapter, which identifies the CHA port that is mounted on the storage subsystem.<br>You can identify an actual CHA port by referencing this number from the storage subsystem management program. |
| `Status` | | Status of the path<br>▪ `Online`: Online<br>▪ `Offline(C)`: Offline status caused by a command operation<br>▪ `Offline(E)`: Offline due to an error<br>▪ `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.)<br>Paths that are `Offline(E)` or `Online(E)` require corrective action. The appropriate action can be determined by referring to 5.3. |
| `Type`# | `Type` | Attribute of the path<br>`Own`: Owner path<br>`Non`: Non-owner path<br>When connecting to 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.<br>To reset the `IO-Count` value to zero, execute the `dlnkmgr` command's `clear` operation. Executing the `clear` operation also resets the number of I/O errors (`IO-Errors`) to zero. For details about 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.<br>To reset the `IO-Errors` value to zero, execute the `dlnkmgr` command's `clear` operation. Executing the `clear` operation also clears the number of I/O operations (`IO-Count`) to zero.<br>For details about the `clear` operation, see section 6.2. |
| `DNum`# | – | Dev number, where 0 (fixed) is displayed.<br>This item pertains to a partition number. |
| `HDevName`# | – | Host device name.<br>The name of the logical device file (minus the partition number), `/dev/sddlm`[*a-p*][*a-p*], for the HDLM device. |

| Displayed Item | | Description |
|---|---|---|
| No summary displayed | Summary displayed | |
| | | [a-p] indicates a single letter from a to p. |
| IEP# | - | Information about the intermittent error.<br>This item is displayed only when you specify -iem with the -path parameter.<br>One of the following values is displayed for each path:<br>▪ -<br>Indicates that intermittent error monitoring is disabled or the monitoring time for an intermittent error is out of range.<br>▪ A value of at least 0<br>Indicates the number of errors that occurred during intermittent error monitoring.<br>▪ *<br>Indicates that an intermittent error occurred (automatic failback does not check the path). |

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 | Linux representation |
|---|---|
| Host port number (example: 0000) | Host ID (host port number) |
| Bus number (example: 0000) | Channel number (bus number) |
| Target ID (example: 000000000000003A) | Target ID |
| Host LU number(example: 0005) | Lun (host LU number) |

To obtain information about each element, use the HDLM-configuration definition utility (dlmcfgmgr). For details about this utility, see section 7.3.

**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<br>(Displays the following for the model name) |
| Thunder 9200 | product identifier# | 9200 |
| Thunder 9500V Series | | 9500V |
| TagmaStore AMS | | AMS |
| TagmaStore WMS | | WMS |

| Model names of storage subsystems | Product ID | |
|---|---|---|
| | Without the `-stname` parameter | With the `-stname` parameter (Displays the following for the model name) |
| Lightning 9900 Series | Emulation type# | 9910/9960 |
| Lightning 9900V Series | | 9970/9980 |
| TagmaStore USP | | USP |
| Hitachi Universal Storage Platform V | | USP_V |

#

If the `-c` parameter is specified together with the `-path` parameter, when the number of characters exceeds 10, any characters after the 7th character are displayed as 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 without selecting items to be displayed or displaying a summary. 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 `-hdev` or `-pathid`, the information about all LUs recognized by HDLM is displayed. For details on the contents of each displayed item, see Table 6.14. The subsequent parameters are:

`-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 partition number), `/dev/sddlm`[*a*-*p*][*a*-*p*], for the HDLM device. ([*a*-*p*]: a single letter from `a` to `p`)

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 display the LU information without selecting items to be displayed:

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -lu
Product       : 9500V
SerialNumber  : 0115
LUs           : 5

iLU  HDevName Device   PathID Status
0000 sddlmac  /dev/sds 000002 Online
              /dev/sdd 000007 Online
              /dev/sdi 000012 Online
              /dev/sdn 000017 Online
0001 sddlmad  /dev/sdt 000003 Online
              /dev/sde 000008 Online
              /dev/sdj 000013 Online
              /dev/sdo 000018 Online
0002 sddlmae  /dev/sdu 000004 Online
              /dev/sdf 000009 Online
              /dev/sdk 000014 Online
              /dev/sdp 000019 Online
0003 sddlmaa  /dev/sdq 000000 Online
              /dev/sdv 000005 Online
              /dev/sdg 000010 Online
              /dev/sdl 000015 Online
0004 sddlmab  /dev/sdr 000001 Online
              /dev/sdw 000006 Online
              /dev/sdh 000011 Online
              /dev/sdm 000016 Online
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` option are added to the items to be displayed by the `-lu` option and displayed.

When the value of the `-item` parameter is omitted or `all` is specified, all the items that can be displayed are displayed.

Table 6.12 shows the correspondence between the items that can be displayed by the `dlnkmgr view -lu -item` command and the parameter values for displaying each item.

**Table 6.12 Correspondence between the items displayed by the dlnkmgr view -lu -item command and the values of the -item parameter**

| Items displayed by the dlnkmgr view -lu -item command | Values specified after the -item parameter |
|---|---|
| SLPR | slpr |
| PathName | pn |
| ChaPort | cp |
| CLPR | clpr |
| Type | type |
| IO-Count | ic |

| Items displayed by the dlnkmgr view -lu -item command | Values specified after the -item parameter |
|---|---|
| `IO-Errors` | `ie` |
| `DNum` | `dnu` |
| `IEP` | `iep` |
| `HCTL` | `hctl` |
| 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 subsequent sub-parameters are:

`-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 partition number), `/dev/sddlm`[*a*-*p*][*a*-*p*], for the HDLM device. ([*a*-*p*]: a single letter from `a` to `p`)

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 display LU information by using iLUs as a sorting key, and selecting the items to be displayed. Execute the following command to display `SLPR`, `PathName`, `ChaPort`, `CLPR`, `Type`, `IO-Count`, `IO-Errors`, `DNum`, `IEP`, and HCTL:

```
#/opt/DynamicLinkManager/bin/dlnkmgr view -lu -item
Product      : USP
SerialNumber : 0376
LUs          : 4

iLU  SLPR HDevName Device   PathID PathName                           ChaPort CLPR Status     Type IO-Count   IO-Errors  DNum IEP HCTL
0023   12 sddlmad  /dev/sdd 000000 0002.0000.0000000000000000.0000 0B         1 Offline(E) Own          0          0    0 *   2.0.0.0
                   /dev/sdh 000004 0002.0000.0000000000000001.0000 1B         1 Offline(E) Own          0          0    0 3   2.0.1.0
                   /dev/sdl 000008 0003.0000.0000000000000000.0000 1B         1 Online     Own         60          0    0 2   3.0.0.0
                   /dev/sdp 000009 0003.0000.0000000000000001.0000 0B         1 Online     Own         68          0    0 3   3.0.1.0
0024   12 sddlmaa  /dev/sde 000001 0002.0000.0000000000000000.0001 0B         1 Offline(E) Own          0          0    0 *   2.0.0.1
                   /dev/sdi 000005 0002.0000.0000000000000001.0001 1B         1 Offline(E) Own          0          0    0 3   2.0.1.1
                   /dev/sdm 000010 0003.0000.0000000000000000.0001 1B         1 Online     Own     155817          0    0 2   3.0.0.1
                   /dev/sdq 000011 0003.0000.0000000000000001.0001 0B         1 Online     Own     155815          0    0 3   3.0.1.1
0025   12 sddlmab  /dev/sdf 000002 0002.0000.0000000000000000.0002 0B         1 Offline(E) Own          0          0    0 *   2.0.0.2
                   /dev/sdj 000006 0002.0000.0000000000000001.0002 1B         1 Offline(E) Own          0          0    0 -   2.0.1.2
                   /dev/sdn 000012 0003.0000.0000000000000000.0002 1B         1 Online     Own          0          0    0 -   3.0.0.2
                   /dev/sdr 000013 0003.0000.0000000000000001.0002 0B         1 Online     Own     269035          0    0 -   3.0.1.2
0026   12 sddlmac  /dev/sdg 000003 0002.0000.0000000000000000.0003 0B         1 Offline(E) Own          0          0    0 *   2.0.0.3
                   /dev/sdk 000007 0002.0000.0000000000000001.0003 1B         1 Offline(E) Own          0          0    0 3   2.0.1.3
                   /dev/sdo 000014 0003.0000.0000000000000000.0003 1B         1 Online     Own      10868          0    0 2   3.0.0.3
                   /dev/sds 000015 0003.0000.0000000000000001.0003 0B         1 Online     Own      11124          0    0 3   3.0.1.3
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

-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

The following example shows how to display a summary of LU information (without selecting items to be displayed):

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -lu -c
Product S/N  LUs iLU  HDevName Paths  OnlinePaths
9500V   0115   5 0000 sddlmac      4            4
                 0001 sddlmad      4            4
                 0002 sddlmae      4            4
                 0003 sddlmaa      4            4
                 0004 sddlmab      4            4
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 selecting items to be displayed)

-lu -c -item

The items specified with the -item option are added to the items to be displayed by the -lu -c option and displayed.

When the value of the -item parameter is omitted, all the items that can be displayed are displayed.

Table 6.13 shows the correspondence between the item that can be displayed by the dlnkmgr view -lu -c -item command and the parameter value for displaying the item.

**Table 6.13 Correspondence between the item displayed by the dlnkmgr view -lu -c -item command and the value of the -item parameter**

| Item displayed by the dlnkmgr view -lu -c -item command | Value specified after the -item parameter |
|---|---|
| SLPR | slpr |

The subsequent sub-parameter is:

-t

Does not display the title for each information item.

Example

The following example shows how to display a summary of LU information, and selecting the items to be displayed. Execute the following command to display SLPR:

```
#/opt/DynamicLinkManager/bin/dlnkmgr view -lu -c -item
Product       S/N  LUs iLU  SLPR HDevName Paths  OnlinePaths
9500V         0375   4 0010    - sddlmai       4           1
                       0011    - sddlmaj       4           1
                       0012    - sddlmaa       4           1
                       0013    - sddlmab       4           1
KAPL01001-I The HDLM command completed normally. Operation name = view, completion
time = yyyy/mm/dd hh:mm:ss
#
```

## Items of 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 `-path -item` parameter.
- Summary displayed: The user specifies the `-lu -c` parameter.

**Table 6.14 LU information**

| Displayed item | | Description |
|---|---|---|
| **No summary displayed** | **Summary displayed** | |
| `Product` | | Model ID of the storage subsystem |
| `SerialNumber` | `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 in the storage subsystem.<br>This number combined with the storage subsystem name (shown in `DskName`) identifies the LU that is accessed by a path.<br>▪ For the Lightning 9900 Series, Lightning 9900V Series, and TagmaStore USP, indicated by a hexadecimal number. The first two characters of `iLU` are the internal CU number within the CU, and the last two characters are the internal LU number within the CU.<br>▪ For the 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.<br>▪ 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 functionality for the storage subsystem for the target LU is not supported. |
| `HDevName`#1 | `HDevName` | Host device name.<br>The name of the logical device file (minus the partition number), `/dev/sddlm`[*a-p*][*a-p*], for the HDLM device.<br>[*a-p*] indicates a single letter from `a` to `p`. |
| `Device` | - | SCSI device associated with an HDLM device. If the SCSI device is not connected when a host starts, or the assigned LU is released, this column |

| Displayed item | | Description |
|---|---|---|
| **No summary displayed** | **Summary displayed** | |
| | | displays a hyphen (-). |
| PathID | – | AutoPATH_ID indicated by a decimal number.<br>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 HDLM-configuration definition utility (dlmcfgmgr). |
| 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:<br>▪ Host port number (hexadecimal number)<br>▪ Bus number (hexadecimal number)<br>▪ Target ID (hexadecimal number)<br>▪ Host LU number (hexadecimal number)<br>For details about each element of the path name and its representation in Linux, see Table 6.10. |
| ChaPort#1 | – | Port number of the channel adapter, which identifies the CHA port that is mounted on the storage subsystem.<br>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 belongs, indicated by a decimal number from 0 to 31. Note that a hyphen (-) is displayed if the following items are subject to display:<br>▪ CHA ports in the storage subsystem that do not support cache logical partition functionality<br>▪ Paths connected to the Snapshot Image of the Copy-on-write Snapshot of the TagmaStore AMS/WMS series |
| Status | – | Status of the path<br>Online: Online<br>Offline(C): Offline status caused by a command operation<br>Offline(E): Offline due to an error<br>Online(E): Failure has occurred (If none of the paths accessing one Dev have an Online status, one of those paths is changed to the Online(E) status.)<br>Paths that are Offline(E) or Online(E) require corrective action. The appropriate action can be determined by referring to 5.3. |
| Type#1 | – | Attribute of the path<br>Own: Owner path<br>Non: Non-owner path<br>When connecting to 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.<br>To reset the IO-Count value to zero, execute the dlnkmgr |

| Displayed item | | Description |
| --- | --- | --- |
| No summary displayed | Summary displayed | |
| | | command's clear operation. Executing the clear operation also resets the number of I/O errors (IO-Errors) to zero. For details about the clear operation, see section 6.2. |
| 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.<br>To reset the IO-Errors value to zero, execute the dlnkmgr command's clear operation. Executing the clear operation also clears the number of I/O operations (IO-Count) to zero.<br>For details about the clear operation, see section 6.2. |
| DNum#1 | - | Dev number, where 0 (fixed) is displayed.<br>This item pertains to a partition number. |
| 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.<br>One of the following values is displayed for each path:<br>▪ -: Indicates that intermittent error monitoring is disabled or the monitoring time for an intermittent error is out of range.<br>▪ A value of at least 0: Indicates the number of errors that occurred during intermittent error monitoring.<br>▪ *: Indicates that an intermittent error occurred (automatic failback does not check the path). |
| HCTL#1 | - | SCSI device configuration information. The following information is shown in "*a.b.c.d*" format:<br>*a*: Host port number (hexadecimal number)<br>*b*: Bus number (hexadecimal number)<br>*c*: Target ID (hexadecimal number)<br>*d*: Host LU number (hexadecimal number) |
| - | 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 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 are offline. In this case, you should check the offline 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 corresponding information about an HDLM device, SCSI device, and LDEV

`-drv`

Using this parameter displays the PathID, the HDLM device, SCSI device for the HDLM device, and information on the LDEV in the storage subsystem (storage subsystem model name, serial number, and LU number, separated by period).

For details on the contents of each display item, see Table 6.15.

`-t`

Does not display the title for each information item.

**Table 6.15 To display corresponding information about an HDLM device, SCSI device, and LDEV**

| Item | Description |
|---|---|
| `PathID` | AutoPATH_ID indicated by a decimal number.<br>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 HDLM-configuration definition utility (`dlmcfgmgr`). |
| `HDevName` | Host device name.<br>The name of the logical device file (minus the partition number), `/dev/sddlm`[*a-p*][*a-p*], for the HDLM device.<br>[*a-p*] indicates a single letter from `a` to `p`. |
| `Device` | SCSI device associated with an HDLM device. If the SCSI device is not connected when a host starts, or the assigned LU is released, this column displays a hyphen (`-`). |
| `LDEV` | The model ID, serial number, and iLU number for the storage subsystem, separated by periods. Information about LDEV can identify an actual device that is managed by HDLM. |

Example

To display corresponding information about an HDLM device, SCSI device, and LDEV

```
# /opt/DynamicLinkManager/bin/dlnkmgr view -drv
PathID HDevName     Device           LDEV
000000 sddlmaa      /dev/sdag        9970/9980.15001.05B7
000001 sddlmaa      /dev/sdq         9970/9980.15001.05B7
000002 sddlmab      /dev/sdr         9970/9980.15001.05B0
000003 sddlmac      /dev/sds         9970/9980.15001.05B1
000004 sddlmad      -                9970/9980.15001.05B2
000005 sddlmae      /dev/sdu         9970/9980.15001.05B3
000006 sddlmaf      /dev/sdv         9970/9980.15001.05B4
000007 sddlmag      /dev/sdw         9970/9980.15001.05B5
000008 sddlmah      /dev/sdx         9970/9980.15001.05B6
```

```
KAPL01001-I The HDLM command completed normally. Operation name = view, completion time =
yyyy/mm/dd hh:mm:ss
#
```

### 6.7.2.5 To display the format of the view operation

`-help`

Use this parameter to display the `view` operation format.

Example

The following example shows how to display the format of the `view` operation.

```
# /usr/DynamicLinkManager/bin/dlnkmgr view -help
view:
  Format
    dlnkmgr view -sys [ -sfunc | -msrv | -adrv | -pdrv | -lic ] [-t]
    dlnkmgr view -path [ -hdev HostDeviceName ] [-stname] [-iem]
                                                 [-srt {pn | lu | cp}] [-t]
    dlnkmgr view -path
                 -item [pn] [dn] [lu] [cp] [type] [ic] [ie] [dnu] [hd] [iep]
                  [ -hdev HostDeviceName ] [-stname] [-srt {pn | lu | cp}] [-t]
    dlnkmgr view -path -c [-stname] [-srt {lu | cp}] [-t]
    dlnkmgr view -lu [ -hdev HostDeviceName | -pathid AutoPATH_ID ] [-t]
    dlnkmgr view -lu
                 -item [ [slpr] [pn] [cp] [clpr] [type]
                                         [ic] [ie] [dnu] [iep] [hctl] | all ]
                            [ -hdev HostDeviceName | -pathid AutoPATH_ID ] [-t]
    dlnkmgr view -lu -c [-t]
    dlnkmgr view -lu -c -item [slpr] [-t]
    dlnkmgr view -drv [-t]
KAPL01001-I The HDLM command completed normally. Operation name = view, 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

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 section 7.2.

- The `dlmcfgmgr` HDLM-configuration definition utility

  This utility enables you to add, update, or delete information that HDLM requires to manage paths. For details on the `dlmcfgmgr` utility, see section 7.3.

- The `dlmlfk` utility for supporting LifeKeeper

  This utility enables or disables the use of HDLM devices from LifeKeeper. For details on the `dlmlfk` utility, see section 7.4.

- The `dlmmkinitrd` utility for supporting a boot disk

  This utility creates an initial RAM disk image file for using an HDLM device as a boot disk. For details on the `dlmmkinitrd` utility, see section 7.5.

- The `dlmpr` utility for clearing HDLM persistent reservation

  When multiple hosts share an LU, if the persistent reservation of the LU is not canceled for some reason, this utility clears the Reservation Key to cancel the persistent reservation. For details on the `dlmpr` utility, see section 7.6.

- The `dlmsetopt` utility for setting HDLM driver option

  This utility changes the settings for the HDLM filter driver. Restart the host machine to apply the changed set values. For details on the `dlmsetopt` utility, see section 7.7.

- The `installgetras` utility for collecting HDLM installation error information

  If an installation error occurs, this utility collects files that contain information to be submitted to your HDLM vendor or maintenance company. For details on the `installgetras` utility, see section 7.8.

***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, and libraries. The collected information is archived in a file and saved to the directory that you specified. The following files are output:

- `cluster.tar.gz`

  This file contains compressed cluster information.

- `hbsa.tar.gz`

  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.gz`

  This file contains compressed HDLM information and system information.

For details about the information that is stored in each file, see section 7.2.3.

When you want to collect information other than that in 7.2.3, define the information to collect in the *information- collection- definition file*. Information defined in *information- collection- definition file is compressed into* `getras.tar.gz`.

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 file-that-defines-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 file-that-defines-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.

`-f` *file-that-defines-information-to-be-collected*

Use this parameter when you want to specify certain directories or files to be collected. In this parameter, specify the *file-that-defines-information-to-be-collected*, which defines the files and directories you want to collect.

Use an absolute path to specify *file-that-defines-information-to-be-collected.*

Figure 7.1 shows an example of the coding in this file. In this example, the information specified to be collected is as follows: the OS files that are not collected by the `DLMgetras` utility by default, and the PostgreSQL files when PostgreSQL is an upper-level application program.

| Coding | Description |
|---|---|
| `# System Memory Information`<br>`/proc/kcore` | Collects information about the system's physical memory |
| `# kernel symbols`<br>`/var/log/ksyms.*` | Collects the kernel symbol information that was exported |
| `# PostgreSQL Database`<br>`/usr/local/pgsql/data` | Collects the user data directory for the database (PostgreSQL). |

**Figure 7.1 Coding Example of a File that Defines the Information to Be Collected**

Rules for coding a file that defines 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 the hash mark (#) are handled as comment lines.
- If the 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.

***Notes:***

- Because `DLMgetras` first stores error information in the specified output directory 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 the existing files; 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.

- When the `DLMgetras` utility is executed but the file to be obtained does not exist, the message `KAPL10033-W` is output to `getras.log`. However, the `DLMgetras` utility operations are not affected.

- If you execute the `DLMgetras` utility in the EM64T/AMD64 environment while running SUSE LINUX Enterprise Server 9, SUSE LINUX Enterprise Server 10 or Red Hat Enterprise Linux AS4/ES4, an error message (KAPL05023-E) will be output to the syslog file. However, HDLM operations are not affected.

### 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.

- `cluster.tar.gz`

  This file contains cluster information. Table 7.1 shows the detailed error information stored in the `cluster.tar.gz` file. Table 7.2 lists and describes the detailed information stored in the `cluster.tar.gz` file, which is recorded by the OS and cluster commands when the `DLMgetras` utility is executed.

- `hbsa.tar.gz`

  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.gz`

  This file contains HDLM information and system information.

  When you execute the `DLMgetras` utility, specifying the file that defines information to be collected, the `getras.tar.gz` file contains the information stored in the file that defines information to be collected.

  Table 7.3 lists the detailed information stored in the `getras.tar.gz` file. Table 7.4 lists and describes the detailed information stored in the `getras.tar.gz` file, which is recorded by the OS and HDLM commands when the `DLMgetras` utility is executed.

  Table 7.3 lists the collected information with or without specifying the file that defines information to be collected.

**Table 7.1 Information Stored in the cluster.tar.gz File**

| Output directory#1 | Files | Explanation |
|---|---|---|
| `/etc` or `/etc/cluster` | `cluster.conf`#2 or `cluster.xml`#2 | RHCM configuration file (when using Red |

| Output directory#1 | Files | Explanation |
|---|---|---|
| | | Hat Enterprise Linux 3 or Red Hat Enterprise Linux 4) |
| /etc/ha.d | authkeys | HeartBeat setting information (when using SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10) |
| | ha.cf | |
| | haresources | |
| /etc/default | LifeKeeper | LifeKeeper standard settings file |
| /etc/ha.d/resource.d | All subdirectories and files | SuSE HeartBeat cluster resource information (when using SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10) |
| /etc/VRTSvcs/conf/config | All subdirectories and files#2 | VCS setting file |
| /opt/LifeKeeper/subsys/scsi/resources/hostadp | device_info | List of devices used by LifeKeeper |
| | device_info.* | device_info for another node in a cluster |
| | device_pattern | Device pattern information recognized by LifeKeeper |
| /opt/LifeKeeper/subsys/scsi/resources/disk | device_pattern | Device pattern information recognized by LifeKeeper |
| | disk.reserve | List of disks for which LifeKeeper has currently put on reserve |
| /opt/LifeKeeper/subsys/scsi/resources/disk/disk.mapping | All subdirectories and files | Information for the correspondence between sg devices and character devices |
| /opt/LifeKeeper/config | core.inittab | LifeKeeper configuration information |
| | equivalencies | |
| | event.lastid | |
| | flg | |
| | lkGUIapp.kdelnk | |
| | lkgui.cert | |
| | lkgui.inittab | |
| | LK_INITDONE | |
| | lk.inittab | |
| | LK_START_TIME | |
| | LK_XREFDONE | |
| | networks | |
| | systems | |
| | UNAME | |
| | varconfig | |
| /opt/LifeKeeper/locale | *.properties | property file for LifeKeeper |

| Output directory#1 | Files | Explanation |
|---|---|---|
| /var/log | ha-log | HeartBeat log file (when using SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10) |
| | ha-debug | HeartBeat debug log file (when using SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10) |
| /var/VRTSvcs/log | All subdirectories and files#2 | VCS log |

#1

Each output directory is created in the directory that is specified when the user expands the `cluster.tar.gz` file.

#2

This information is collected only if the software for which information is to be collected is installed.

**Table 7.2 Information Stored in the cluster.tar.gz File, Which Is Recorded by the OS and HDLM Commands When the DLMgetras Utility Is Executed**

| Executed command | Files | Explanation |
|---|---|---|
| /bin/rpm -qlv clumanager or /bin/rpm -qlv rgmanager | RHCM_rpm-qlv.txt# | RHCM package file information (when using Red Hat Enterprise Linux 3 or Red Hat Enterprise Linux 4) |
| /bin/rpm -qlv heartbeat | HA_rpm-qlv.txt | HeartBeat package file information |
| /bin/rpm -qlv VRTSvcs | VCS_rpm-qlv.txt# | VCS package file information |
| /etc/init.d/heartbeat status | HA_status.txt | HeartBeat cluster status information (when using SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10) |
| /opt/LifeKeeper/bin/app_list | LK_app_list.txt | List of applications |
| /opt/LifeKeeper/bin/dep_list | LK_dep_list.txt | List of dependencies |
| /opt/LifeKeeper/bin/eqv_list | LK_eqv_list.txt | List of equivalent relationships |
| /opt/LifeKeeper/bin/flg_list | LK_flg_list.txt | List of flags |
| /opt/LifeKeeper/bin/ins_list | LK_ins_list.txt | List of resource instances |
| /opt/LifeKeeper/bin/lcdstatus | LK_lcdstatus.txt | LCD status |
| /opt/LifeKeeper/bin/lcduname | LK_lcduname.txt | LCD name |
| /opt/LifeKeeper/bin/lk_log log | LK_log_log.txt | LifeKeeper log |
| /opt/LifeKeeper/bin/lk_log TTYLCM | LK_log_TTYLCM.txt | Log of TTY CommPath events and status changes |

| Executed command | Files | Explanation |
|---|---|---|
| /opt/LifeKeeper/bin/lk_log LCM | LK_log_LCM.txt | Log of TCP CommPath events and status changes |
| /opt/LifeKeeper/bin/lk_log LCD | LK_log_LCD.txt | Log of information and status changes related to the LifeKeeper database |
| /opt/LifeKeeper/bin/lk_log remote_execute | LK_log_ remote_execute.txt | Remote request log |
| /opt/LifeKeeper/bin/lk_log SNMP | LK_log_SNMP.txt | SNMP trap log |
| /opt/LifeKeeper/bin/lk_log GUI | LK_log_GUI.txt | Log of GUI server connection events and status changes |
| /opt/LifeKeeper/bin/lktest | LK_lktest.txt | LifeKeeper status display |
| /opt/LifeKeeper/bin/lmdiag | LK_lmdiag.txt | Diagnostic information |
| /opt/LifeKeeper/bin/lmhostid | LK_lmhostid.txt | Host ID |
| /opt/LifeKeeper/bin/net_list | LK_net_list.txt | List of network paths |
| /opt/LifeKeeper/bin/sys_list | LK_sys_list.txt | List of systems in a cluster |
| /opt/LifeKeeper/bin/typ_list | LK_typ_list.txt | List of resource layers |
| /opt/LifeKeeper/bin/sys_getstate -s hostname | LK_getstate.txt | Displays the system status as DEAD, ALIVE, or UNKNOWN. |
| /opt/LifeKeeper/bin/sys_getdescr -s hostname | LK_getdescr.txt | Displays the system status as text. |
| /opt/LifeKeeper/lkadm/bin/getpriority | LK_getpriority.txt | Priority |
| /opt/LifeKeeper/lkadm/bin/scsisdev | LK_scsisdev.txt | SCSI device information |
| /opt/LifeKeeper/lkadm/bin/shared_devices | LK_shared_devices.txt | Shared device information |
| /opt/LifeKeeper/lkadm/bin/test_lk hostname | LK_test_lk.txt | LifeKeeper test |
| /opt/VRTSvcs/bin/haclus -display | VCS_haclus-display.txt[#] | VCS cluster configuration information |
| /opt/VRTSvcs/bin/haclus -notes | VCS_haclus-notes.txt[#] | VCS management information |
| /opt/VRTSvcs/bin/hagrp -display | VCS_hagrp-display.txt[#] | VCS group information |
| /opt/VRTSvcs/bin/hagrp -dep | VCS_hagrp-dep.txt[#] | VCS group dependent information |
| /opt/VRTSvcs/bin/hagrp -state | VCS_hagrp-state.txt[#] | VCS group status information |
| /opt/VRTSvcs/bin/hares -display | VCS_hares-display.txt[#] | VCS resource information |

| Executed command | Files | Explanation |
|---|---|---|
| /opt/VRTSvcs/bin/hares -dep | VCS_hares-dep.txt# | VCS resource dependent information |
| /opt/VRTSvcs/bin/hares -state | VCS_hares-state.txt# | VCS resource status information |
| /opt/VRTSvcs/bin/hastatus -summary | VCS_hastatus-summary.txt# | VCS cluster status information |
| /opt/VRTSvcs/bin/hasys -display | VCS_hasys-display.txt# | VCS node information |
| /opt/VRTSvcs/bin/hasys -state | VCS_hasys-state.txt# | VCS node status information |
| /sbin/cluconfig -l | RHCM_cluconfig-l.txt# | RHCM cluster configuration information (when using Red Hat Enterprise Linux 3 or Red Hat Enterprise Linux 4) |
| /sbin/clustat | RHCM_clustat.txt# | RHCM cluster status information (when using Red Hat Enterprise Linux 3 or Red Hat Enterprise Linux 4) |
| /sbin/clustonith -L | RHCM_Clustonith-L.txt# | RHCM utility information (when using Red Hat Enterprise Linux 3 or Red Hat Enterprise Linux 4) |
| /sbin/clustonith -vSl | RHCM_Clustonith-vSl.txt# | RHCM utility list information (when using Red Hat Enterprise Linux 3 or Red Hat Enterprise Linux 4) |
| /usr/sbin/stonith -L | HA_stonith-L.txt | HeartBeat utility list information (when using SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10) |
| /usr/sbin/stonith -vSl | HA_stonith-vSl.txt | HeartBeat utility list information (when using SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10) |

#

This information is collected only if the software for which information is to be collected is installed.

**Table 7.3 Information Stored in the getras.tar.gz File**

| Output directory#1 | Files | Explanation |
|---|---|---|
| The output destination directory specified when the DLMgetras utility is executed | getras.log | DLMgetras utility log file |
| IA32, EM64T, or AMD64: /boot/grub | menu.lst | Boot selection list (when using SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10) |
| | grub.conf | Boot loader setting file (when using Red Hat Enterprise Linux 3 or Red Hat Enterprise Linux 4) |
| IPF (Red Hat Enterprise Linux3, Red Hat Enterprise Linux4): /boot/efi/efi/redhat<br>IPF (SUSE LINUX Enterprise Server 9, SUSE LINUX Enterprise Server 10): | elilo.conf | Boot loader setting file |

| Output directory#1 | Files | Explanation |
|---|---|---|
| /boot/efi/efi/SuSE | | |
| /etc | exports | Information about exports of the NFS file system |
| | evms.conf | Default value information file (when using SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10) |
| | fstab | Auto-mount information for the file system at startup |
| | hba.conf | HBA settings file |
| | inittab | Boot process script |
| | lilo.conf | Boot loader settings file (when using IA32, EM64T, or AMD64) |
| | lpfc.conf | Emulex driver settings file |
| | lvmtab | LVM settings information file (when using Red Hat Enterprise Linux 3) |
| | modprobe.conf | Load module information file (when using Red Hat Enterprise Linux 4, SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10) |
| | modprobe.conf.local | Load module information file(when using SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10) |
| | modules.conf | Parameter information, etc. (when using Red Hat Enterprise Linux 3) |
| | mtab | Mount information |
| | qla*.conf | QLogic driver settings file |
| | raidtab | md device configuration file |
| | raw | Character-type device information (when using SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10) |
| | rc.sysinit | Start-up initialization script (when using Red Hat Enterprise Linux 3 or Red Hat Enterprise Linux 4) |
| | syslog.conf | syslogd settings file (when using Red Hat Enterprise Linux 3, Red Hat Enterprise Linux 4, or SUSE LINUX Enterprise Server 9) |
| /etc/hotplug | blacklist | Information about modules that have not been loaded by a hot plug |
| /etc/init.d | boot | Script for starting and initializing (when using SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10) |
| | boot.* | Script for starting and initializing (when using SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10) |
| | halt | Shutdown script |

| Output directory#1 | Files | Explanation |
|---|---|---|
| /etc/init.d/boot.d | all subdirectories and files | Script for starting and initializing (when using SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10) |
| /etc/lvm | .cache | Valid LVM information file (when using Red Hat Enterprise Linux 4, SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10) |
| | lvm.conf | LVM configuration settings file (when using Red Hat Enterprise Linux 4, SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10) |
| /etc/lvm/backup | All subdirectories and files | LVM related file (when using Red Hat Enterprise Linux 4, SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10) |
| /etc/opt/DynamicLinkManager | .dlmfdrv.conf | Configuration definition file |
| | .dlmfdrv1.conf<br>.dlmfdrv2.conf<br>.dlmfdrv3.conf | Backup files for the configuration definition file |
| | .dlmfdrv.unconf | File that defines non-HDLM-managed disks |
| | .dlmfdrv1.unconf<br>.dlmfdrv2.unconf<br>.dlmfdrv3.unconf | Backup files for the file that defines non-HDLM-managed disks |
| | dlmmgr.xml | Manager setting information |
| | dlmwebagent.properties | Configuration file for HDLM remote access interface |
| | hdlmboot*.log | HDLM boot log |
| | hdlm_kernel_version | Information about the kernel version used when HDLM was installed |
| /etc/sysconfig | boot | Boot process information file (when using SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10) |
| | hardware | Hardware information (when using SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10) |
| | hotplug | Hot plug option settings file (when using SUSE LINUX Enterprise Server 9) |
| | hwconf | Hardware information (when using Red Hat Enterprise Linux 3 or Red Hat Enterprise Linux 4) |
| | irqbalance | Interrupt processing information |
| | kernel | Kernel configuration file (when using SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10) |
| | lvm | Boot LVM information file (when using SUSE LINUX Enterprise Server 9 or SUSE LINUX |

| Output directory#1 | Files | Explanation |
|---|---|---|
| | | Enterprise Server 10) |
| | rawdevices | Character-type device information (when using Red Hat Enterprise Linux 3 or Red Hat Enterprise Linux 4) |
| | scsidev | SCSI device settings file (when using SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10) |
| | sysctl | sysrq information file (when using SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10) |
| /etc/sysconfig/hdlm-scripts | All subdirectories and files | HDLM-related script |
| /etc/syslog-ng | syslog-ng.conf | syslog-ng settings file (when using SUSE LINUX Enterprise Server 10) |
| /etc/udev | udev.conf | udev configuration file#2<br>(when using Red Hat Enterprise Linux 4, SUSE LINUX Enterprise Server 9, or SUSE LINUX Enterprise Server 10) |
| | udev.rules | udev rule file<br>(when using SUSE LINUX Enterprise Server 9) |
| /etc/udev/rules.d | *.rules file | udev rule file<br>(when using Red Hat Enterprise Linux 4, or SUSE LINUX Enterprise Server 10) |
| /etc/cron* | All subdirectories and files | cron file |
| /lib/modules/*kernel-name* | modules.dep | Module dependency information |
| /opt/hitachi/HNTRLib/mmap | hntrmmap.mm | Memory mapped file (HNTRLib) |
| /opt/hitachi/HNTRLib2/mmap | hntr2mmap.mm | Memory mapped file (HNTRLib2) |
| /proc | cmdline | Parameter information for kernel startup |
| | config.gz | Kernel configuration information (when using SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10) |
| | cpuinfo | CPU information |
| | devices | Device information |
| | diskstats | Disk statistics (when using Red Hat Enterprise Linux 4, SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10) |
| | dma | DMA-related information |
| | filesystems | File system list |
| | interrupts | Information about interruption to a processor |
| | iomem | I/O memory map information |
| | ioports | I/O port information |
| | kallsyms | Kernel symbol information (when using Red |

| Output directory#1 | Files | Explanation |
|---|---|---|
| | | Hat Enterprise Linux 4, SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10) |
| | ksyms | Kernel symbol information (when using Red Hat Enterprise LINUX 3) |
| | loadavg | Information about average load time |
| | locks | File lock information |
| | mdstat | md device information |
| | meminfo | Memory information |
| | misc | MISC Driver information |
| | modules | Load module information |
| | mounts | Mount information |
| | partitions | Information about the partitions recognized by the kernel |
| | pci | PCI device list (when using Red Hat Enterprise Linux 3, or Red Hat Enterprise Linux 4) |
| | slabinfo | Slab cache information |
| | stat | Statistics |
| | swaps | SWAP information |
| | uptime | System's operating time information |
| | version | Kernel version information |
| /proc/fs | All subdirectories and files | File system information |
| /proc/irq | All subdirectories and files | IRQ information |
| /proc/lvm | All subdirectories and files | LVM-related information (when using Red Hat Enterprise Linux 3) |
| /proc/net | All subdirectories and files | Network device information |
| /proc/scsi | All subdirectories and files | SCSI device information |
| /proc/sysvipc | All subdirectories and files | SystemV IPC object information |
| /proc/sys/dev | All subdirectories and files | Device-related parameter information |
| /proc/sys/fs | All subdirectories and files | File system parameter information |
| /proc/sys/kernel | All subdirectories and files | Kernel parameter information |
| /proc/sys/vm | All subdirectories and files | Virtual memory parameter information |
| /sys/class/scsi_host/host* | All files | SCSI host information (when using Red Hat Enterprise Linux 4, SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10) |
| /var/log | boot.* | Startup log and message information (when using SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10) |

| Output directory#1 | Files | Explanation |
|---|---|---|
| | `boot.log*` | Boot log (when using Red Hat Enterprise Linux 3 or Red Hat Enterprise Linux 4) |
| | `cron*` | Cron log file (when using Red Hat Enterprise Linux 3 or Red Hat Enterprise Linux 4) |
| | `ksyms.*` | Kernel symbol information logs (when using Red Hat Enterprise Linux 3) |
| | `messages*` | Syslog file (including the cluster logs) |
| | `rpmpkgs*` | Package information (when using Red Hat Enterprise Linux 3 or Red Hat Enterprise Linux 4) |
| | `cron*` | cron log file (when using Red Hat Enterprise Linux 3 or Red Hat Enterprise Linux 4) |
| `/var/opt/DynamicLinkManager/log` | `dlmmgr1.log`<br>:<br>`dlmmgr16.log` | HDLM manager log (including the driver log) |
| | `dlmwebagent[1-N].log`#3 | HDLM remote access interface log |
| | `dlminquiry.log` | Inquiry information log |
| | `dlmcfgmgr1.log`,<br>`dlmcfgmgr2.log` | `dlmcfgmgr` utility log |
| | `.dlmomtrace1.log`<br>`.dlmomtrace2.log`<br>:<br>`.dlmomtrace5.log` | On-memory trace log |
| | `hdlmtr1.log`<br>:<br>`hdlmtr64.log` | Trace file |
| `/var/opt/DynamicLinkManager/log/mmap` | `hdlmtr.mm` | Trace management file |
| `/var/opt/hitachi/HNTRLib/spool`<br>(by default)#4 | `hntr1.log`,<br>`hntr2.log`,<br>:<br>`hntr16.log`<br>(by default) | Integrated trace file (HNTRLib) |
| `/var/opt/hitachi/HNTRLib2/spool`<br>(by default)#4 | `hntr21.log`,<br>`hntr22.log`,<br>:<br>`hntr216.log`<br>(by default)#5 | Integrated trace file (HNTRLib2) |
| `/var/spool/cron` | All files | User-defined cron entry information (when using Red Hat Enterprise Linux 3 or Red Hat Enterprise Linux 4) |
| `/var/spool/cron/tabs` | All files | User-defined cron entry information (when using SUSE LINUX Enterprise Server 9 or |

| Output directory[#1] | Files | Explanation |
|---|---|---|
| | | SUSE LINUX Enterprise Server 10)) |
| /var/tmp/hdlminstlog | All subdirectories and files | Installation log information |

#1

Each output directory is created in the directory that is specified when the user expands the getras.tar.gz file.

#2

Information about the rules files specified in udev_rules= in the udev.conf file is also collected.

If a file name is specified, information about that file is collected.

If a directory name is specified, information about the *.rules files in that directory is collected.

#3

The value *N* depends on the setting in the dlmwebagent.properties file. The default is 8.

#4

If a user-defined output directory is specified in the Hitachi Network Objectplaza Trace Library utility, this is the specified directory.

#5

If a user-defined prefix for the integrated trace file is specified in the Hitachi Network Objectplaza Trace Library utility, the hntr prefix is replaced by the specified prefix appended with a 2, then the file number (1 to 16).

**Table 7.4 Information Stored in the getras.tar.gz File, Which Is Recorded by the OS and HDLM Commands When the DLMgetras Utility Is Executed**

| Executed command | Files | Explanation |
|---|---|---|
| /bin/df | bdf.txt | Disk usage of the file system |
| /bin/dmesg | dmesg.txt | System diagnostic message |
| /bin/ls -altR /dev | dev_dsk.txt | List of disks to be handled as a block-type device |
| /bin/ls -altR /dev/raw | dev_raw.txt | List of disks to be handled as a character-type device |
| /bin/ls -altR /etc/rc.d | rc.txt | List of script files (when using Red Hat Enterprise Linux 3 or Red Hat Enterprise Linux 4) |
| /bin/ls -altR /etc/init.d | init.txt | List of script files (when using SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10)) |
| /bin/ls -altR /opt/DynamicLinkManager | DLMfilelist.txt | List of files in the /opt/DynamicLinkManager directory |

| Executed command | Files | Explanation |
|---|---|---|
| /bin/ls -altRZ /dev | dev_dsk_security.txt | List of disks, containing security context information, that are handled as block-type devices (when using Red Hat Enterprise Linux 4) |
| /bin/ls -altRZ /dev/raw | dev_raw_security.txt | List of disks, containing security context information, that are handled as character-type devices (when using Red Hat Enterprise Linux 4) |
| /bin/ls -altRZ /etc/rc.d | rc_security.txt | List of script files containing security context information (when using Red Hat Enterprise Linux 4) |
| /bin/ls -altRZ /opt/DynamicLinkManager | DLMfilelist_security.txt | List of files containing security context information in the /opt/DynamicLinkManager directory (when using Red Hat Enterprise Linux 4) |
| /bin/mount -v | mount-v.txt | File system mount information |
| /bin/ps -elf | ps-elf.txt | Process information (when using Red Hat Enterprise Linux 3) |
| /bin/ps -elF | ps-elF.txt | Process information (when using Red Hat Enterprise Linux 4, SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10) |
| /bin/ps -elL | ps-elL.txt | Thread information displayed in the LWD column (when using Red Hat Enterprise Linux 4, SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10) |
| /bin/ps -elT | ps-elT.txt | Thread information displayed in the SPID column (when using Red Hat Enterprise Linux 4, SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10) |
| /bin/ps -elm | ps-elm.txt | Thread information displayed after process information (when using Red Hat Enterprise Linux 4, SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10) |
| /bin/ps -elZ | ps-elZ.txt | Process information that contains security context information (when using Red Hat Enterprise Linux 4) |
| /bin/rpm -qai | rpm-qai.txt | List of installed packages |
| /bin/rpm -qlv HDLM | HDLM_rpm-qlv.txt | HDLM package file information |
| /bin/rpm -V HDLM | rpm-V.txt | HDLM package revised information |
| /bin/uname -a | uname-a.txt | OS version |
| /opt/DynamicLinkManager/bin/dlmpr -k | dlmpr-k.txt | HDLM persistent reservation information (when using a kernel that supports LifeKeeper) |
| /opt/DynamicLinkManager/bin/dlmgetomtrace | dlmgetomtrace.dmp | Trace information of HDLM functions |
| /opt/DynamicLinkManager/bin/dlnkmgr view -lu -item slpr pn cp clpr type ic ie dnu hctl | dlmmgr-lu.txt | HDLM LU information |

| Executed command | Files | Explanation |
|---|---|---|
| /opt/DynamicLinkManager/bin/dlnkmgr view -lu -item all | dlmmgr-lu-all.txt | HDLM LU information (all items, including the number of times an intermittent error occurred) |
| /opt/DynamicLinkManager/bin/dlnkmgr view -path | dlmmgr-path.txt | HDLM path information |
| /opt/DynamicLinkManager/bin/dlnkmgr view -path -iem | dlmmgr-path-iem.txt | HDLM path information (including the number of times an intermittent error occurred) |
| /opt/DynamicLinkManager/bin/dlnkmgr view -drv | dlmmgr-drv.txt | Correspondence between an HDLM device, SCSI device, and LDEV information |
| /opt/DynamicLinkManager/bin/dlnkmgr view -sys,<br>/opt/DynamicLinkManager/bin/dlnkmgr view -sys -sfunc,<br>/opt/DynamicLinkManager/bin/dlnkmgr view -sys -msrv,<br>/opt/DynamicLinkManager/bin/dlnkmgr view -sys -adrv,<br>/opt/DynamicLinkManager/bin/dlnkmgr view -sys -pdrv,<br>/opt/DynamicLinkManager/bin/dlnkmgr view -sys -lic | dlmmgr-sys.txt | HDLM program information |
| /sbin/chkconfig --list | chkconfig.txt | Service settings |
| /sbin/e2label | e2label.txt | Label information (all the SCSI device names and the corresponding label names are displayed in each line. If there is not a label, the hyphen (-) is displayed instead. If the Linux file system other than the ext file system is used, a label name is not recorded for label information.) |
| /sbin/dlmcfgmgr -v | dlmcfgmgr-v.txt | Configuration information and management status of HDLM devices (when using Red Hat Enterprise Linux 3) |
| /sbin/dlmcfgmgr -v -udev | dlmcfgmgr-v.txt | Configuration information and management status of HDLM devices (when using Red Hat Enterprise Linux 4, SUSE LINUX Enterprise Server 9, or SUSE LINUX Enterprise Server 10). |
| /sbin/fdisk -l | fdisk-l.txt | Disk information (when using IA32, EM64T, or AMD64) |
| /sbin/ifconfig -a | ifconfig-a.txt | Information about all NICs |
| /sbin/lilo -t | lilo-t.txt | Boot loader information (when using IA32, EM64T, or AMD64) |
| /sbin/lsmod | lsmod.txt | List of load modules |
| /sbin/lspci -v | lspci-v.txt | PCI information |
| /sbin/lvmdiskscan | lvmdiskscan.txt | lvm disk information (when using Red Hat Enterprise Linux 3, SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10) |
| /sbin/lvs | lvs.txt | Logical volume information(when using SUSE |

| Executed command | Files | Explanation |
|---|---|---|
| | | LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10) |
| /sbin/parted -s *HDLM-device* print | parted-s.txt | HDLM device partition information (when using Red Hat Enterprise Linux 3 or Red Hat Enterprise Linux 4 on IPF) |
| /usr/sbin/parted -s *HDLM-device* print | parted-s.txt | HDLM device partition information (when using SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10 on IPF) |
| /sbin/pvs | pvs.txt | Physical volume information (when using SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10) |
| /sbin/pvscan | pvscan.txt | Physical volume information |
| /sbin/runlevel | runlevel.txt | Run level information |
| /sbin/sysctl -A | sysctl.txt | Kernel parameter information |
| /sbin/vgdisplay -v | vgdisplay-v.txt | Volume group information (when using Red Hat Enterprise Linux 3, SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10) |
| /sbin/vgs | vgs.txt | Volume group information (when using SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10) |
| ulimit -a | ulimit-a.txt | Limit values of the available system resources (data segment, stack segment, and file descriptor) for the process |
| /usr/bin/env | env.txt | Environment variables information |
| /usr/bin/free | free.txt | Memory information |
| /usr/bin/getconf PAGESIZE | getconfPAGESIZE.txt | Memory page size information |
| /usr/bin/ident dlnkmgr dlmmgr DLMgetras libdlm.so libhdlmhcc-* libhdlmhccmp-* dlmcfgmgr sddlmadrv.o sddlmfdrv.o | whatlist.txt | Program build number (dlnkmgr, dlmmgr, DLMgetras, libdlm.so, libhdlmhcc-* libhdlmhccmp-*, dlmcfgmgr, sddlmadrv.o, and sddlmfdrv.o) (when using Red Hat Enterprise LINUX 3) |
| /usr/bin/ident dlnkmgr dlmmgr DLMgetras libdlm.so libhdlmhcc-* libhdlmhccmp-*dlmcfgmgr sddlmadrv.ko sddlmfdrv.ko | whatlist.txt | Program build number (dlnkmgr, dlmmgr, DLMgetras, libdlm.so libhdlmhcc-*, libhdlmhccmp-*dlmcfgmgr, sddlmadrv.ko, and sddlmfdrv.ko) (when using Red Hat Enterprise Linux 4, SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10) |
| /usr/bin/iostat | iostat.txt | Device statistics |
| /usr/bin/iostat -p | iostat-p.txt | Block device statistics information (when using Red Hat Enterprise Linux 4, SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10) |
| /usr/bin/lsdev | lsdev.txt | List of drivers (when using Red Hat Enterprise Linux 3, SUSE LINUX Enterprise Server 9, or SUSE LINUX Enterprise Server 10) |

| Executed command | Files | Explanation |
|---|---|---|
| /usr/bin/procinfo | procinfo.txt | Machine information (when using Red Hat Enterprise Linux 3, SUSE LINUX Enterprise Server 9, or SUSE LINUX Enterprise Server 10) |
| /usr/bin/raw -qa | raw-qa.txt | Character-type device information (when using Red Hat Enterprise Linux 3 or Red Hat Enterprise Linux 4) |
| /usr/sbin/raw -qa | raw-qa.txt | Character-type device information (when using SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10) |
| /usr/bin/uptime | uptime.txt | System's operating time information |
| /usr/bin/udevinfo -d | udevinfo-d.txt | The udev database information (when using Red Hat Enterprise Linux 4, SUSE LINUX Enterprise Server 9, or SUSE LINUX Enterprise Server 10) |
| /usr/sbin/dmidecode | dmidecode.txt | DMI information (when using Red Hat Enterprise LINUX 3) |
| /usr/sbin/lsscsi | lsscsi.txt | SCSI device list information (when using SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10) |
| /usr/sbin/lvdiskscan | lvmdiskscan.txt | lvm disk information (when using Red Hat Enterprise LINUX 4) |
| /usr/sbin/lvs | lvs.txt | Logical volume information (when using Red Hat Enterprise LINUX 4) |
| /usr/sbin/pvs | pvs.txt | Physical volume information (when using Red Hat Enterprise LINUX 4) |
| /usr/sbin/sestatus -v | sestatus-v.txt | SELinux information (when using Red Hat Enterprise LINUX 4) |
| /usr/sbin/vgdisplay -v | vgdisplay-v.txt | Volume group information (when using Red Hat Enterprise LINUX 4) |
| /usr/sbin/vgs -v | vgs.txt | Volume group information (when using Red Hat Enterprise LINUX 4) |

## 7.3 dlmcfgmgr (HDLM-Configuration Definition Utility)

For management by HDLM, the HDLM-configuration definition utility obtains the SCSI device information stored by Linux by detecting the information on the path between the host and storage subsystem to register in HDLM, and performs creation, update, or deletion of the HDLM devices. Also, this utility sets an HDLM device as a management target or excludes it from being managed.

When an existing device is updated during execution of this utility, the definition of the HDLM device is inherited.

Whenever you make any changes to the storage subsystem configuration (adding or deleting an LU or path), LU partitions, or hardware configuration on the host, restart the host or execute the `dlmcfgmgr` utility to re-configure the HDLM devices.

The operation log is obtained while the `dlmcfgmgr` utility is executed.

### 7.3.1 Format

```
/sbin/dlmcfgmgr
        [-s]
        {-r
        |-o {logical-device-file-name-of-the-HDLM-device ...|all}
        |-i {logical-device-file-name-of-the-HDLM-device ...|all}
        |-v [-udev]
        |-u {logical-device-file-name-of-the-HDLM-device ...|all}}
```

### 7.3.2 Parameters

The following table lists and describes each parameter name and its functionality:

**Table 7.5 Functionality of the Parameters of the dlmcfgmgr Utility**

| Parameter | Functionality |
|---|---|
| `-s` | Executes the `dlmcfgmgr` utility without displaying the confirmation message. |
| `-r` | Registers the HDLM device path while the host is running. |
| `-o` | Excludes a management-target HDLM device from being managed. |
| `-i` | Returns an HDLM device, which was excluded from being managed by HDLM, to being a management-target HDLM device. |
| `-v` `[-udev]` | Displays the management status and configuration information of all the HDLM devices recognized by HDLM. |
| `-u` | Checks the status of all the paths for the specified HDLM device. |

`-s`

Specify this parameter when you do not want to display the confirmation messages during execution of the `dlmcfgmgr` utility. When you execute this utility with this parameter specified, the confirmation message for each HDLM device is not displayed.

Use this parameter when you want to skip a response to the confirmation message (to eliminate the manual intervention). For example, use this parameter when executing a command using a shell script or batch file.

When you specify the `-v` parameter, the message confirming that the utility is executed is not displayed even if you omit the `-s` parameter.

`-r`

Specify this parameter when you register, in HDLM, the path that HDLM has not recognized. When the `dlmcfgmgr` utility is executed, the logical device file of the HDLM device is created in the `/dev` directory. This enables a user to use an LDEV of the storage subsystem as an HDLM device.

This parameter is used to create a new HDLM device definition; for example, when changing the configuration of the storage subsystem (for example, add or delete an LU), or the host-side hardware configuration.

The path definition information registered in HDLM is inherited even though the path is disconnected when the host starts.

If the SCSI device of the new path that HDLM detected is already registered (that is, when the path between the host and an LU of the storage subsystem already exists and a new path is added), the same HDLM device mapped with the path already registered is allocated. If the SCSI device of the new path that HDLM detected is not yet registered (that is, when you define a new path from the host to an LU of the storage subsystem), the HDLM device with the lowest available letter is allocated to the path. Examples of allocating a new HDLM device are shown in Table 7.6.

**Table 7.6 Example: Allocation of New HDLM Devices**

| Status before executing the dlmcfgmgr utility (**`dlmcfgmgr -r`**) | Status after executing the dlmcfgmgr utility (**`dlmcfgmgr -r`**) |
|---|---|
| none | /dev/sddlmaa |
| /dev/sddlmaa | /dev/sddlmaa<br>/dev/sddlmab# |
| /dev/sddlmaa<br>/dev/sddlmab | /dev/sddlmaa<br>/dev/sddlmab<br>/dev/sddlmac# |
| /dev/sddlmaa<br>/dev/sddlmac | /dev/sddlmaa<br>/dev/sddlmab#<br>/dev/sddlmac |
| /dev/sddlmab<br>/dev/sddlmac | /dev/sddlmaa#<br>/dev/sddlmab<br>/dev/sddlmac |

#

A newly assigned logical device file name of the HDLM device

Every time the host starts, the `dlmcfgmgr` utility (`dlmcfgmgr -r`) is automatically executed. This utility can be executed a number of times if necessary after starting the host. When this utility is executed, this utility checks the current HDLM settings. If a new path is detected, this utility registers the path dynamically so that the path can be used. The path status for the newly detected path or the existing path other than the Offline(C) path is changed to the current status. In this case, the path health checking and the automatic failback are executed at the same time.

If a new LU was recognized when the number of HDLM devices created reached the upper limit, HDLM displays the message KAPL10357-E and does not add an HDLM device. In this case, execute the `dlmcfgmgr` utility with the `-u` parameter to delete the unused HDLM device. This releases a usable name and you can create a logical device file for the HDLM device for a new LU.

Even if you execute the `dlmcfgmgr` utility (`dlmcfgmgr -r`) to add a path to an HDLM device that is not managed by HDLM, that HDLM device will not be managed by HDLM.

`-o {`*logical-device-file-name-of-the-HDLM-device* `...|all}`

Use this parameter to exclude the management-target HDLM device from being managed.

You can specify one or more logical device file names of an HDLM device (`/dev/sddlm[`*a-p*`][`*a-p*`]`) as a parameter value. For `[`*a-p*`]`, specify an alphabetic character from `a` to `p`.

If you want to exclude all the HDLM devices from being managed, specify `all`. Specifying `all` together with a logical device file name of the HDLM device causes an error.

Specifying the `-o` parameter without a parameter value also causes an error. If you specify a logical device file name of an HDLM device that does not exist, non-HDLM device, or HDLM device that is already defined as a non-HDLM-managed device, the specification is ignored.

The `dlmcfgmgr` utility with this parameter can be executed a number of times while Linux is running. The specified HDLM device can be excluded from being managed immediately after execution of the utility, except when the specified HDLM device is in use. If the specified HDLM device is in use, an error occurs.

You can use this parameter to exclude the HDLM devices that are no longer required to be managed by HDLM. The excluded HDLM devices are not displayed with the `dlnkmgr` command's `view` operation and the HDLM device files for these devices are not created. Therefore, these devices become unavailable to the user.

`-i {`*logical-device-file-name-of-the-HDLM-device* `...|all}`

Use this parameter to reset the HDLM device that has been excluded from being managed to an HDLM management target.

You can specify one or more logical device file names of an HDLM device (`/dev/sddlm[`*a-p*`][`*a-p*`]`) as a parameter value. For `[`*a-p*`]`, specify an alphabetical character from `a` to `p`.

If all the HDLM devices are the target, specify `all`. Specifying `all` together with a logical device file name of an HDLM device causes an error.

Specifying the `-i` parameter without a parameter value causes an error. If you specify a logical device file name of an HDLM device that does not exist, non-HDLM device, or HDLM device that is already defined as a management target, the specification is ignored.

The `dlmcfgmgr` utility with this parameter can be executed a number of times while Linux is running. The specified HDLM device becomes available immediately after execution of the utility.

You can use this parameter to reset an HDLM device that must be managed again to a management target. The HDLM devices reset to a management target are displayed with the `dlnkmgr` command's `view` operation, and the HDLM device files for the devices are created. Therefore, these devices become available to the user again.

`-v [-udev]`

Use this parameter to display the management status and configuration information of all the HDLM devices recognized by HDLM.

Specify this parameter when you want to get the current status of management, configuration, or other information for the HDLM device.

With this parameter, the message confirming that the utility is executed is not displayed at execution of the utility even if you omit the `-s` parameter. Unlike the `-o`, `-i`, or `-u` parameter, the utility with this parameter specified can operate during the I/O processing to and from the HDLM device.

If you want to also check the `udev` name assigned to the HDLM device, specify the `-udev` parameter. The `-udev` parameter can be specified only when Red Hat Enterprise Linux AS/ES 4, SUSE LINUX Enterprise Server 9, or SUSE LINUX Enterprise Server 10 is used.

`-u {`*logical-device-file-name-of-the-HDLM-device* `...|all}`

Use this parameter to check the status of all the paths associated with an HDLM device.

You can specify one or more logical device file names of an HDLM device (`/dev/sddlm[`*a-p*`][`*a-p*`]`) as a value of the `-u` parameter. For `[`*a-p*`]`, specify an alphabetical character from `a` to `p`. Specifying the `-u` parameter without the parameter value cases an error.

If all the HDLM devices are the target, specify `all`. Specifying `all` together with a logical device file name of an HDLM device causes an error. If you specify the logical device file name for an HDLM device that does not exist or a non-HDLM device, the specification is ignored.

The `dlmcfgmgr` utility with this parameter can be executed a number of times while Linux is running. The re-configured HDLM device becomes available immediately after execution of the utility.

You can use this parameter to delete the information for an unavailable path (in the disconnected state) associated with an HDLM device at host boot time and to unregister this information from HDLM for re-configuration of the HDLM device paths.

In general, the path information for the HDLM device is added at HDLM startup time once it is recognized. Even if the recognized HDLM device is not available because this path is disconnected or the LU is deallocated, the path information for the unavailable HDLM device will not be deleted automatically. Because the unused HDLM device path information is preserved, newly defined LUs and added paths cannot be used in some cases. This parameter is useful to eliminate such a problem.

If you specify this parameter with `all`, an HDLM device, which was simply disconnected at host boot time but is currently being used, may be assumed to be deleted. Before executing the `dlmcfgmgr` utility with this parameter set to a value of `all`, check the status of the HDLM device to be deleted by executing this utility with the `-v` parameter.

## Examples

Example 1

To register the path that HDLM has not identified, by confirming the execution of the utility:

In response to the confirmation message, enter `y` to execute the utility and `n` not to execute the utility.

```
# /sbin/dlmcfgmgr -r
KAPL10339-I This operation will change the configuration of HDLM devices. Do
you want to continue? [y/n] :
```

Example 2

To exclude the HDLM device `sddlmaa` from being managed without confirming the execution of the utility:

```
# /sbin/dlmcfgmgr -s -o /dev/sddlmaa
KAPL10341-I The HDLM device configurations have been changed.
KAPL10302-I /sbin/dlmcfgmgr completed normally.
```

Example 3

To delete the information for all the unavailable paths and unregister the path information from HDLM after confirming the execution of the utility:

In response to the confirmation message, enter `y` to execute the utility and `n` not to execute the utility.

```
# /sbin/dlmcfgmgr -u all
KAPL10339-I This operation will change the configuration of HDLM devices. Do
you want to continue? [y/n] : y
KAPL10341-I The HDLM device configurations have been changed.
KAPL10302-I /sbin/dlmcfgmgr completed normally.
```

Example 4

To display the management status and configuration information of an HDLM device:

```
# /sbin/dlmcfgmgr -v
HDevName         Management      Device     Host    Channel Target  Lun
/dev/sddlmaa    configured      /dev/sda        2       0       0       0
                                /dev/sdb        2       0       0       1
/dev/sddlmab    unconfigured    /dev/sdc        2       0       1       0
                                -               2       0       1       1
KAPL10302-I /sbin/dlmcfgmgr completed normally.
#
```

Table 7.7 lists and describes the items displayed.

**Table 7.7 Information for the Management Status and Configuration Information of an HDLM Device**

| Item | Description |
|---|---|
| `HDevName` | Indicates the logical device file name of the HDLM device, using an absolute path. This is indicated only for the first path for the relevant LU. |
| `Management` | Indicates the current management status of the HDLM device.<br>`configured`: indicates that the HDLM device is to be managed.<br>`unconfigured`: indicates that the HDLM device is excluded from being managed. |
| `Device` | Indicates the paths associated with the HDLM device.<br>The logical device file of the SCSI device is indicated.<br>The path whose allocation to an LU is canceled or path that was disconnected at boot time is temporarily registered into HDLM and indicated by a hyphen (–). |
| `Host` | Indicates the host ID (host port number). |
| `Channel` | Indicates the channel number (bus number). |
| `Target` | Indicates the target ID (target ID). |
| `Lun` | Indicates the Lun (host LU number). |
| `Udev` | Indicates the udev name.<br>The name of a device whose name has not been customized by using the `udev` function is indicated by a hyphen (–). |

## Reference information

If you do not specify a parameter, the syntax of the `dlmcfgmgr` utility is displayed.

Example when Red Hat Enterprise Linux AS/ES 3 is used:

```
# /sbin/dlmcfgmgr
KAPL10319-W usage: /sbin/dlmcfgmgr [-s] { -r | -o {special-file-name ... | all}
                                | -i {special-file-name ... | all} | -v
                                | -u {special-file-name ... | all}}
#
```

Example when Red Hat Enterprise Linux AS/ES 4, SUSE LINUX Enterprise Server 9, or SUSE LINUX Enterprise Server 10 is used:

```
# /sbin/dlmcfgmgr
KAPL10319-W usage: /sbin/dlmcfgmgr [-s] { -r | -o {special-file-name ... | all}
                                   | -i {special-file-name ... | all} | -v [-udev]
                                   | -u {special-file-name ... | all}}
#
```

***Notes:*** **(common to OSs)**

- When an LDEV is allocated to a path and then another LDEV is allocated to the same path (host port number, bus number, target ID, and host LU number), the HDLM device name for the new LDEV may sometimes be the same as that for the old LDEV. For example, when an LDEV is deallocated while a host is running, another LDEV can be allocated to the same path. In this case, if an application preserves the settings to use that HDLM device name, a program such as a higher-level program product of HDLM accesses that HDLM device without recognizing that the LDEV was changed. This may result in data corruption. Check that the HDLM device name you previously used is nowhere to be found before deallocating the LDEV.
- You cannot execute two instances of the `dlmcfgmgr` utility at the same time.
- A logical device file name of the HDLM device file that can be specified with the `-i`, `-o`, or `-u` parameter of the `dlmcfgmgr` utility can consist of up to 4,095 characters. If a name is more than 4,096 characters long, the error message KAPL10358-E is displayed.
- If an LU is allocated and another LU is reallocated using the same host LU number, HDLM does not handle them as different LUs. Therefore, restart the host whenever a configuration change is made.
- When two or more HDLM devices are specified with the `-s` parameter, the confirmation message is not displayed each time the processing for each HDLM device is executed. For example, when this utility executes an operation specified with the `-o` parameter for `/dev/sddlmaa`, `/dev/sddlmab`, and `/dev/sddlmac`, the confirmation message KAPL10339-I is not displayed for the processing of each device. This message is displayed once at the beginning of the operation. If you enter `y`, the processing for all the specified HDLM devices (`/dev/sddlmaa`, `/dev/sddlmab`, and `/dev/sddlmac`) is executed continuously.
- After changing partition configuration directly from a SCSI device for LUs corresponding to the HDLM devices not to be managed, if the HDLM device is returned to management target, execute the following operations:

  a) `Execute the dlmcfgmgr` utility `with the -i parameter specified.`

  `Execute the dlmcfgmgr` utility `with -r parameter specified.`
- If you execute the `dlmcfgmgr` utility with the `-i` parameter specified, the HDLM device that is not a management target can be changed to a managed status. At this time, if an error occurs anywhere on the paths to an HDLM device and you change that HDLM device to a managed status, the message KAPL08026-E might be output to the syslog. However, HDLM operations are not affected.
- When the `dlmcfgmgr` utility is executed with `-r` parameter specified, the I/O and I/O error counts for existing paths are reset to 0.

- If one or more HDLM devices are mounted when you execute the `dlmcfgmgr` utility with the `-r` parameter specified, an error message (KAPL05023-E) is output to the syslog. However, HDLM operations are not affected.
- If an HDLM device that is managed on the active host is changed to be managed on the standby node, and then the same HDLM device is excluded from management on the standby node then, during the cluster creation, an error will occur when the active node is failed over. When the `dlmcfgmgr` utility is executed with the `-o` parameter specified to exclude an HDLM device from management on the standby node, make sure that you exclude the same HDLM device from management on the active node.
- When an HDLM device is bound to a raw device, do not attempt to exclude that HDLM device from management by executing the `dlmcfgmgr` utility with the `-o` parameter specified. If the HDLM device is not being managed, an error will occur when the raw device is accessed.

***Notes:*** **(when using SUSE LINUX Enterprise Server 9, SUSE LINUX Enterprise Server 10 or Red Hat Enterprise Linux AS4/ES4)**

- Do not change the set values of the `sysfs` interface that treats the SCSI devices after creating an HDLM device by executing the `dlmcfgmgr` utility.
- After you create an HDLM device by executing the `dlmcfgmgr` utility and unload the HBA driver, perform the following operations. If you load the HBA driver without performing the operations, the number of HDLM paths that cannot be used might increase.

  a) Update the HDLM device with information about unavailable SCSI devices by executing the `dlmcfgmgr` utility with the `-r` parameter specified.

  Delete the HDLM device by executing the `dlmcfgmgr` utility with `all` specified in the `-u` parameter.
- If you execute the `dlmcfgmgr` utility with the `-r` parameter after dynamically deleting an SCSI device, the HDLM device will become unavailable. If you then add the SCSI device dynamically and execute the `dlmcfgmgr` utility with the `-r` parameter specified, the HDLM partition information will be lost from `/proc/partitions`. To restore the information, follow these steps:

  a) Delete the SCSI device you just added.

  Dynamically add the SCSI device once again.

  Execute the `online` operation to place the device path online.

  Execute the `dlmcfgmgr` utility with the `-r` parameter specified.
- If you execute the `dlmcfgmgr` utility with the `-r` parameter specified while all paths connecting to the HDLM device are disconnected, the HDLM partition information will be lost from `/proc/partitions`. To restore the HDLM partition information, follow these steps:

  a) Connect the path and place the device online.

  Execute `blockdev --rereadpt` *HDLM-device* to update the partition information.

***Notes:*** **(when using SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10)**

- If you execute the `DLMgetras` utility for collecting HDLM error information while running SUSE LINUX Enterprise Server 9 (IPF) or SUSE LINUX Enterprise Server 10 (IPF), an error message (KAPL10034-E) will be output to `getras.log` at copy processing of `/proc/sys/fs/binfmt_misc/register`. However, HDLM operations are not affected.

***Notes:*** **(when using Red Hat Enterprise Linux AS4/ES4)**

- if you perform either of the operations below, the message KAPL05023-E will be output to the syslog. However, HDLM operations are not affected.
  - The `dlmcfgmgr` utility is executed with the `-r` parameter specified while one or more unpartitioned SCSI devices exist
  - The `dlmcfgmgr` utility is executed with the `-i` parameter specified for an unpartitioned HDLM device

## 7.4 dlmlfk (Utility for Supporting LifeKeeper)

This utility enables or disables the use of HDLM devices from LifeKeeper.

### 7.4.1 Format

```
/opt/DynamicLinkManager/bin/dlmlfk { -c | -u }
```

### 7.4.2 Parameters

`-c`

Enables the use of HDLM devices from LifeKeeper.

`-u`

Disables the use of HDLM devices from LifeKeeper.

***Notes:***

- Stop LifeKeeper before you execute the `dlmlfk` utility.
- The `dlmlfk` utility cannot be executed in an environment where LifeKeeper has not been installed.
- When starting or stopping LifeKeeper, you must not execute the `dlmlfk` utility for supporting LifeKeeper from the command line or a custom script. The following shows examples you are not allowed to execute:
  - The command `dlmlfk -c(-u)` must not be executed at the same time as LifeKeeper's start command `lkstart` or stop command `lkstop`.
  - The code `dlmlfk -c(-u)` must not be included in any script executed at host startup or shutdown.

## 7.5 dlmmkinitrd (Utility for Supporting a Boot Disk)

This utility creates an initial RAM disk image file for using an HDLM device as a boot disk.

### 7.5.1 Format

```
/opt/DynamicLinkManager/bin/dlmmkinitrd
        [-v]
        [-f]
        [-fstab fstab-name]
        initrd-image
        kernel-version
```

### 7.5.2 Parameters

`-v`

Specify this parameter to output a log to the console during execution of the `dlmmkinitrd` utility.

`-f`

Specify this parameter to create an initial RAM disk image file that has the same name as an existing file. If you execute the `dlmmkinitrd` utility with this option specified, the created initial RAM disk image file will be overwritten.

`-fstab` *fstab-name*

Specify the `fstab` file used to determine the root device file system. The default is `/etc/fstab`.

*initrd-image*

Specify the full path name or relative path name of the initial RAM disk image file to be created by using the `dlmmkinitrd` utility.

*kernel-version*

Specify the kernel used for creating an initial RAM disk image.

## 7.6 dlmpr (Utility for Clearing HDLM Persistent Reservation)

When multiple hosts share an LU, if the persistent reservation of the LU is not canceled for some reason, this utility clears the Reservation Key to cancel the persistent reservation.

### 7.6.1 Format

`/opt/DynamicLinkManager/bin/dlmpr {{-k | -c} [sdn] [sdn] ... [-a] | -h}`

### 7.6.2 Parameters

`-k`

Specify this parameter to display the Reservation Key. The following explains the items displayed when the `-k` parameter is specified:

Reservation Key

An asterisk (`*`) is displayed at the end of the Reservation Key for a Reservation Key of another host.

If the Reservation Key is not set, `[0x0000000000000000]` is displayed.

Regist Key

The registered Keys are displayed.

Key Count

The number of registered Keys is displayed.

`-c`

Specify this parameter to clear the Reservation Key.

`sd`*n*

Specify the SCSI device (`sd`*n*) for which you want to clear the Reservation Key. You can specify more than one SCSI device.

If you omit this parameter, the utility clears the Reservation Keys for all SCSI devices.

`-a`

When multiple SCSI devices (`sd`*n*) are specified, even if an error occurs during processing, the processing continues for all SCSI devices.

`-h`

Displays the format of the utility for clearing HDLM persistent reservation.

***Note:***

- The `dlmpr` utility can only be used on a kernel that supports LifeKeeper.

The `dlmpr` utility is not installed for a kernel that does not support LifeKeeper. For details about kernels that support LifeKeeper, see section Table 3.4.

- Before executing the `dlmpr` utility, make sure that the LU for which persistent reservation will be cleared is not being accessed.

  If you perform the following operation while both of the following conditions exist, an error occurs when Reservation-key is read, and `0x????????????????` is displayed.

  Conditions

  - Multiple paths are configured to the HDLM device.

  - A persistent reservation has been set by the cluster software.

  Operation

  1. Clear persistent reservation, specifying the SCSI device.

  2. Execute the `dlmpr` utility with the `-k` parameter specified to display the PersistentReserve information for a SCSI device associated with a different path.

If the above display error occurs, re-execute the `dlmpr` utility with the `-k` parameter specified.

## Example

To check the Reservation Keys, and then clear the Reservation Keys other than those for the local host:

1. Execute the `dlmpr` utility to display the Reservation Keys for `sda`, `sdb`, `sdc`, `sdd`, `sde`, and `sdf`.

```
# /opt/DynamicLinkManager/bin/dlmpr -k sda sdb sdc sdd sde sdf
self Reservation Key : [0xaaaaaaaaaaaaaaaa]
sda  Reservation Key : [0xaaaaaaaaaaaaaaaa]
          Regist Key : [0xaaaaaaaaaaaaaaaa], Key Count : 1
sdb  Reservation Key : [0xbbbbbbbbbbbbbbbb]*
          Regist Key : [0xaaaaaaaaaaaaaaaa], Key Count : 2
          Regist Key : [0xbbbbbbbbbbbbbbbb], Key Count : 2
sdc  Reservation Key : [0xbbbbbbbbbbbbbbbb]*
          Regist Key : [0xbbbbbbbbbbbbbbbb], Key Count : 4
sdd  Reservation Key : [0xaaaaaaaaaaaaaaaa]
          Regist Key : [0xaaaaaaaaaaaaaaaa], Key Count : 4
sde  Reservation Key : [0x0000000000000000]
sdf  Reservation Key : [0x0000000000000000]
          Regist Key : [0xaaaaaaaaaaaaaaaa], Key Count : 1
          Regist Key : [0xbbbbbbbbbbbbbbbb], Key Count : 1
```

2. Execute the `dlmpr` utility to clear the Reservation Keys for other hosts (marked by an asterisk (*)).

```
# /opt/DynamicLinkManager/bin/dlmpr -c sdb sdc
```

3. The confirmation message appears. Enter `y` to clear. Otherwise, enter `n`.

```
KAPL10641-I Reservation Key will now be cleared. Is this OK? [y/n]:y
KAPL10642-I Reservation Key of sdb was cleared.
KAPL10642-I Reservation Key of sdc was cleared.
```

## 7.7 dlmsetopt (Utility for Setting HDLM Driver Option)

This utility changes the settings for the HDLM filter driver. Restart the host machine to apply the changed settings. Note that you do not need to restart the host machine every time you execute the `dlmsetopt` utility. If you execute the `dlmsetopt` utility more than once, you only need to restart the host once after you have finished all of the executions.

### 7.7.1 Format

`/opt/DynamicLinkManager/bin/dlmsetopt {-r` *retrycount* `| -inqt` *InquiryTimeout* `| -inqr` *InquiryRetry* `| -h}`

### 7.7.2 Parameters

`-r`

Specify the number of retries when memory is allocated in the filter driver. This memory is allocated for processing I/O requests from users. If memory allocation fails, the system retries at intervals of 0.05 seconds. If retrying for a long time results in I/O not being executed, set an appropriate retry count. If that retry count is exceeded and memory allocation fails, the KAPL05708-E message is output to syslog and to the HDLM manager log file.

*retrycount*

Specify a numeric value from -1 to 2,147,483,646.

-1 : The system retries until memory is allocated. This is the initial value of the system.

0 : The system does not retry.

1 to 2,147,483,646: The system retries the specified number of times.

If the following values are specified, the message KAPL12559-E is displayed and then the processing terminates:

- Non-numeric value
- Value less than -1
- Value more than 2,147,483,646

`-inqt`

Specify, in seconds, the timeout value of the `SCSI INQUIRY` command. In HDLM, the default timeout value is 30 seconds. Note that the timeout value specified by using the -inqt parameter only applies to the `SCSI INQUIRY` command executed from HDLM. This value has no effect on the `SCSI INQUIRY` command executed from another application.

*InquiryTimeout*

Specify a numeric value of -1 or in the range from 1 to 3,600.

-1: This sets the timeout value to 30 seconds, which is HDLM's default value.

1 to 3,600: Timeout value (in seconds)

If the following values are specified, the message KAPL12553-W is displayed and then the processing ends:

- 0
- Value of -1 or less
- Value more than 3,600
- Non-numeric value.

`-inqr`

Specify the number of retries of the `SCSI INQUIRY` command. In HDLM, the default number of retries is 1. Note that the number of retries specified by using the `-inqr` parameter only applies to the `SCSI INQUIRY` command executed from HDLM. This value has no effect on the `SCSI INQUIRY` command executed from another application.

*InquiryRetry*

Specify a numeric value in the range from -1 to 2,147,483,646.

-1: This sets the number of retries to 1, which is HDLM's default value.

0: The system does not retry.

1 to 2,147,483,646: The system retries the specified number of times.

If the following values are specified, the KAPL12553-W message is displayed and then processing ends:

- Value of -1 or less
- Value more than 2,147,483,646
- Non-numeric value

`-h`

Displays the format of the `dlmsetopt` utility for setting the HDLM driver option.

## Example

Example 1

In this example, the utility sets the retry count to 100000:

```
# /opt/DynamicLinkManager/bin/dlmsetopt -r 100000
KAPL12554-I The utility for setting HDLM driver option has started.
KAPL12555-I The utility for setting HDLM driver option completed normally.
KAPL12558-I Please restart the computer so that the option settings take effect.
```

Example 2

In this example, the utility displays Help:

```
# /opt/DynamicLinkManager/bin/dlmsetopt -h
KAPL12554-I The utility for setting HDLM driver option has started.
Usage: dlmsetopt {-r retrycount | -inqt InquiryTimeout | -inqr InquiryRetry | -h}
       retrycount:     -1           = Infinite
                        0            = No Retry
                       1-2147483646 = Retry Count
       InquiryTimeout: -1           = Default Timeout(30(s))
                       1-3600       = Inquiry Timeout Value
       InquiryRetry:    -1           = Default Retry(1)
                        0            = No Retry
                       1-2147483646 = Inquiry Retry Count
KAPL12555-I The utility for setting HDLM driver option completed normally.
```

***Notes:***

- The contents specified in the `dlmsetopt` utility are written to the `options sddlmfdrv` line of the conf files. Do not attempt to edit this line using an editor such as vi. For details on the conf files, see Table 7.8.

**Table 7.8 conf files used by the dlmsetopt utility**

| OS | File Name |
|---|---|
| Red Hat Enterprise Linux AS4 | `/etc/modprobe.conf` |
| Red Hat Enterprise Linux ES4 | |
| SUSE LINUX Enterprise Server 9 | |
| SUSE LINUX Enterprise Server 10 | |
| Red Hat Enterprise Linux AS3 | `/etc/modules.conf` |
| Red Hat Enterprise Linux ES3 | |

- If you use the `dlmsetopt` utility in an environment where an HDLM device is used for the boot disk, after you change the filter driver settings you must re-create the initial RAM disk image file by executing the `dlmmkinitrd` utility. If you change the name of the initial RAM disk image file, change the settings file for the boot loader so that the re-created initial RAM disk image file is used during boot-up. For details on this procedure, see Table 7.9.

**Table 7.9 Setting file names for boot loaders**

| Boot loader | File Name |
|---|---|
| LILO | `/etc/lilo.conf` |
| GRUB | `/etc/grub.conf` |
| ELILO | `/etc/elilo.conf` |

## 7.8 installgetras (Utility for Collecting HDLM Installation Error Information)

This utility collects the error log information needed to analyze errors that occurred during HDLM installation. The collected information is compiled in a file and output to a specified directory. The following file is output:

- `installgetras.tar.gz`

This is a compressed file that contains information related to HDLM installation processing.

For details on the information included in this output file, see section 7.8.3.

### 7.8.1 Format

*mount-point-of-HDLM-installation-media*/`installgetras` *directory-to-which-collected-information-is-output*

### 7.8.2 Parameters

*directory-to-which-collected-information-is-output*

Specify the output directory for the information collected by `installgetras`. The collected information is compiled in the `installgetras.tar.gz` file and output to the specified directory.

You cannot specify the root directory (/) for the directory to which collected information is output.

### 7.8.3 Error Information to be Collected

The following shows the information collected when HDLM is being installed:

- `installgetras.tar.gz`

  This file contains information related to errors that occurred during HDLM installation. Table 7.10 describes the detailed information contained in `installgetras.tar.gz`.

**Table 7.10 Information Contained in installgetras.tar.g**

| Output directory# | File | Description |
|---|---|---|
| *output-directory-of-collected-information-specified-whe n-executing-the- -installgetras-utility* | `installgetras.log` | Log file of the `installgetras` utility |
| `/var/log` | `messages` | OS syslog file |
| `/var/tmp/hdlminstlog` | All subdirectories and files | Log information created during HDLM installation |

\#

This directory is created in *directory-to-which-collected-information-is-output* specified when the `installgetras` utility was executed.

# 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 8.17 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

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.<br>01: HDLM command<br>03: HDLM API<br>04: HDLM manager<br>05: HDLM driver (filter component)<br>06: HDLM alert driver<br>07: HDLM driver (core logic component)<br>08: HDLM management target<br>09: HDLM installation program<br>10: The following utilities:<br>▪ DLMgetras utility (KAPL100*nn*)<br>▪ dlmcfgmgr utility (KAPL103*nn*)<br>▪ dlmpr utility (KAPL106*nn*)<br>11: HDLM remote access interface<br>12: dlmlfk utility (KAPL120*nn*), dlmmkinitrd utility (KAPL123*nn*), dlmsetopt utility (KAPL125*nn*)<br>13: installgetras utility (KAPL134*nn*) |
| *nnn* | Message serial number for the module |
| *l* | Message level<br>C: Critical<br>E: Error<br>W: Warning<br>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 a message contains two or more variables, they are displayed as *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<br>The HDLM command completed successfully.<br>*aa...aa*: `clear`, `help`, `offline`, `online`, `set`, `view`<br>*bb...bb*: the year of grace/month/day hour:minute:second<br>Action<br>None. |
| KAPL01002-I | The HDLM command started. Operation name = *aa...aa* | Details<br>The HDLM command was executed.<br>*aa...aa*: `clear`, `offline`, `online`, `set`, `view`<br>Action<br>None. |
| KAPL01003-W | No operation name is specified. | Details<br>The operation name is missing.<br>Action<br>Specify the operation name, and then retry. |
| KAPL01004-W | The operation name is invalid. Operation name = *aa...aa* | Details<br>*aa...aa*: Specified operation name<br>Action<br>Execute the `help` operation of the HDLM command (`dlnkmgr`) to check the operation name, and then retry. For details on the `help` operation, see section 6.3. |
| KAPL01005-W | A parameter is invalid. Operation name = *aa...aa*, parameter = *bb...bb* | Details<br>*aa...aa*: `clear`, `set`, `online`, `offline`, `view`<br>*bb...bb*: Specified parameter<br>Action<br>Execute `help` *operation-name* of the HDLM command (`dlnkmgr`) to check the parameter, and then retry. For details on the `help` operation, see section 6.3. |
| KAPL01006-W | A necessary parameter is not specified. Operation name = *aa...aa* | Details<br>The specified operation does not contain the necessary parameter.<br>*aa...aa*: `clear`, `set`, `offline`, `view`<br>Action<br>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 section 6.3. |
| KAPL01007-W | A duplicate parameter is specified. Operation name = *aa...aa*, parameter = *bb...bb* | Details<br>*aa...aa*: `clear`, `offline`, `online`, `set`, `view`<br>*bb...bb*: Duplicate parameter<br>Action<br>Delete the duplicate parameter, and then retry. |

| | | |
|---|---|---|
| KAPL01008-W | A necessary parameter value is not specified. Operation name = *aa...aa*, parameter = *bb...bb* | Details<br>*aa...aa*: `offline`, `online`, `set`, `view`<br>*bb...bb*: Parameter name<br>Action<br>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<br>*aa...aa*: `offline`, `online`, `set`, `view`<br>*bb...bb*: Parameter name<br>*cc...cc*: Specified parameter value<br>*dd...dd*: Specifiable parameter value range<br>Action<br>Specify the correct value for the parameter, and then retry. |
| KAPL01012-E | Could not connect the HDLM manager. Operation name = *aa...aa* | Details<br>In the `view -sys -sfunc` operation, information must be collected from the HDLM manager but the manager cannot be accessed.<br>*aa...aa*: `view`<br>Action<br>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 section 6.7. |
| KAPL01013-E | An error occurred in internal processing of the HDLM command. Operation name = *aa...aa* details = *bb...bb* | Details<br>An error whose cause does not seem to be a user operation occurred during command processing.<br>*aa...aa*: `clear`, `offline`, `online`, `set`, `view`<br>*bb...bb*: Function name and contents of error<br>Action<br>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 section 7.2. |
| KAPL01014-W | No authority to execute the HDLM command. Operation name = *aa...aa* | Details<br>You are not authorized to execute the HDLM command as the administrator.<br>*aa...aa*: `clear`, `offline`, `online`, `set`, `view`<br>Action<br>Execute the command as a user with root privileges. |
| KAPL01015-W | The target HBA was not found. Operation name = *aa...aa* | Details<br>The path having the host port number and path number specified in the `-hba` parameter could not be found.<br>*aa...aa*: `offline`, `online`<br>Action<br>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 section 6.7. |

| | | |
|---|---|---|
| KAPL01016-W | The target CHA port was not found. Operation name = *aa...aa* | Details<br>The path ID indicated by -pathid and required by the `-cha` parameter is not an object of HDLM management.<br>*aa...aa*: `offline`, `online`<br>Action<br>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 section 6.7. |
| KAPL01018-W | The target device was not found. Operation name = *aa...aa* | Details<br>The specified host device name could not be found.<br>*aa...aa*: `view`<br>Action<br>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 section 6.7. |
| KAPL01019-W | The target path was not found. Operation name = *aa...aa* | Details<br>*aa...aa*: `offline`, `online`, `view`<br>▪ `offline`/`online` operation<br>The specified path does not exist.<br>▪ `view` operation<br>The paths have not been configured because creation of the HDLM environment or configuration changes to the HDLM operating environment have not finished.<br>Action<br>▪ `offline`/`online` operation<br>Use the `view` operation of the HDLM command (`dlnkmgr`) to check the specification, and then retry. For details on the `view` operation, see section 6.7.<br>▪ `view` operation<br>Refer to Chapter 3. Creating an HDLM Environment or 0, 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 section 7.2. |
| KAPL01021-E | Cannot execute the HDLM command due to insufficient memory. | Details<br>Memory required for HDLM command processing could not be allocated.<br>Action<br>Stop unnecessary applications, increase the amount of free memory, and then re-execute the HDLM command. |

| | | |
|---|---|---|
| KAPL01023-W | The last Online path for the device cannot be placed Offline(C). | Details<br>The path specified in the offline operation cannot be placed Offline(C) because that is the last path for the LU.<br>Action<br>Use the `view` operation of the HDLM command (`dlnkmgr`) to check the status of the path. For details on the `view` operation, see section 6.7. |
| KAPL01024-W | The specified parameters cannot be specified at the same time. Operation name = *aa...aa*, parameters = *bb...bb* | Details<br>*aa...aa*: `offline`, `online`, `set`, `view`<br>*bb...bb*: Parameters that cannot be specified at the same time<br>Action<br>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 section 6.3. |
| KAPL01036-E | The Offline path cannot be placed online. PathID = *aa...aa* | Details<br>*aa...aa*: Path ID (decimal number)<br>Action<br>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<br>A path that cannot be placed Online was detected during multi-path online processing.<br>To ignore this path and perform online processing for the next path, enter `y`.<br>To cancel processing, enter `n`.<br>*aa...aa*: Path ID (decimal number)<br>Action<br>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 section 6.5. |
| KAPL01040-W | The entered value is invalid. Re-enter [y/n]: | Details<br>A value other than y and n was entered. Enter y or n.<br>Action<br>Enter `y` or `n`. |
| KAPL01041-E | The entered value is invalid. The operation stops. Operation name = *aa...aa* | Details<br>Command processing will be aborted because an incorrect response was made three times in reply to the request.<br>*aa...aa*: `clear`, `offline`, `online`, `set`<br>Action<br>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<br>*aa...aa*: `view`<br>*bb...bb*: Parameter name<br>*cc...cc*: Duplicated parameter value<br>Action<br>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<br>*aa...aa*: `offline`, `online`, `set`, `view`<br>*bb...bb*: Parameter name<br>*cc...cc*: Parameter value<br>Action<br>Execute `help` *operation-name* of the HDLM command (`dlnkmgr`) to check the parameter value, and then retry. For details on the `help` operation, see section 6.3. |
| KAPL01048-W | Help information cannot be found. Operation name = *aa...aa*. | Details<br>The specified operation is not an operation of the HDLM command.<br>*aa...aa*: Specified operation name<br>Action<br>Use the `help` operation of the HDLM command (`dlnkmgr`) to check the operation name. And then retry. For details on the `help` operation, see section 6.3. |
| KAPL01049-I | Would you like to execute the operation? Operation name = *aa...aa* [y/n]: | Details<br>The `clear`/`set` operation will be started. To continue the operation, enter `y`. To cancel the operation, enter `n`.<br>*aa...aa*: `clear`, `set`<br>Action<br>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 section 6.2. For details on the `set` operation, see section 6.6. |
| KAPL01050-I | The currently selected paths will be changed to the Online status. Is this OK? [y/n]: | Details<br>The `online` operation will be started. To continue the `online` operation, enter `y`. To cancel the operation, enter `n`.<br>Action<br>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 section 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<br>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`.<br>Action<br>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 section 6.7. For details on the `online` operation, see section 6.5. |
| KAPL01052-I | The currently selected paths will be changed to the Offline(C) status. Is this OK? [y/n]: | Details<br>The `offline` operation will be started. To continue the `offline` operation, enter `y`. To cancel the operation, enter `n`.<br>Action<br>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 section 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<br>The `offline` operation will be started. To continue the `offline` operation, enter `y`. To cancel the operation, enter `n`.<br>Action<br>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 section 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<br>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`.<br>*aa...aa*: Path ID (decimal number)<br>Action<br>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 section 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<br>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`.<br>*aa...aa*: `CHA port`, `HBA`<br>Action<br>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<br>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.<br>*aa...aa*: `CHA port`, `HBA`<br>Action<br>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<br>All multiple paths will place online status because the `-hba` or `-cha` parameter was specified. To continue this operation, enter `y`. To discontinue this operation, enter `n`.<br>*aa...aa*: `CHA port`, `HBA`<br>Action<br>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<br>A parameter value was specified in a parameter that does not need any parameter value.<br>*aa...aa*: `clear`, `offline`, `online`, `set`, `view`<br>*bb...bb*: Parameter name<br>*cc...cc*: Parameter value<br>Action<br>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 section 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<br>A parameter value conflicts with the specification of another parameter.<br>*aa...aa*: Parameter name<br>*bb...bb*: Parameter name<br>*cc...cc*: Parameter value<br>*dd...dd*: `view`, `set`<br>Action<br>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 section 6.3. |
| KAPL01060-I | The user terminated the operation. Operation name = *aa...aa* | Details<br>The command processing will be aborted because n was entered in reply to the acknowledgment.<br>*aa...aa*: `online`, `offline`, `set`, `clear`<br>Action<br>None. |
| KAPL01061-I | *aa...aa* path(s) were successfully placed *bb...bb*; *cc...cc* path(s) were not. Operation name = *dd...dd* | Details<br>This message indicates the number of the paths processed in the `online`/`offline` operation.<br>*aa...aa*: Number of paths where `online`/`offline` operation is successful(decimal number)<br>*bb...bb*: Online or Offline(C)<br>*cc...cc*: Number of paths where `online`/`offline` is unsuccessful(decimal number)<br>*dd...dd*: `online`, `offline`<br>Action<br>None. For details on the `online` operation, see section 6.5. For details on the `offline` operation, see section 6.4. |
| KAPL01063-I | The target path(s) are already *aa...aa*. | Details<br>As a result of executing the `online`/`offline` operation, the specified path is already placed Online/Offline(C).<br>*aa...aa*: Online or Offline(C)<br>Action<br>Use the `view` operation of the HDLM command (`dlnkmgr`) to check the status of the path. For details on the `view` operation, see section 6.7. For details on the `online` operation, see section 6.5. For details on the `offline` operation, see section 6.4. |

| | | |
|---|---|---|
| KAPL01068-I | Enter a license key: | Details<br>The license key will now be renewed. Enter a license key.<br>Action<br>None. |
| KAPL01069-W | The entered license key is invalid. | Details<br>The entered license key is invalid.<br>Action<br>Enter a valid license key. |
| KAPL01070-E | The entered license key is invalid. Renewal of the license key will now stop. | Details<br>The license key renewal will be aborted because an invalid license key was entered three times.<br>Action<br>Obtain a valid license key, and then retry. |
| KAPL01071-I | The permanent license was installed. | Details<br>The license was renewed into a permanent license.<br>Action<br>None. |
| KAPL01072-I | The emergency license was installed. The license expires on *aa...aa*. | Details<br>A license was renewed into the emergency license.<br>*aa...aa*: The year of grace (4 numeric characters)/Month (01-12)/Day (01-31)<br>Action<br>Install a permanent license by the expiration day. |
| KAPL01073-E | The temporary license expired. | Details<br>The temporary license expired. Register a permanent license.<br>Action<br>Register a permanent license. |
| KAPL01074-E | The emergency license expired. | Details<br>The emergency license expired. Register a permanent license.<br>Action<br>Register a permanent license. |
| KAPL01075-E | A fatal error occurred in HDLM. The system environment is invalid. | Details<br>The license information file is missing.<br>Action<br>Re-install HDLM. |
| KAPL01076-I | The permanent license has been installed. | Details<br>You need not install a license because a permanent license has already been installed.<br>Action<br>None. |

| | | |
|---|---|---|
| KAPL01079-W | The intermittent error monitoring function cannot be set up because automatic failback is disabled. | Details<br>The intermittent error monitoring function cannot be set up because automatic failback is disabled.<br>Action<br>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<br>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.<br>Action<br>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<br>The format of the license key file is invalid.<br>*aa...aa*: `/var/tmp/hdlm_license`<br>Action<br>Save the correct license key file in the designated, and then re-execute.<br>`/var/tmp/hdlm_license` |
| KAPL01082-E | There is no installable license key in the license key file. File name = *aa...aa* | Details<br>There is no HDLM-installable license key in the license key file.<br>*aa...aa*: `/var/tmp/hdlm_license`<br>Action<br>Make sure that the license key file is correct, and then re-execute.<br>`/var/tmp/hdlm_license` |
| KAPL01083-I | There is no license key file. File name = *aa...aa* | Details<br>There is no license key file in the designated directory:<br>*aa...aa*: `/var/tmp/hdlm_license`<br>Action<br>When the message that prompts you to enter the license key is displayed, enter the license key.<br>Alternatively, cancel the HDLM command, save the correct license key file in the designated directory, and then re-execute the HDLM command.<br>*aa...aa*: `/var/tmp/hdlm_license` |
| KAPL01084-W | An attempt to delete the license key file has failed. File name = *aa...aa* | Details<br>*aa...aa*: `/var/tmp/hdlm_license`<br>Action<br>If the following 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<br>*aa...aa*: `view`<br>*bb...bb*: Parameter name<br>*cc...cc*: Parameter values cannot be specified at the same time<br>Action<br>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 section 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<br>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.<br>Do not perform the following operations:<br>▪ Simultaneous executions of the HDLM command's `set -lic` operation<br>▪ Execution of the HDLM command's `set -lic` operation simultaneously with an update of the license for an upgrade or re-installation |
| KAPL01095-E | An attempt to acquire the HDLM version information has failed. details = *aa....aa* | Details<br>*aa...aa*: Code showing the reason for the error<br>Action<br>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<br>*aa...aa*: Code showing the reason for the error<br>Action<br>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<br>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`.<br>Action<br>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<br>This message indicates the executed command line.<br>Action<br>None. |

| | | |
|---|---|---|
| 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<br>If you want to set the load balancing type for the system, enter `y`. If you want to terminate the processing, enter `n`. |
| KAPL01111-E | The version of the kernel supported by the installed HDLM does not match the currently running kernel version. | Details<br>The version of the kernel to which installed HDLM has adjusted and the kernel that is running now is not corresponding. There is a possibility of installing the package of the kernel after installing HDLM.<br>Action<br>Install HDLM that adjusts to the kernel that is running now. |
| KAPL01113-E | The HDLM file is invalid. File name = *aa...aa* | Details<br>A file necessary for processing HDLM cannot be found. Alternatively, the content of a file is invalid.<br>*aa...aa*: Name of file where error was detected.<br>Action<br>Reinstall an HDLM. |

## 8.3 HDLM API Messages

| | | |
|---|---|---|
| KAPL03001-I | HDLM API information - *aa...aa* | Details<br>This information is required for determining the cause of the problem (if any).<br>*aa...aa*: Trace information<br>Action<br>None. |
| KAPL03003-E | HDLM API Error information - *aa...aa* | Details<br>This information is required for determining the cause of the problem (if any).<br>*aa...aa*: API trace error information<br>Action<br>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 section 7.2. |
| KAPL03004-C | A critical error occurred in the HDLM API. (*aa...aa*) | Details<br>This information is required for determining the cause of the problem (if any).<br>*aa...aa*: API trace error information<br>Action<br>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 section 7.2. |
| KAPL03006-E | An access to the HDLM driver causes an error. (*aa...aa*) | Details<br>This information is required for determining the cause of the problem (if any).<br>*aa...aa*: API trace error information<br>Action<br>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 section 7.2. |

| | | |
|---|---|---|
| KAPL03007-E | An error occurred during communication with the HDLM manager. (*aa...aa*) | Details<br>This information is required for determining the cause of the problem (if any).<br>*aa...aa*: API trace error information<br>Action<br>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 section 7.2. |
| KAPL03008-E | An error occurred during log input to the HDLM alert driver. (*aa...aa*) | Details<br>This information is required for determining the cause of the problem (if any).<br>*aa...aa*: API trace error information<br>Action<br>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 section 7.2. |
| KAPL03999-E | An unexpected error occurred. | Details<br>A conflict occurred in the versions of the modules that HDLM uses internally.<br>Action<br>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 section 7.2. |

## 8.4 HDLM Manager Messages

| | | |
|---|---|---|
| KAPL04001-I | HDLM manager started. | Action<br>None. |
| KAPL04002-E | Could not start the HDLM manager. | Details<br>HDLM manager failed to start because the environment is incorrect for the manager to run properly.<br>Action<br>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 section 7.2. |
| KAPL04003-E | The startup parameter is invalid. | Details<br>The parameter held internally by the HDLM manager is incorrect.<br>Action<br>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 section 7.2. |
| KAPL04004-I | HDLM manager will now terminate. | Action<br>None. |
| KAPL04008-E | Cannot open the option definition file (*aa...aa*). | Details<br>HDLM manager cannot start normally (unable to open the option definition file).<br>*aa...aa*: Option definition file name<br>Action<br>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<br>HDLM manager cannot start normally (some of the definitions in the option definition file are invalid).<br>Action<br>If the KAPL04033-W message is output after this message, execute `dlnkmgr view -sys -sfunc` and check the option settings.<br>For options with setting values that have returned to default values, use the `dlnkmgr set` operation to reset the values.<br>If the KAPL04033-W message is not output, restart HDLM Manager.<br>If the same error occurs, re-install HDLM. For details on the `view` operation, see section 6.7. For details on the `set` operation, see section 6.6. |

| | | |
|---|---|---|
| KAPL04010-E | Could not open the error log file. | Details<br>HDLM manager cannot start normally (unable to open the error log file `/var/opt/DynamicLinkManager/log/dlmmgr[1-16].log`).<br>Action<br>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<br>The log information could not be output to the error log file `/var/opt/DynamicLinkManager/log/dlmmgr[1-16].log`.<br>Action<br>Check that the disk has sufficient free space. |
| KAPL04012-E | Could not create a communication pipe. RC = *aa...aa* | Details<br>HDLM manager cannot start normally (unable to create a pipe file to be used in communication with HDLM commands).<br>*aa...aa*: OS error code(decimal number)<br>Action<br>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 section 7.2. |
| KAPL04013-E | Input is impossible via the communication pipe. RC = *aa...aa* | Details<br>Data could not be read from the pipe file during the communication with the HDLM command.<br>*aa...aa*: OS error code(decimal number)<br>Action<br>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 section 7.2. |
| KAPL04014-E | Output is impossible via the communication pipe. RC = *aa...aa* | Details<br>Data could not be written to the pipe file during the communication with the HDLM command.<br>*aa...aa*: OS error code(decimal number)<br>Action<br>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 section 7.2. |
| KAPL04019-E | Could not collect the error information. RC = *aa...aa* | Details<br>An attempt to read the log information from the alert driver failed.<br>*aa...aa*: API return code(decimal number)<br>Action<br>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 section 7.2. |

| | | |
|---|---|---|
| KAPL04021-I | HDLM manager information - *aa...aa* | Details<br>This information is required for determining the cause of the problem (if any).<br>*aa...aa*: HDLM manager trace information<br>Action<br>None. |
| KAPL04022-W | HDLM manager warning information - *aa...aa* | Details<br>This information is required for determining the cause of the problem (if any).<br>*aa...aa*: HDLM manager trace warning information<br>Action<br>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 section 7.2. |
| KAPL04023-E | HDLM manager error information - *aa...aa* | Details<br>This information is required for determining the cause of the problem (if any).<br>*aa...aa*: HDLM manager trace error information<br>Action<br>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 section 7.2. |
| KAPL04024-C | A critical error occurred in the HDLM manager. (*aa...aa*) | Details<br>This information is required for determining the cause of the problem (if any).<br>*aa...aa*: HDLM manager trace error information<br>Action<br>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 section 7.2. |
| KAPL04025-C | A memory shortage occurred in the HDLM manager. | Details<br>Sufficient memory to process the HDLM manager could not be obtained.<br>Action<br>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<br>*aa...aa*: Expiration day<br>*bb...bb:* The year of grace (4 numeric characters)/Month (01-12)/Day (01-31)<br>Action<br>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<br>*aa...aa*: Expiration day<br>*bb...bb:* The year of grace (4 numeric characters)/Month (01-12)/Day (01-31)<br>Action<br>Install the emergency license by the expiration day. |
| KAPL04028-E | The temporary license expired. | Action<br>Install the permanent license. |
| KAPL04029-E | The emergency license expired. | Action<br>Install the permanent license. |
| KAPL04030-E | The temporary license has already expired. | Action<br>Install the permanent license. |
| KAPL04031-E | The emergency license has already expired. | Action<br>Install the permanent license. |
| KAPL04032-C | A fatal error occurred in HDLM. The system environment is invalid | Details<br>A part of the HDLM configuration file is missing.<br>Action<br>Re-install HDLM. |
| KAPL04033-W | The option definition file was re-created. | Details<br>An option definition file was re-created using the default values. The specified values are set when some of the options have been read.<br>Action<br>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 section 6.6. |
| KAPL04034-E | An attempt to create the option definition file has failed. | Details<br>An attempt to re-create an option definition file (`/etc/opt/DynamicLinkManager/dlmmgr.xml`) using the default values has failed.<br>Action<br>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<br>*aa...aa*: Total number of paths<br>Action<br>None. |
| KAPL04036-I | The path health check for the path *aa....aa* was executed. Number of error paths = *bb...bb* | Details<br>*aa...aa*: Number of paths targeted for the path health check.<br>*bb...bb*: Number of error paths by the path health check.<br>Action<br>None. |

| | | |
|---|---|---|
| KAPL04042-I | HDLM SNMP TRAP information - *aa...aa* | Details<br>There were no error paths as a result of executing the path health check.<br>*aa...aa*: Start or Stop.<br>Action<br>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<br>*aa...aa*: Trap ID<br>*bb...bb*: Destination IP address of the trap<br>*cc...cc*: Destination port number of the trap<br>*dd...dd*: Community Name given to the trap<br>*ee...ee*: Send data<br>Action<br>None. |

## 8.5 HDLM Driver(Filter Component) Messages

| | | |
|---|---|---|
| KAPL05003-I | The HDLM driver (filter component) was successfully attached to Disk (*a...aa*), Partition (*bb...bb*). | Details<br>The path corresponding to disk *aa...aa* and partition *bb...bb* was successfully registered in the core logic.<br>*aa...aa*: Disk sequence number (decimal number)<br>*bb...bb*: Fixed at 0 (decimal number)<br>Action<br>None. |
| KAPL05008-E | Could not allocate memory. (*aa...aa*:*bb...bb*) | Details<br>*aa...aa*: File ID, line number, (hexadecimal number)<br>*bb...bb*: Memory capture size (hexadecimal number)<br>Action<br>Execute the utility for setting HDLM driver option (`dlmsetopt`) and increase the number of retries for allocating memory. For details on the `dlmsetopt` utility, see section 7.7. Increase system memory. Stop any unnecessary processes. If the problem is not solved, 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 section 7.2. |
| KAPL05011-E | Could not attach the HDLM driver (filter component) to Disk (*aa...aa*), Partition (*bb...bb*). (*cc...cc*:*dd...dd*) | Details<br>The path corresponding to the disk *aa...aa* and partition *bb...bb* could not be registered in the core logic.<br>*aa...aa*: Disk sequence number (decimal number)<br>*bb...bb*: Partition number (decimal number)<br>*cc...cc*: Error code (hexadecimal number)<br>*dd...dd*: Filter driver management table address (hexadecimal number)<br>Action<br>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<br>The HDLM filter driver was successfully attached and registered as a path.<br>*aa...aa*: Address of table for managing the HDLM filter driver (hexadecimal number)<br>*bb...bb*: Path ID (hexadecimal number)<br>Action<br>None. |
| KAPL05018-W | The FO processing in the path (*aa...aa*) failed. (*bb...bb*:*cc...cc*) | Details<br>An attempt at FO processing for path *aa...aa* has failed.<br>*aa...aa*: Core logic path identifier for the failed FO (hexadecimal number)<br>*bb...bb*: Error code (hexadecimal number)<br>*cc...cc*: 0 (fixed)<br>Action<br>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<br>FO processing for the path *aa...aa* was successful. I/O requests was processed in the path *bb...bb*.<br>*aa...aa*: Path ID (hexadecimal number)<br>*bb...bb*: Path ID (hexadecimal number)<br>Action<br>None. |
| KAPL05020-I | Processing of IOCTL(*aa...aa*) will now start. | Details<br>An IOCTL request was received.<br>*aa...aa*: IOCTL code (hexadecimal number)<br>Action<br>None. |
| KAPL05021-I | Processing of IOCTL(*aa...aa*) completed normally. | Details<br>The processing for the requested IOCTL operation was successful.<br>*aa...aa*: IOCTL code (hexadecimal number)<br>Action<br>None. |

| | | |
|---|---|---|
| KAPL05023-E | Could not process the IOCTL(*aa...aa*).<br>(*bb...bb*:*cc...cc*) | Details<br>An attempt to process the requested ioctl call has failed.<br>*aa...aa*: ioctl code (hexadecimal number)<br>*bb...bb*: Error code (hexadecimal number)<br>*cc...cc*: ID of the process that issued the I/O<br><br>Action<br>Check the following. After the check, if the problem is not solved, 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 then report the error and detail code. For details on the `DLMgetras` utility, see section 7.2.<br>1. Make sure that the command is executed by a user with root permission.<br>2. The memory size is insufficient. Stop the unnecessary process.<br>3. Make sure that the ioctl is not issued to the device management file. |
| KAPL05701-I | HDLM version *aa...aa* Build *bb...bb cc...cc dd...dd ee...ee* | Details<br>The version of HDLM is shown.<br>*aa...aa*: HDLM version (character string)<br>*bb...bb*: HDLM build number (decimal number)<br>*cc...cc*: Days-and-months A.D (character string)<br>*dd...dd*: Time in seconds (character string)<br>*ee...ee*: Internal code (hexadecimal number) (This is information required for the investigation at the time the problem was generated.)<br>Action<br>None. |

| | | |
|---|---|---|
| KAPL05702-W | Path could not be placed online since partition information could not be read. (Host: *aa...aa* Channel: *bb...bb* ID: *cc...cc* LUN: *dd...dd*) | Details<br>Because partition information was not able to be read from the device, the path was not made online.<br>*aa...aa*: Host Port Number<br>*bb...bb*: Bus Number<br>*cc...cc*: Target ID<br>*dd...dd*: Host LU Number<br>Action<br>When using Block device<br>1. Delete the HDLM device of corresponding LU from the resource of LifeKeeper.<br>2. Execute umount to the corresponding HDLM device.<br>3. Connect cacle.<br>4. Make the state of the path online by the command.<br>5. Execute mount to the corresponfing HDLM device.<br>6. Add the HDLM device of corresponding LU from the resource of LifeKeeper.<br>When using raw device<br>1. Delete the HDLM device of corresponding LU from the resource of LifeKeeper.<br>2. Connect cacle.<br>3. Make the state of the path online by the command.<br>4. Add the HDLM device of corresponding LU from the resource of LifeKeeper. |

| | | |
|---|---|---|
| KAPL05704-E | The LDEV information for an HDLM device has changed. HCTL value = *aa...aa bb...bb cc...cc dd...dd* | Details<br>The LDEV information for the HCTL value corresponding to an HDLM device has changed.<br>LDEV information (storage subsystem model name, serial number, and LU number, separated by period)<br>*aa...aa*: Host ID (host port number) (decimal number)<br>*bb...bb*: Channel number (bus number) (decimal number)<br>*cc...cc*: Target ID (target ID) (decimal number)<br>*dd...dd*: LUN (host LU number) (decimal number)<br>Action<br>Please perform the following steps.<br>1. Note down the HCTL values of the SCSI device as shown in the error messages.<br>2. Obtain the corresponding HDLM device name for the HCTL value noted in step 1 by using the following command:<br>`# dlmcfgmgr -v`<br>3. Unregister the HDLM devices by using the following command:<br>`# dlmcfgmgr -u` *logical-device-file-name-of-the-HDLM-device*<br>4. Reconfigure the HDLM by using the following command:<br>`# dlmcfgmgr -r` |
| KAPL05708-E | The system retried to allocate memory *aa...aa* times, but the retries failed. (*bb...bb*:*cc...cc*) | Details<br>*aa...aa*: Memory allocation retry count (decimal number)<br>*bb...bb*: File ID, line number (hexadecimal number)<br>*cc...cc*: Memory allocation size (hexadecimal number)<br>Action<br>Execute the utility for setting HDLM driver option (`dlmsetopt`) and increase the number of retries for allocating memory. For details on the `dlmsetopt` utility, see section 7.7. Increase system memory. Stop any unnecessary processes. If the problem is not solved, 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 section 7.2. |
| KAPL05709-I | The system will retry to allocate memory *aa...aa* times. | Details<br>*aa...aa*: Memory allocation retry count (decimal number)<br>Action<br>None. |

| | | |
|---|---|---|
| KAPL05711-I | The timeout value and retry count of SCSI INQUIRY has been set. (timeout value = *aa...aa*, retry count = *bb...bb*) | Details<br>*aa...aa*: SCSI INQUIRY timeout value<br>*bb...bb*: SCSI INQUIRY retry count<br>Action<br>None. |

## 8.6 HDLM Alert Driver Messages

| | | |
|---|---|---|
| KAPL06004-E | Could not allocate memory. (*aa...aa*:*bb...bb*) | Details<br>An attempt to reserve memory to save alert information has failed.<br>*aa...aa*: Program line (hexadecimal number)<br>*bb...bb*: Target memory size (hexadecimal number)<br>Action<br>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 section 7.2. |
| KAPL06007-I | IOCTL processing (*aa...aa*) will now start. | Details<br>The ioctl request *aa...aa* was accepted.<br>*aa...aa*: ioctl code (hexadecimal number)<br>Action<br>None. |
| KAPL06008-I | IOCTL processing (*aa...aa*) completed normally. | Details<br>The ioctl request *aa...aa* was processed normally.<br>*aa...aa*: ioctl code (hexadecimal number)<br>Action<br>None. |
| KAPL06009-I | Invalid IOCTL(*aa...aa*) was received. The processing is canceled. | Details<br>A request having an invalid IOCTL code *aa...aa* was issued to the alert driver.<br>*aa...aa*: IOCTL code (hexadecimal number)<br>Action<br>None. |

| | | |
|---|---|---|
| KAPL06010-E | Could not process the IOCTL(*aa...aa*). (*bb...bb*:*cc...cc*) | Details<br>An IOCTL request was not processed normally.<br>*aa...aa*: IOCTL code (hexadecimal number)<br>*bb...bb*: Error number (hexadecimal number)<br>*cc...cc*: 0 (fixed)<br>Action<br>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 section 7.2. |
| KAPL06013-E | Could not write log information into the log buffer. (*aa...aa*:*bb...bb*) | Details<br>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.<br>*aa...aa*: Message code (hexadecimal number)<br>*bb...bb*: Buffer size (hexadecimal number)<br>Action<br>Check whether any other error occurred. The information that could not be written is discarded.<br>Review the actual memory size when another error does not occur.<br>When the actual memory size is insufficient, increase the actual memory size.<br>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 section 7.2. |

| | | |
|---|---|---|
| KAPL06014-E | Could not write emergency information into the emergency information buffer. (*aa...aa*:*bb...bb*) | Details<br>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.<br>*aa...aa*: Message code (hexadecimal number)<br>*bb...bb*: Buffer size (hexadecimal number)<br>Action<br>Check whether any other error occurred. The information that could not be written is discarded.<br>Review the actual memory size if another error does not occur.<br>If the actual memory size is insufficient, increase it.<br>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 section 7.2. |
| KAPL06100-E | The version of the kernel supported by the installed HDLM does not match the currently booting kernel version. | Details<br>The version of the kernel supported by the installed HDLM is different from the currently running kernel version.<br>It may be that you booted the kernel that is different from the kernel when you installed HDLM.<br>Action<br>Install an HDLM that supports the currently running kernel. |
| KAPL06101-E | The HDLM file is invalid. File name = *aa...aa* | Details<br>A file necessary for processing HDLM cannot be found or the content of a file is invalid. Therefore, the Alert driver cannot be loaded.<br>*aa...aa*: Name of file where error was detected.<br>Action<br>Reinstall an HDLM. |

## 8.7 HDLM Driver (Core Logic Component) essages

| | | |
|---|---|---|
| KAPL07819-I | Data for maintenance: *aa...aa bb...bb cc...cc dd...dd*. | Details<br>This message is generated by the core logic for maintenance.<br>*aa...aa*: Detailed information 1 (decimal number)<br>*bb...bb*: Internal function number of the core logic (decimal number)<br>*cc...cc*: Detailed information 2 (decimal number)<br>*dd...dd*: Detailed information 3 (decimal number)<br>Action<br>None. |

## 8.8 HDLM Management Target Messages

| | | |
|---|---|---|
| KAPL08019-E | The path (*aa...aa*) detected an error (*bb...bb*). (*cc...cc*) | Details<br>An error occurred in the path because of a condition such as a disconnection.<br>*aa...aa*: Path identifier (hexadecimal number)<br>*bb...bb*: Error code (hexadecimal number)<br>▪ When the path error was detected by a path health checking or the online operation<br>Displays 0x000F0000 (Fixed).<br>▪ When the path error was detected through an I/O error<br>Displays the OS error code.<br>*cc...cc*: 0x00000000 (fixed)<br>Action<br>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<br>A physical or logical error occurred in the path.<br>*aa...aa*: OS error code (hexadecimal number)<br>▪ When the path error was detected by a path health checking or the online operation<br>Displays 0x000F0000 (Fixed).<br>▪ When the path error was detected through an I/O error<br>Displays the OS error code.<br>*bb...bb*: Path ID (same as PathID of `view -path`) (decimal number)<br>*cc...cc*: Host port number (same as PathName of `view -path`) (hexadecimal number)<br>*dd...dd*: Bus number (hexadecimal number)<br>*ee...ee*: Target ID (hexadecimal number)<br>*ff...ff*: HLU number (same as PathName of `view -path`) (hexadecimal number)<br>*gg...gg*: Dev number (same as DNum of `view -path`) (decimal number)<br>*hh...hh*: Host device name<br>Action<br>There could be an error in the path. See section 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<br>*aa...aa*: Path ID (same as PathID of `view -path`) (decimal number)<br>*bb...bb*: Host port number (same as PathName of `view -path`) (hexadecimal number)<br>*cc...cc*: Bus number (same as PathName of `view -path`) (hexadecimal number)<br>*dd...dd*: Target ID (same as PathName of `view -path`) (hexadecimal number)<br>*ee...ee*: HLU number (same as PathName of `view -path`) (hexadecimal number)<br>*ff...ff*: Device number (same as DNum of `view -path`) (decimal number)<br>*gg...gg*: Host Dev name (same as HDevName of `view -path`)<br>Action<br>None. |
| KAPL08026-E | An error occurred on all the paths of the LU. PathID = *aa...aa* | Details<br>An error occurred in the last path of one LU because of a condition such as a disconnection.<br>*aa...aa*: Path ID (same as PathID of `view -path`) (decimal number)<br>Action<br>Errors are detected in all the paths connected to the LUs. See section 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<br>A path was excluded from the items subject to automatic failback because the system judged that an intermittent error was occurring in that path.<br>*aa...aa*: Path ID (same as PathID of `view -path`) (decimal number)<br>Action<br>An intermittent error was occurring. The path may have a problem. For details on the action to be taken, see section 5.3 and switch the path shown in the message into Online. |

## 8.9 HDLM Installation Program Messages

| | | |
|---|---|---|
| KAPL09001-E | There is no system management permission. | Details<br>The current user does not have the administrator permission to install HDLM.<br>Action<br>Re-execute the installation program as a user with root permission. |
| KAPL09003-E | Cannot install in this system. | Details<br>HDLM cannot be installed on this system.<br>Action<br>Confirm the following:<br>▪ The system is a system for which HDLM is supported.<br>▪ The package is compatible with the system. |
| KAPL09005-E | Could not stop the HDLM manager. | Details<br>An attempt to stop the HDLM manager service failed.<br>Action<br>Stop the HDLM manager manually, and then re-execute the installation program or the uninstallation program. |
| KAPL09011-E | Cannot find a license key file /etc/opt/DynamicLinkManager/dlm.lic_key | Details<br>The license key file, `/etc/opt/DynamicLinkManager/dlm.lic_key`, cannot be found in the specified directory.<br>Action<br>Create a license key file, and re-execute the installation program. |
| KAPL09013-E | Some HDLM drivers could not be removed. | Details<br>Deleting several HDLM drivers failed because these HDLM drivers were being used for HDLM upgrade installation, re-installation, or uninstallation.<br>Action<br>Check whether an HDLM device is mounted. If an HDLM device has been mounted, unmount the device, and then remove the HDLM drivers again. For details on the `dlmcfgmgr` utility, see section 7.3. |
| KAPL09019-E | An attempt to cancel the registration of the bundle PP name of Hitachi Network Objectplaza Trace Library 2 failed. | Action<br>Manually cancel the registration of the bundle program product (PP) name and uninstall HNTRLib2. If the attempt to cancel the registration of the bundle PP name and to uninstall HNTRLib2 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<br>An attempt to uninstall HNTRLib2 has failed.<br>Action<br>Manually uninstall HNTRLib2. If the attempt to uninstall HNTRLib2 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<br>An attempt to register the PP name of HNTRLib 2 has failed.<br>Action<br>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<br>A target file to copy to the directory of HiCommand products other than HDLM could not be found among the files related to HDLM.<br>Action<br>Re-execute the HDLM installation program. |
| KAPL09024-E | An attempt to copy a file or directory related to HDLM has failed. | Details<br>An attempt to copy a file related to HDLM to the directory of HiCommand products other than HDLM has failed.<br>Action<br>Re-execute the HDLM installation program. |
| KAPL09026-I | Hitachi Network Objectpraza Trace Library 2 wasn't uninstalled because it was being used for other products. | Details<br>Hitachi Network Objectplaza Trace Library 2 was not uninstalled because it was being used for other products.<br>Action<br>None. |
| KAPL09028-E | An attempt to install Hitachi Network Objectplaza Trace Library 2 failed. | Details<br>An attempt to install HNTRLib2 failed.<br>Action<br>Contact your HDLM vendor or the maintenance company if there is a maintenance contact of HDLM. |
| KAPL09035-E | The HDLM driver could not be loaded. | Details<br>The HDLM driver could not be loaded.<br>Action<br>Carry out the following process:<br>Load the filter driver.<br>Execute dlmcfgmgr -r.<br>Start DLMManager. |
| KAPL09036-E | The HDLM manager could not be started. | Details<br>The HDLM manager could not be started.<br>Action<br>Execute the `/etc/init.d/DLMManager` command with the start parameter to start the HDLM manager. |
| KAPL09037-E | The status of the HDLM manager could not be checked. | Details<br>An attempt to acquire the status of the HDLM manager has failed.<br>Action<br>If the HDLM manager does not start, execute the `/etc/init.d/DLMManager` command with the start parameter to start the HDLM manager. |

| | | |
|---|---|---|
| KAPL09038-E | The HDLM configuration could not be deleted. | Details<br>Deleting the HDLM configuration failed because a managed HDLM device was being used.<br>Action<br>Stop applications that are using the managed HDLM device, unmount the managed HDLM device, and then retry uninstallation. |
| KAPL09039-E | The HDLM devices could not be configured. | Details<br>The HDLM devices could not be configured.<br>Action<br>Execute the HDLM-configuration definition utility (`dlmcfgmgr`) with the `-r` parameter specified. For details on the `dlmcfgmgr` utility, see section 7.3. |
| KAPL09040-E | The file could not be edited. File name = *aa...aa* | Details<br>An attempt to edit a file has failed.<br>*aa...aa*: File name<br>Action<br>Make sure you can edit the file *aa...aa*. |
| KAPL09043-I | The installation of *aa...aa* completed successfully. | Details<br>HDLM was installed.<br>*aa...aa*: Name of installed file<br>Action<br>None. |
| KAPL09044-I | The uninstallation of *aa...aa* completed successfully. | Details<br>HDLM was uninstalled.<br>*aa...aa*: Name of uninstalled file.<br>Action<br>None. |
| KAPL09045-E | An attempt to create a file or directory related to HDLM has failed. | Details<br>Installation could not be executed because an attempt to create an HDLM-related file or directory failed.<br>Action<br>Retry installation. |
| KAPL09046-E | The RPM file is not in the correct path. | Details<br>Installation could not be executed because there was no RPM file in a suitable path.<br>Action<br>Confirm that the installation medium is correct, and then retry installation. |
| KAPL09049-W | An attempt to delete a file or directory related to HDLM has failed. Name = *aa...aa* | Details<br>An attempt to delete a HDLM-related file or directory has failed.<br>*aa...aa*: File name or a directory name<br>Action<br>Delete the file or directory shown in *aa...aa*. |

| | | |
|---|---|---|
| KAPL09050-E | The driver module dependencies could not be updated. | Details<br>Installation or uninstallation failed because the dependency information for a driver module could not be updated.<br>Action<br>Contact your HDLM vendor or the maintenance company if there is a maintenance contact of HDLM. |
| KAPL09051-E | You do not have permission to execute the file. File name = *aa...aa* | Details<br>Installation or uninstallation failed because there was no execution permission for a file.<br>*aa...aa*: File name<br>Action<br>Make sure you can execute the file shown in *aa...aa*. |
| KAPL09076-I | The permanent license was installed. | Details<br>The permanent license was installed.<br>Action<br>None. |
| KAPL09077-I | The temporary license was installed. The license expires on *aa...aa*. | Details<br>The temporary license was installed.<br>*aa...aa*: The year of grace (4 numeric characters)/Month (01-12)/Day (01-31)<br>Action<br>Install the permanent license by the expiration day. |
| KAPL09078-I | The emergency license was installed. The license expires on *aa...aa*. | Details<br>The emergency license was installed.<br>*aa...aa*: The year of grace (4 numeric characters)/Month (01-12)/Day (01-31)<br>Action<br>Install the permanent license by the expiration day. |
| KAPL09079-I | The permanent license has been installed. | Action<br>None. |
| KAPL09080-I | The temporary license has been installed. The license expires on *aa...aa*. | Details<br>The temporary license has been installed.<br>*aa...aa*: The year of grace (4 numeric characters)/Month (01-12)/Day (01-31)<br>Action<br>Install the permanent license by the expiration day. |
| KAPL09081-I | The emergency license has been installed. The license expires on *aa...aa*. | Details<br>The emergency license has been installed.<br>*aa...aa*: The year of grace (4 numeric characters)/Month (01-12)/Day (01-31)<br>Action<br>Install the permanent license by the expiration day. |
| KAPL09082-W | The temporary license expired. | Action<br>Enter the permanent license key. |

| | | |
|---|---|---|
| KAPL09083-W | The emergency license expired. | Action<br>Enter the permanent license key. |
| KAPL09087-E | The entered license key is invalid. Renewal of the license key will now stop. | Details<br>License renewal will now stop because an invalid license key was entered.<br>Action<br>Obtain a valid license key, and then retry. |
| KAPL09088-E | The entered license key is invalid. The HDLM installation will now terminate. | Action<br>Obtain a valid license key, and then retry. |
| KAPL09090-W | This operation will now be continued without updating the license. | Details<br>This operation will be continued without updating the license.<br>Action<br>Install the permanent license later. |
| KAPL09091-E | A fatal error occurred in HDLM. The system environment is invalid. | Details<br>A part of the HDLM configuration file is missing.<br>Action<br>Re-install HDLM. |
| KAPL09092-I | The installation was stopped because a stop request was received. | Details<br>The installation was aborted, because a stop request was received.<br>Action<br>Installation was terminated before all processing could be completed. If the HDLM RPM has been installed, then use the rpm command to perform uninstallation. |
| KAPL09093-I | *aa...aa* will be installed. Is this OK? [y/n]: | Details<br>The message checks whether you want to install *aa...aa*.<br>Action<br>Please enter "`y`" to install, and "`n`" to abort the installation. |
| KAPL09094-W | The entered value is invalid. Re-enter [y/n]: | Details<br>The message checks whether you want to correct an invalid value.<br>Action<br>Please enter "`y`" or "`n`". |
| KAPL09095-E | The entered value is invalid. The installation has been stopped. | Details<br>The entered value is invalid. The installation has been stopped.<br>Action<br>To install HDLM, please re-execute the HDLM installation program. |
| KAPL09096-I | The user stopped the installation. | Details<br>Installation was aborted by the user.<br>Action<br>None. |

| | | |
|---|---|---|
| KAPL09097-E | A package installable on this system was not found. | Details<br>A suitable package for installation on this system could not be found.<br>Action<br>Please check whether HDLM is supported for this system. For details about installing the package on a supported OS, see section 3.1.1.1. |
| KAPL09098-E | An attempt to install *aa...aa* has failed. Code = *bb...bb* | Details<br>An attempt to install *aa...aa* has failed.<br>*aa...aa*: Program name<br>*bb...bb*: Code (decimal number).<br>Action<br>Please re-install HDLM after checking the error information and resolving the problems. |
| KAPL09112-E | The license key file is invalid. File name = *aa...aa* | Details<br>The format of the license key file is invalid.<br>*aa...aa*: `/var/tmp/hdlm_license`<br>Action<br>Save the correct license key file in the designated directory, and then re-execute.<br>`/var/tmp/hdlm_license` |
| KAPL09113-E | There is no installable license key in the license key file. File name = *aa...aa* | Details<br>There is no HDLM-installable license key in the license key file.<br>*aa...aa*: `/var/tmp/hdlm_license`<br>Action<br>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<br>There is no license key file in the designated directory.<br>*aa...aa*: `/var/tmp/hdlm_license`<br>Action<br>Save the correct license key file in the designated directory, and then re-execute installation.<br>`/var/tmp/hdlm_license` |
| KAPL09115-W | An attempt to delete the license key file has failed. File name = *aa...aa* | Details<br>An attempt to delete the license key file has failed.<br>*aa...aa*: `/var/tmp/hdlm_license`<br>Action<br>If a license key file exists, delete it.<br>`/var/tmp/hdlm_license` |

| | | |
|---|---|---|
| KAPL09116-W | The command could not be installed. (command = *aa...aa*) | Details<br>The output HDLM command cannot be used.<br>*aa...aa*: Command name<br>Action<br>The output command can be executed by using a different name. If you want to use the output name, use the output command to overwrite or re-install. |
| KAPL09121-E | *aa...aa* is not present at *bb...bb*. | Details<br>*aa...aa*: Installer name<br>*bb...bb*: Installer path<br>Action<br>Install HDLM from the CD-ROM.<br>When you are attempting to install HDLM from the directory to which the CD-ROM was copied, make sure that all the contents of the CD-ROM have been copied to that directory, and then retry the installation. |
| KAPL09122-E | A Script related to LifeKeeper failed : *aa...aa* | Details<br>A Script related to LifeKeeper failed : *aa...aa*<br>*aa...aa*: The failed command<br>Action<br>1. Make sure the file `/etc/default/LifeKeeper` exists. If it does not exist, install LifeKeeper, and then re-install HDLM.<br>2. Make sure the directory `/etc/opt/DynamicLinkManager` exists. If it does not exist, HDLM might not be installed successfully. Re-install HDLM. |
| KAPL09123-W | The configuration of LifeKeeper with HDLM could not be retrieved properly. | Details<br>The configuration of HDLM with LifeKeeper could not be retrieved correctly.<br>Action<br>1. Make sure the user has execution permission for the file `/opt/LifeKeeper/bin/lktest`.<br>2. Make sure the file `/etc/default/LifeKeeper` exists. If it does not exist, install LifeKeeper, and then re-install HDLM. |

| | | |
|---|---|---|
| KAPL09124-W | The configuration of LifeKeeper with HDLM could not be removed completely. | Details<br>The configuration of HDLM with LifeKeeper could not be deleted completely.<br>Action<br>When an uninstallation is performed:<br>If there is a directory whose name is `DLMLFK` under the `/etc/opt/DynamicLinkManager` directory, delete the directory. LifeKeeper might not be installed successfully. Re-install LifeKeeper.<br>When the `/etc/opt/DynamicLinkManger` directory exists after performing an upgrade installation:<br>If there is a directory whose name is `DLMLFK` under the `/etc/opt/DynamicLinkManager` directory, delete the directory. LifeKeeper might not be installed successfully. Re-install LifeKeeper.<br>When the `/etc/opt/DynamicLinkManger` directory does not exists after performing an upgrade installation:<br>HDLM might not be installed successfully. Re-install HDLM. |
| KAPL09125-E | A file related to HDLM does not have execute permission. File Name : *aa...aa* | Details<br>*aa...aa*: File name<br>Action<br>Grant execution permission to the `dlmlfk` commands, and then re-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<br>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.<br>Do not perform the following operation:<br>Execution of the HDLM command's set -lic operation simultaneously with an update of the license for an upgrade installation |
| KAPL09142-E | HDLM *aa...aa* cannot be performed. Wait a while, and then perform *aa...aa* again. Error Code =*bb...bb* | Details<br>HDLM cannot be installed or uninstalled.<br>*aa...aa* : "installation" or "uninstallation"<br>*bb...bb* : Internal code (decimal number)<br>Action<br>Wait a while, and then reperform the installation or uninstallation. |
| KAPL09143-E | HDLM *aa...aa* cannot be performed. Error Code = *bb...bb* | Details<br>HDLM cannot be installed or uninstalled.<br>*aa...aa*: "installation" or "uninstallation"<br>*bb...bb*: Internal code (decimal number)<br>Action<br>Contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM. |

| | | |
|---|---|---|
| KAPL09146-E | A directory required by HDLM not found. Directory = *aa...aa* | Details<br>*aa...aa*: Referenced directory name<br>Action<br>Check the directory displayed in the message. Re-install HDLM after creating the directory displayed in the message. |
| KAPL09147-E | An attempt to create a file related to HDLM has failed. File name = *aa...aa* | Details<br>*aa...aa*: File that an attempt was made to create<br>Action<br>Remove unnecessary files and secure free space on the file system. Check the write permissions for the directory. Re-install HDLM. |
| KAPL09160-E | HDLM does not support this kernel. | Details<br>HDLM does not support the current kernel, or it does not support the combination of the current CPU and the kernel.<br>In addition, environments in which a 32-bit kernel is installed on an AMD64 CPU are not supported except for #3 and #7 as shown in Table 3.2.<br>Action<br>Make sure the current kernel and the running server's CPU are of the correct type. |
| KAPL09161-E | The mount point was not found. Mount point = *aa...aa* | Details<br>The mount point that is specified in the message does not exist.<br>*aa...aa*: mount point<br>▪ For Red Hat Enterprise Linux AS3 or Red Hat Enterprise Linux ES3:<br>`/mnt/cdrom`<br>▪ For Red Hat Enterprise Linux AS4 or Red Hat Enterprise Linux ES4:<br>`/media/cdrom, /media/cdrecorder`<br>▪ For SUSE LINUX Enterprise Server 9<br>`/media/cdrom, /media/cdrecorder, /media/dvd, /media/dvdrecorder, /media/dvdram`<br>▪ For SUSE LINUX Enterprise Server 10<br>`/media/media-volume-ID`<br>Action<br>Check the mount point where HDLM was mounted. Mount the HDLM CD-ROM in the mount point that is specified in the message. |

| | | |
|---|---|---|
| KAPL09162-E | The HDLM package was not found in the mount point. Mount point = *aa...aa* | Details<br>The CD-ROM of HDLM has not been mounted in the mount point that is specified in the message.<br>*aa...aa*: mount point<br>▪ For Red Hat Enterprise Linux AS3 or Red Hat Enterprise Linux ES3:<br>`/mnt/cdrom`<br>▪ For Red Hat Enterprise Linux AS4 or Red Hat Enterprise Linux ES4:<br>`/media/cdrom, /media/cdrecorder`<br>▪ For SUSE LINUX Enterprise Server 9<br>`/media/cdrom, /media/cdrecorder, /media/dvd, /media/dvdrecorder, /media/dvdram`<br>▪ For SUSE LINUX Enterprise Server 10<br>`/media/media-volume-ID`<br>Action<br>Check the mount point where HDLM was mounted. Mount the HDLM CD-ROM in the mount point that is specified in the message.<br>Check the CD-ROM mounted in the mount point. |
| KAPL09163-E | An HDLM package installable on this kernel was not found in the mount point. Mount point = *aa...aa* | Details<br>The HDLM for the current kernel has not been stored in the HDLM CD-ROM being mounted in the mount point that is specified in the message.<br>*aa...aa*: mount point<br>▪ For Red Hat Enterprise Linux AS3 or Red Hat Enterprise Linux ES3:<br>`/mnt/cdrom`<br>▪ For Red Hat Enterprise Linux AS4 or Red Hat Enterprise Linux ES4:<br>`/media/cdrom, /media/cdrecorder`<br>▪ For SUSE LINUX Enterprise Server 9<br>`/media/cdrom, /media/cdrecorder, /media/dvd, /media/dvdrecorder, /media/dvdram`<br>▪ For SUSE LINUX Enterprise Server 10<br>`/media/media-volume-ID`<br>Action<br>Check the HDLM CD-ROM, or check the current kernel. |
| KAPL09177-I | HDLM version: *aa...aa* | Details<br>*aa...aa*: The version of HDLM to be installed.<br>Action<br>None. |
| KAPL09501-E | HDLM is not installed on this system. | Details<br>The SP cannot be applied because HDLM is not installed on this system.<br>Action<br>Check whether HDLM has been correctly installed. |

| | | |
|---|---|---|
| KAPL09505-E | *aa...aa* cannot be applied to the installed *bb...bb*. | Details<br>This HDLM or SP (*aa...aa*) cannot be used for an upgrade installation or re-installation on the already installed HDLM or installed SP (*bb...bb*).<br>*aa...aa*: The version of HDLM to be installed, or the version of the HDLM SP to be installed<br>*bb...bb*: The installed version of HDLM, or the installed version of the HDLM SP to be installed<br>Action<br>An upgrade installation or re-installation cannot be performed on the already installed HDLM or SP.<br>When installing HDLM: Uninstall the installed HDLM or SP, and then perform installation.<br>When installing a SP: Obtain, and then install, a SP 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<br>The same version as that of SP to be installed is installed. Installation of SP is stopped.<br>*aa...aa*: Version of the SP to be installed<br>Action<br>You do not have to install the SP. Keep using the already installed HDLM. |
| KAPL09510-E | Service Pack *aa...aa* cannot be installed. A newer version has already been installed. | Details<br>A newer version is already installed. Installation of SP (*aa...aa*) is stopped.<br>*aa...aa*: Version of SP to be installed<br>Action<br>You do not have to install the SP. Keep using the already installed HDLM. |

## 8.10 Messages from the DLMgetras Utility for Collecting HDLM Error Information

| | | |
|---|---|---|
| KAPL10001-W | No parameter has been specified. | Details<br>No parameter (directory to which collected information is output) has been specified.<br>Action<br>Check the parameters of the `DLMgetras` utility for collecting HDLM error information, and then retry. For details on the `DLMgetras` utility, see section 7.2. |
| KAPL10002-W | Too many parameters have been specified. | Details<br>Four or more parameters have been specified.<br>Action<br>Check the parameters of the `DLMgetras` utility for collecting HDLM error information, and then retry. For details on the `DLMgetras` utility, see section 7.2. |
| KAPL10003-W | The first parameter has not been set to a directory. Value = *aa...aa* | Details<br>A value which is not a directory name is set for the first parameter. The first parameter must be set to a directory to which collected information is output.<br>*aa...aa*: First parameter<br>Action<br>Check the parameters of the `DLMgetras` utility for collecting HDLM error information, and then retry. For details on the `DLMgetras` utility, see section 7.2. |
| KAPL10004-W | The parameter contains an incorrect value. Value = *aa...aa* | Details<br>The first parameter must be a directory. The second parameter must be `-f`.<br>*aa...aa*: Invalid parameter<br>Action<br>Check the parameters of the `DLMgetras` utility for collecting HDLM error information, and then retry. For details on the `DLMgetras` utility, see section 7.2. |
| KAPL10005-W | The number of parameters is insufficient. | Details<br>The `-f` parameter exists but the file for defining the information to be collected does not exist. The number of parameters is insufficient.<br>Action<br>Check the parameters of the `DLMgetras` utility for collecting HDLM error information, and then retry. For details on the `DLMgetras` utility, see section 7.2. |

| | | |
|---|---|---|
| KAPL10006-W | The file for defining the information to be collected does not exist, or cannot be read. Value = *aa...aa* | Details<br>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.<br>*aa...aa*: Name of the file for defining the information to be collected<br>Action<br>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<br>A directory is specified for the `-f` parameter.<br>*aa...aa*: Third parameter<br>Action<br>Check the parameters of the `DLMgetras` utility for collecting HDLM error information, and then retry. For details on the `DLMgetras` utility, see section 7.2. |
| KAPL10008-W | You lack write permission for the specified directory. Value = *aa...aa* | Details<br>You do not have write permission for the specified directory, or the creation of a subdirectory of the specified directory failed.<br>*aa...aa*: first parameter<br>Action<br>Check the following.<br>1. Check whether you have access permission for the specified directory.<br>2. Check whether the specified directory name is correct.<br>3. Check that the disk has sufficient free space. |
| KAPL10009-W | The specified directory already exists. Do you want to overwrite it? [y/n]: | Details<br>The specified directory already exists. To overwrite the directory, enter `y`. To cancel this operation, enter `n`.<br>Action<br>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 section 7.2. |

| | | |
|---|---|---|
| KAPL10010-W | A root directory has been specified. Line = *aa...aa* | Details<br>The root "/" has been specified as a directory to be collected in the file for defining the information to be collected.<br>*aa...aa*: Line number of the file for defining information to be collected (decimal number)<br>Action<br>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 section 7.2. |
| KAPL10011-W | More than one file or directory has been specified on one line. Line = *aa...aa,* Value = *bb...bb* | Details<br>Two or more file names or directory names exist in the file for defining the information to be collected.<br>*aa...aa*: Line number of the file for defining information to be collected (decimal number)<br>*bb...bb:* Indicated contents in a line<br>Action<br>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 section 7.2. |
| KAPL10012-W | The specified file or directory does not exist. Line = *aa...aa,* Value = *bb...bb* | Details<br>The specified file or directory does not exist in the file for defining the information to be collected.<br>*aa...aa*: Line number of the file for defining information to be collected (decimal number)<br>*bb...bb*: Indicated contents in the line<br>Action<br>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 section 7.2. |

| | | |
|---|---|---|
| KAPL10013-W | You lack read permission for the specified file. Line = *aa...aa,* Value = *bb...bb* | Details<br>You lack read permission for the specified file in the file for defining information to be collected.<br>*aa...aa*: Line number of the file for defining information to be collected (decimal number)<br>*bb...bb*: Indicated contents in a line<br>Action<br>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 section 7.2. |
| KAPL10014-W | You lack read permission for the specified directory. Line = *aa...aa,* Value = *bb...bb* | Details<br>You lack read permission for the specified directory in the file for defining information to be collected.<br>*aa...aa*: Line number of the file for defining information to be collected (decimal number)<br>*bb...bb*: Indicated contents in a line<br>Action<br>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 section 7.2. |
| KAPL10015-W | The file format is invalid. Value = *aa...aa* | Details<br>The file format is invalid. The file format in the file for defining information to be collected is not a text file.<br>*aa...aa*: Third parameter<br>Action<br>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 section 7.2. |

| | | |
|---|---|---|
| KAPL10016-W | The root directory has been specified in the first parameter. | Details<br>The root directory has been specified in the first parameter. A root "/" cannot be specified in a directory to which collected information is output.<br>Action<br>Check the parameters of the `DLMgetras` utility for collecting HDLM error information, and then re-execute. For details on the `DLMgetras` utility, see section 7.2. |
| KAPL10017-W | You lack privileges for executing the utility for collecting HDLM error information. | Details<br>The `DLMgetras` utility for collecting HDLM error information must be executed by a user with root privileges.<br>Action<br>Re-execute as a user with root privileges. For details on the `DLMgetras` utility, see section 7.2. |
| KAPL10020-I | The file has been obtained successfully. File = *aa...aa*, Collection time = *bb...bb*(GMT:*bb...bb*) | Details<br>The file to be collected has been obtained.<br>*aa...aa*: Collected file name<br>*bb...bb*: the year of grace/month/day hour:minute:second<br>Action<br>None. |
| KAPL10021-I | Processing terminated before completion because a signal was received. | Details<br>The process has been terminated by an operation such as `Ctrl` + `c`.<br>Action<br>The utility for collecting HDLM error information terminated before completion.<br>If the directory is unnecessary, delete directory. For details on the `DLMgetras` utility, see section 7.2. |
| KAPL10022-I | The utility for collecting HDLM error information completed normally. | Details<br>Error information has been collected.<br>Action<br>None. For details on the `DLMgetras` utility, see section 7.2. |
| KAPL10030-I | A user terminated the utility for collecting HDLM error information. | Details<br>Processing of the `DLMgetras` utility has been terminated because `n` was sent as a confirmation reply.<br>Action<br>None. For details on the `DLMgetras` utility, see section 7.2. |

| | | |
|---|---|---|
| KAPL10031-W | The entered value is invalid. Continue operation? [y/n]: | Details<br>A value other than y or n has been entered for a [y/n] request. Enter `y` or `n`.<br>Action<br>Enter `y` or `n`. |
| KAPL10032-W | The entered value is invalid. The utility for collecting HDLM error information stops. | Details<br>Processing of the `DLMgetras` utility for collecting HDLM error information will terminate because an invalid response was sent three times to a request.<br>Action<br>Re-execute the `DLMgetras` utility. For details on the `DLMgetras` utility, see section 7.2. |
| KAPL10033-W | The file does not exist. Filename = *aa...aa* | Details<br>The file to collect does not exist.<br>*aa...aa*: Name of file to collect<br>Action<br>The file that should be obtained under normal operation does not exist. HDLM might be operating in an invalid environment. Check the system configuration. |
| KAPL10034-E | The file could not be copied. Filename = *aa...aa*, Details = *bb...bb* | Details<br>Execution of the cp command failed.<br>*aa...aa*: Name of the file you tried to copy<br>*bb...bb*: cp output message<br>Action<br>An error occurred while the collected files were being copied. The environment, in which a user executed the command, might be wrong. Check the system configuration. |
| KAPL10035-E | An attempt to archive the error information failed. Details = *aa...aa* | Details<br>Execution of the tar command failed.<br>*aa...aa*: tar command output message<br>Action<br>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<br>Execution of the `gzip` command failed.<br>*aa...aa*: `gzip` command output message<br>Action<br>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. |
| KAPL10037-W | The file does not exist or does not have executable permissions. Filename = *aa...aa* | Details<br>There is no output file, or the user does not have execution permission.<br>*aa...aa*: file name<br>Action<br>Make sure the specified file exists and the user has execution permission. |
| KAPL10038-W | Lifekeeper is not running. Output of some Lifekeeper commands is not collected. | Details<br>Output results from LifeKeeper commands could not be collected because LifeKeeper is stopped.<br>Action<br>If the information related to LifeKeeper is necessary, start LifeKeeper, and then re-execute the command. |

## 8.11 Messages from the dlmcfgmgr Utility for Managing the HDLM Configuration

| | | |
|---|---|---|
| KAPL10301-I | /sbin/dlmcfgmgr started: *aa...aa* | Details<br>The HDLM-configuration definition utility (`dlmcfgmgr`) has started.<br>*aa...aa*: date and time (for example, Fri Aug 23 19 : 12 : 50 2004)<br>Action<br>None. For details on the `dlmcfgmgr` utility, see section 7.3. |
| KAPL10302-I | /sbin/dlmcfgmgr completed normally. | Details<br>The HDLM-configuration definition utility (`dlmcfgmgr`) has completed successfully.<br>Action<br>None. For details on the `dlmcfgmgr` utility, see section 7.3. |
| KAPL10305-E | A special file could not be created. Filename = *aa...aa* | Details<br>A logical device file for an HDLM device cannot be created.<br>*aa...aa*: Logical device file name of HDLM device<br>Action<br>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 section 7.2. |
| KAPL10306-W | The configuration definition file is invalid. | Details<br>The configuration definition file (`/etc/opt/DynamicLinkManager/.dlmfdrv.conf`) is invalid, therefore, it will be created again.<br>Action<br>`/etc/opt/DynamicLinkManager/.dlmfdrv.conf` to an alias name and restart the host. The logical device file name of the HDLM device might be changed, so that after the host restarts, check the logical device file name. If the logical device file name was changed, modify the path to the upper-level program. |
| KAPL10308-W | The configuration definition file could not be created. | Details<br>The configuration definition file (`/etc/opt/DynamicLinkManager/.dlmfdrv.conf`) could not be created.<br>Action<br>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 section 7.2. |
| KAPL10309-W | The HDLM management-target disk does not exist. | Details<br>The HDLM management-target device does not exist.<br>Action<br>Check the system configuration. |

| | | |
|---|---|---|
| KAPL10312-I | HDLM has created an HDLM device special file. Device = *aa...aa* | Details<br>A logical device file for an HDLM device has been created.<br>*aa...aa*: Logical device file name of HDLM device<br>Action<br>None. |
| KAPL10313-I | HDLM has deleted an HDLM device special file. Device = *aa...aa* | Details<br>A logical device file for an HDLM device has been deleted.<br>*aa...aa*: Logical device file name of HDLM device<br>Action<br>None. |
| KAPL10314-I | HDLM has updated an HDLM device special file. Device = *aa...aa* | Details<br>A logical device file for an HDLM device has been updated.<br>*aa...aa*: Logical device file name of HDLM device<br>Action<br>None. |
| KAPL10316-E | Could not allocate memory. Size = *aa...aa* | Details<br>Securing of memory failed.<br>*aa...aa*: Target memory size (decimal number)<br>Action<br>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 section 7.2. |
| KAPL10318-E | An internal error occurred in the HDLM-configuration definition utility (dlmcfgmgr). Code = *aa...aa* Errno = *bb...bb cc...cc* | Details<br>An internal error occurred in the HDLM-configuration definition utility (`dlmcfgmgr`).<br>*aa...aa*: Line number in which the error occurred (decimal number)<br>*bb...bb*: Error number (errno) (decimal number)<br>*cc...cc*: Detailed information (optional)<br>Action<br>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 section 7.2. For details on the `dlmcfgmgr` utility, see section 7.3. |

| | | |
|---|---|---|
| KAPL10319-W | For Red Hat Enterprise Linux AS/ES3:<br>`usage: /sbin/dlmcfgmgr [-s] {-r \| -o {special-file-name ... \| all} \| -i {special-file-name ... \| all} \| -v \| -u {special-file-name ... \| all}}`<br>For Red Hat Enterprise Linux AS/ES4, SUSE LINUX Enterprise Server 9 or SUSE LINUX Enterprise Server 10:<br>`usage: /sbin/dlmcfgmgr [-s] {-r \| -o {special-file-name ... \| all} \| -i {special-file-name ... \| all} \| -v [-udev] \| -u {special-file-name ... \| all}}` | Details<br>Since the parameter of the HDLM-configuration definition utility (`dlmcfgmgr`) was inaccurate, execution of the `dlmcfgmgr` utility failed. For details on the `dlmcfgmgr` utility, see section 7.3.<br>Action<br>Specify the correct options, and then re-execute the `dlmcfgmgr` utility. |
| KAPL10320-E | The file format is invalid. File name = *aa...aa* | Details<br>The file format is invalid.<br>*aa...aa*: file name<br>Action<br>Re-execute the HDLM-configuration definition utility (`dlmcfgmgr`). 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 section 7.3. |
| KAPL10321-W | The specified HDLM device is invalid. Device = *aa...aa* | Details<br>Executing the HDLM-configuration definition utility (`dlmcfgmgr`) failed because the logical device file name of the HDLM device was incorrect.<br>*aa...aa*: Logical device file name of HDLM device<br>Action<br>Specify a correct logical device file name of an HDLM device, and then re-execute the `dlmcfgmgr` utility. For details on the `dlmcfgmgr` utility, see section 7.3. |
| KAPL10322-E | A file that defines HDLM-unmanaged disks could not be created. File name = *aa...aa*, Errno = *bb...bb* | Details<br>Execution of the HDLM-configuration definition utility (`dlmcfgmgr`) failed because an exclusion disk definition file (*aa...aa*) could not be created.<br>*aa...aa*: File name<br>*bb...bb*: Error number (decimal number)<br>Action<br>Contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM. For details on the `dlmcfgmgr` utility, see section 7.3. |

| | | |
|---|---|---|
| KAPL10323-E | The registration of the block device driver for HDLM failed. All dynamic major numbers are exhausted. | Details<br>Registration of the block device driver of HDLM failed because all dynamic major numbers were used.<br>Action<br>Unload (remove from the kernel) any unnecessary driver, and then re-execute the HDLM-configuration definition utility (`dlmcfgmgr`). For details on the `dlmcfgmgr` utility, see section 7.3. |
| KAPL10324-I | The device configuration of the system has changed. Device = *aa...aa* | Details<br>The device composition of a system to a device (*aa...aa*) was changed.<br>*aa...aa*: HDLM device name<br>Action<br>None. |
| KAPL10325-E | A regular expression cannot be established. The files will not be deleted. Errno = *aa...aa* | Details<br>A file has not been deleted because a regular expression could not be used.<br>*aa...aa*: Error number (decimal number)<br>Action<br>Contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM. |
| KAPL10326-E | You lack permission for executing the HDLM-configuration definition utility. | Details<br>The current user does not have authority to execute the HDLM-configuration definition utility (`dlmcfgmgr`).<br>Action<br>Re-execute the `dlmcfgmgr` utility as a user with root permission. For details on the `dlmcfgmgr` utility, see section 7.3. |
| KAPL10327-W | The entered value is invalid. | Details<br>The value input into the execution check message of the HDLM-configuration definition utility (`dlmcfgmgr`) is invalid.<br>Action<br>Enter the correct y/n value after the KAPL10339-I message is output. For details on the `dlmcfgmgr` utility, see section 7.3. |
| KAPL10328-E | The entered value is invalid. The HDLM-configuration definition utility processing will now stop. | Details<br>An invalid value was entered 3 or more times for the message confirming the execution of the HDLM-configuration definition utility (`dlmcfgmgr`). Processing is interrupted.<br>Action<br>Re-execute the `dlmcfgmgr` utility. For details on the `dlmcfgmgr` utility, see section 7.3. |
| KAPL10329-E | A file could not be opened. File name = *aa...aa*, Errno = *bb...bb* | Details<br>The file could not be opened.<br>*aa...aa*: File name<br>*bb...bb*: Error number (decimal number)<br>Action<br>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. |

| | | |
|---|---|---|
| KAPL10330-E | A symbolic link cannot be created. File name = *aa...aa*, Errno = *bb...bb* | Details<br>An attempt to create a symbolic link has failed.<br>*aa...aa*: File name<br>*bb...bb*: Error number (decimal number)<br>Action<br>Contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM. |
| KAPL10331-E | The registering of the block device driver for HDLM failed. Errno = *aa...aa* | Details<br>An attempt to register the block device driver of HDLM has failed.<br>*aa...aa*: Error number (decimal number)<br>Action<br>Contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM. |
| KAPL10332-E | The registering of the partition information for HDLM device(s) failed. Errno = *aa...aa* | Details<br>An attempt to register the partition information on HDLM has failed.<br>*aa...aa*: Error number (decimal number)<br>Action<br>Contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM. |
| KAPL10333-E | An internal error occurred in the file operation. File name = *aa...aa* | Details<br>An internal error occurred during file operation.<br>*aa...aa*: File name<br>Action<br>Contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM. |
| KAPL10334-E | Ioctl *aa...aa* failed with errno = *bb...bb*. | Details<br>An attempt to issue an IOCTL call has failed.<br>*aa...aa*: IOCTL name<br>*bb...bb*: Error number (decimal number)<br>Action<br>Contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM. |
| KAPL10335-E | A directory cannot be created. Directory = *aa...aa*, Errno = *bb...bb* | Details<br>An attempt to create a directory has failed.<br>*aa...aa*: Directory name<br>*bb...bb*: Error number (decimal number)<br>Action<br>Confirm the access permission. If there is no problem, contact your HDLM vendor or maintenance company if you have a maintenance contract for HDLM. |
| KAPL10336-W | A file cannot be deleted. File name = *aa...aa* | Details<br>An attempt to delete a file has failed.<br>*aa...aa*: File name<br>Action<br>Delete the file shown in File name. |

| | | |
|---|---|---|
| KAPL10337-W | A special device exists with a different major number. The device will now be unlinked. Device = *aa...aa* | Details<br>The device has unlinked because major numbers in the logical device files do not match.<br>*aa...aa*: Logical device file name of the HDLM device<br>Action<br>Execute the HDLM-configuration definition utility (`dlmcfgmgr`) with the `-v` option, and then check the configuration of the HDLM device. For details on the `dlmcfgmgr` utility, see section 7.3. |
| KAPL10338-W | A directory cannot be deleted. Directory = *aa...aa* | Details<br>An attempt to delete a directory has failed.<br>*aa...aa*: Directory name<br>Action<br>Delete the directory shown in Directory. |
| KAPL10339-I | This operation will change the configuration of HDLM devices.<br>Do you want to continue? [y/n]: | Details<br>This message checks whether you want to change the configuration definition of an HDLM device.<br>Action<br>Enter `y` to continue. Enter `n` to cancel. |
| KAPL10340-E | Several processes failed. See the *aa...aa* file. | Details<br>Some processes in the execution of the HDLM-configuration definition utility (`dlmcfgmgr`) have failed. Check the file *aa...aa*.<br>*aa...aa*: File name<br>Action<br>See the `/var/opt/DynamicLinkManager/log/dlmcfgmgr1.log` file and check the message output before this message. For details on the `dlmcfgmgr` utility, see section 7.3. |
| KAPL10341-I | The HDLM device configurations have been changed. | Details<br>The configuration of a HDLM device was changed by execution of the HDLM-configuration definition utility (`dlmcfgmgr`).<br>Action<br>None. For details on the `dlmcfgmgr` utility, see section 7.3. |
| KAPL10343-I | HDLM has detected and registered a new HDLM device. HDLM device = *aa...aa* | Details<br>By executing the HDLM-configuration definition utility (`dlmcfgmgr`), a new HDLM device has been detected and registered into the HDLM device configuration definition.<br>*aa...aa*: Logical device file name of HDLM device<br>Action<br>None. For details on the `dlmcfgmgr` utility, see section 7.3. |

| | | |
|---|---|---|
| KAPL10344-I | HDLM has detected and registered a new path to an already registered HDLM device. HDLM device = *aa...aa*, Device = *bb...bb* (Host: *cc...cc* Channel: *dd...dd* ID: *ee...ee* Lun: *ff...ff*) | Details<br>By executing the HDLM-configuration definition utility (`dlmcfgmgr`), a newly detected path has been added to the registered HDLM device.<br>*aa...aa*: HDLM device name<br>*bb...bb*: SCSI device name<br>*cc...cc*: Host ID (decimal number)<br>*dd...dd*: Channel ID (decimal number)<br>*ee...ee*: Target ID (decimal number)<br>*ff...ff*: LUN (decimal number)<br>Action<br>None. For details on the `dlmcfgmgr` utility, see section 7.3. |
| KAPL10345-I | HDLM has unregistered the existing path to an HDLM device. HDLM device = *aa...aa*, Device = *bb...bb* (Host: *cc...cc* Channel: *dd...dd* ID: 0 Lun: *ee...ee*) | Details<br>By executing the HDLM-configuration definition utility (`dlmcfgmgr`), the existing path has been unregistered from the registered HDLM device.<br>*aa...aa*: HDLM device name<br>*bb...bb*: SCSI device name<br>*cc...cc*: Host ID (decimal number)<br>*dd...dd*: Channel ID (decimal number)<br>*ee...ee*: Target ID (fixed value: `0`) (decimal number)<br>*ff...ff*: LUN (decimal number)<br>Action<br>None. For details on the `dlmcfgmgr` utility, see section 7.3. |
| KAPL10346-I | HDLM has unregistered an existing HDLM device. HDLM device = *aa...aa* | Details<br>By executing the HDLM-configuration definition utility (`dlmcfgmgr`), the registered HDLM device has been unregistered.<br>*aa...aa*: Logical device file name of HDLM device<br>Action<br>None. For details on the `dlmcfgmgr` utility, see section 7.3. |
| KAPL10347-I | The HDLM device is no longer under management. HDLM device = *aa...aa* | Details<br>By executing the HDLM-configuration definition utility (`dlmcfgmgr`), the HDLM device has been removed from HDLM management.<br>*aa...aa*: Logical device file name of HDLM device<br>Action<br>None. For details on the `dlmcfgmgr` utility, see section 7.3. |
| KAPL10348-I | The HDLM device is now under management. HDLM device = *aa...aa* | Details<br>By executing the HDLM-configuration definition utility (`dlmcfgmgr`), the HDLM device has been placed under HDLM management.<br>*aa...aa*: Logical device file name of HDLM device<br>Action<br>None. For details on the `dlmcfgmgr` utility, see section 7.3. |

| | | |
|---|---|---|
| KAPL10349-I | HDLM will remove unavailable paths or devices. HDLM device = *aa...aa* | Details<br>By executing the HDLM-configuration definition utility (`dlmcfgmgr`), a path and a logical device file of an HDLM device has been deleted.<br>*aa...aa*: Logical device file name of HDLM device<br>Action<br>None. For details on the `dlmcfgmgr` utility, see section 7.3. |
| KAPL10350-E | The HDLM driver(s) is not loaded. | Details<br>An attempt to execute the HDLM-configuration definition utility (`dlmcfgmgr`) failed because the HDLM driver was not loaded.<br>Action<br>Load (install in the kernel) an HDLM driver, and then re-execute the `dlmcfgmgr` utility. For details on the `dlmcfgmgr` utility, see section 7.3. |
| KAPL10351-E | The specified HDLM-configuration definition utility is invalid. Operation name = *aa...aa*, Parameter = *bb...bb* | Details<br>An attempt to execute the HDLM-configuration definition utility (`dlmcfgmgr`) failed because the specified operation name (*aa...aa*:) or parameter (*bb...bb*) was invalid.<br>*aa...aa*: Operation name<br>*bb...bb*: Parameter<br>Action<br>Confirm that the operation name and parameter are correct, and then re-execute the `dlmcfgmgr` utility. For details on the `dlmcfgmgr` utility, see section 7.3. |
| KAPL10352-W | The HDLM device is in use. HDLM device = *aa...aa* | Details<br>An attempt to execute the HDLM-configuration definition utility (`dlmcfgmgr`) failed because the specified HDLM device (*aa...aa*) was using it.<br>*aa...aa*: HDLM device name<br>Action<br>Stop the application using the HDLM device, unmount the HDLM device being mounted, and then re-execute the `dlmcfgmgr` utility. For details on the `dlmcfgmgr` utility, see section 7.3. |
| KAPL10353-W | An error occurred while a backup of the HDLM files was being acquired. File name = *aa...aa* | Details<br>An error occurred during acquisition of the HDLM file (*aa...aa*) backup.<br>*aa...aa*: File name<br>Action<br>Contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM. |
| KAPL10354-E | An attempt to unregister a block device driver for HDLM failed. | Details<br>An attempt to execute the HDLM-configuration definition utility (`dlmcfgmgr`) failed because registration of a HDLM block device driver could not be released.<br>Action<br>Stop the application for using HDLM device, and unmount the mounting HDLM device, and re-execute the `dlmcfgmgr` utility. For details on the `dlmcfgmgr` utility, see section 7.3. |

| | | |
|---|---|---|
| KAPL10355-E | An attempt to unregister partition information for the HDLM device(s) failed. | Details<br>An attempt to execute the HDLM-configuration definition utility (`dlmcfgmgr`) failed because registration of the partition information on a HDLM device could not be released.<br>Action<br>Stop the application for using HDLM device, and unmount the mounting HDLM device, and re-execute the `dlmcfgmgr` utility. For details on the `dlmcfgmgr` utility, see section 7.3. |
| KAPL10356-E | An unused major number could not be released. Major number = *aa...aa* | Details<br>An attempt to release an intact major number (*aa...aa*) has failed.<br>*aa...aa*: Major number (decimal number)<br>Action<br>Restart the host. If the host does not recover, contact your HDLM vendor or maintenance company if you have a maintenance contract for HDLM. |
| KAPL10357-E | The maximum number of HDLM device files has been reached. Maximum number = *aa...aa* | Details<br>An attempt to execute the HDLM-configuration definition utility (`dlmcfgmgr`) failed because the maximum number of logical device files has been created.<br>*aa...aa*: Maximum number of created logical device files of an HDLM device (decimal number)<br>Action<br>Revise the configuration of the HDLM device to reduce the number of HDLM devices used, and then re-execute the `dlmcfgmgr` utility. For details on the `dlmcfgmgr` utility, see section 7.3. |
| KAPL10358-E | The device name is too long. Maximum length = *aa...aa* | Details<br>An attempt to execute the HDLM-configuration definition utility (`dlmcfgmgr`) failed because the name of the logical device file is too long.<br>*aa...aa*: Maximum length of characters for a logical device file name (decimal)<br>Action<br>Check that the logical device file name of the HDLM device is correct, and then re-execute the `dlmcfgmgr` utility. For details on the `dlmcfgmgr` utility, see section 7.3. |
| KAPL10359-E | Multiple instances of the HDLM-configuration definition utility cannot be executed concurrently. | Details<br>An attempt to execute `dlmcfgmgr` failed because two or more instances of `dlmcfgmgr` were executed simultaneously.<br>Action<br>Wait until the running HDLM-configuration definition utility (`dlmcfgmgr`) ends, and then re-execute the HDLM-configuration definition utility (`dlmcfgmgr`). For details on the `dlmcfgmgr` utility, see section 7.3. |

| | | |
|---|---|---|
| KAPL10360-I | HDLM has activated a path for an HDLM device. HDLM device = *aa...aa*, Device = *bb...bb* (Host: *cc...cc* Channel: *dd...dd* ID: *ee...ee* Lun: *ff...ff*) | Details<br>By executing the HDLM-configuration definition utility (`dlmcfgmgr`), the path to the HDLM device has been made active.<br>*aa...aa*: HDLM device name<br>*bb...bb*: SCSI device name<br>*cc...cc*: Host ID (decimal number)<br>*dd...dd*: Channel ID (decimal number)<br>*ee...ee*: Target ID (decimal number)<br>*ff...ff*: LUN (decimal number)<br>Action<br>None. For details on the `dlmcfgmgr` utility, see section 7.3. |
| KAPL10361-I | HDLM has deactivated a path for the HDLM device. HDLM device = *aa...aa*, Device = *bb...bb* (Host: *cc...cc* Channel: *dd...dd* ID: 0 Lun: *ee...ee*) | Details<br>By executing the HDLM-configuration definition utility (`dlmcfgmgr`), the path to the HDLM device has been made inactive.<br>*aa...aa*: HDLM device name<br>*bb...bb*: SCSI device name<br>*cc...cc*: Host ID (decimal number)<br>*dd...dd*: Channel ID (decimal number)<br>*ee...ee*: Target ID (decimal number)<br>*ff...ff*: LUN (decimal number)<br>Action<br>None. For details on the `dlmcfgmgr` utility, see section 7.3. |
| KAPL10362-W | HDLM has detected an active path that cannot be deleted. HDLM device = *aa...aa*, Device = *bb...bb* (Host: *cc...cc* Channel: *dd...dd* ID: *ee...ee* Lun: *ff...ff*) | Details<br>An attempt to execute the HDLM-configuration definition utility (`dlmcfgmgr`) failed because a path to the specified HDLM device was being used.<br>*aa...aa*: HDLM device name<br>*bb...bb*: SCSI device name<br>*cc...cc*: Host ID (decimal number)<br>*dd...dd*: Channel ID (decimal number)<br>*ee...ee*: Target ID (decimal number)<br>*ff...ff*: LUN (decimal number)<br>Action<br>Stop the application that uses the HDLM device, unmount the HDLM device being mounted, and then re-execute the `dlmcfgmgr` utility. For details on the `dlmcfgmgr` utility, see section 7.3. |
| KAPL10363-W | A data mismatch was found. HDLM device = *aa...aa* | Details<br>The HDLM device detected an unsuitable combination of data.<br>*aa...aa*: Logical device file name of HDLM device<br>Action<br>Contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM. |

| | | |
|---|---|---|
| KAPL10364-I | The LDEV for an HDLM device has changed. HDLM will now change the configuration. HDLM device = *aa...aa* | Details<br>The configuration has been modified because an LDEV corresponding to an HDLM device was changed.<br>*aa...aa*: Logical device file name of an HDLM device<br>Action<br>None. |
| KAPL10365-W | The LDEV for an HDLM device has changed. The processing for that HDLM device will be cancelled. HDLM device = *aa...aa* | Details<br>The processing for the HDLM device was interrupted because an LDEV corresponding to an HDLM device was changed.<br>*aa...aa*: The logical device file name of the HDLM device<br>Action<br>Execute the HDLM-configuration definition utility (`dlmcfgmgr`) with the `-v` option, and then check the configuration of the HDLM device. For details on the `dlmcfgmgr` utility, see section 7.3. |
| KAPL10366-I | Unsupported SCSI device was found. Device = *aa...aa* | Details<br>The SCSI device output in this message is not supported.<br>*aa...aa*: Logical device file name of HDLM device<br>Action<br>None. |
| KAPL10911-E | Processing for character devices is stopped due to internal error. | Details<br>Processing related to character devices failed because of an internal error.<br>Action<br>A file related to HDLM might be corrupted. Please re-install HDLM. |
| KAPL10912-I | Failure while opening sg device since the disk does not exist. | Details<br>An attempt to open a character device failed because the disk does not exist.<br>Action<br>Make sure that the SCSI device exists and no path is disconnected, and then re-execute the command. |
| KAPL10913-E | Error occurred in ioctl (SCSI_IOCTL_GET_IDLUN) execution. | Details<br>An error occurred because `ioctl` was executed.<br>Action<br>Make sure that the SCSI device exists, and then reload the filter driver, and then re-execute the command. |
| KAPL10914-E | HDLM character driver is not loaded. | Details<br>A character driver of HDLM has not been loaded.<br>Action<br>Reload the filter driver, and then re-execute the command. |

| | | |
|---|---|---|
| KAPL10915-E | HDLM character driver is loaded. | Details<br>A character driver of HDLM has been loaded.<br>Action<br>If a message was output during system startup, reconfigure HDLM in the following order:<br>1. Stop DLMManager.<br>2. Unload the filter driver.<br>3. Unload the alert driver.<br>4. Load the alert driver.<br>5. Load the filter driver.<br>6. Execute `dlmcfgmgr -r -s`.<br>7. Start DLMManager. |
| KAPL10916-W | Failed to update sg information for path sd = *aa...aa*. | Details<br>An attempt to update the sg device information related to the path (sd = *aa...aa*) has failed.<br>*aa...aa*: SCSI device file name<br>Action<br>Reload the filter driver, and then re-execute the command. |
| KAPL10917-I | Sg information is updated for path sd = *aa...aa*. | Details<br>The sg device information related to the path (sd = *aa...aa*) has been updated.<br>*aa...aa*: SCSI device file name<br>Action<br>None. |
| KAPL10918-W | Network interface not configured. Processing for character devices is stopped. | Details<br>Processing related to the character device stopped because the network interface is not configured.<br>Action<br>Make sure the network has been set up and the service is running, reload the filter driver, and then re-execute the command. |
| KAPL10919-I | A special device exists with a different major/minor number or type. Device = *aa...aa*. | Details<br>*aa...aa*: Device file name (character string)<br>Action<br>There is no impact on operations, so use HDLM as is. |
| KAPL10920-W | Processing for HDLM Config Manager interrupted. | Details<br>Processing of the HDLM-configuration definition utlity (`dlmcfgmgr`) was interrupted.<br>Action<br>Please re-execute the command. |

## 8.12 Messages from the Utility for Clearing HDLM Persistent Reservation

| | | |
|---|---|---|
| KAPL10641-I | Reservation Key will now be cleared. Is this OK? [y/n]: | Details<br>Enter `y` to clear and `n` to not clear the Reservation Key.<br>Action<br>None. |
| KAPL10642-I | Reservation Key of *aa...aa* was cleared. | Details<br>The Reservation Key has been cleared.<br>*aa...aa*: Logical device file name for the HDLM management-target device<br>Action<br>None. |
| KAPL10643-W | A necessary parameter is not specified. | Details<br>A parameter is not specified for the `dlmpr` utility.<br>Action<br>Execute the `dlmpr -h` utility to check the parameter, and then retry execution. For details on the `dlmpr` utility, see section 7.6. |
| KAPL10644-W | The specified parameters cannot be specified at the same time. parameter = *aa...aa* | Details<br>The specified parameters cannot be specified for the `dlmpr` utility at the same time.<br>*aa...aa*: Specified parameter<br>Action<br>Execute the `dlmpr -h` utility to check the parameter, and then retry execution. For details on the `dlmpr` utility, see section 7.6. |
| KAPL10645-W | A parameter value is invalid. parameter = *aa...aa* | Details<br>An invalid parameter value has been specified for the `dlmpr` utility.<br>*aa...aa*: Specified parameter<br>Action<br>Specify the correct value for the parameter, and then retry. For details on the `dlmpr` utility, see section 7.6. |
| KAPL10646-W | A parameter is invalid. parameter = *aa...aa* | Details<br>An invalid parameter has been specified for the `dlmpr` utility.<br>*aa...aa*: Specified parameter<br>Action<br>Execute help of the `dlmpr` utility to check the parameters that can be specified, and then retry. For details on the `dlmpr` utility, see section 7.6. |

| | | |
|---|---|---|
| KAPL10648-E | An internal error occurred in the dlmpr utility. Error Code = *aa...aa* | Details<br>An error not caused by the user has occurred in the `dlmpr` utility.<br>*aa...aa*: Error number (character string)<br>Action<br>When the error code is 1: Make sure that there is sufficient memory.<br>When the error code is 2: Contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM.<br>When the error code is 3: Make sure that a network card is installed.<br>When the error code is 4 or 5: Make sure that the SCSI device file permissions permit reading and writing. For details on the `dlmpr` utility, see section 7.6. |
| KAPL10649-E | *aa...aa* : An attempt to perform Reservation Key clear processing has failed. | Details<br>An attempt to perform Reservation Key clear processing has failed.<br>*aa...aa*: HDLM device (dlmfdrv*n*)<br>Action<br>Contact your HDLM vendor or the maintenance company if there is a maintenance contract for HDLM. |
| KAPL10650-I | *aa...aa* : NO RESERVATION | Details<br>A LU has not been reserved.<br>*aa...aa*: HDLM device (dlmfdrv*n*)<br>Action<br>None. |
| KAPL10651-I | The user terminated the operation. | Details<br>The `dlmpr` utility has been terminated because `n` was sent to a request.<br>Action<br>None. |
| KAPL10652-E | The entered value is invalid. The operation stops. | Details<br>An invalid response was sent three times consecutively to a request.<br>Action<br>Re-execute the `dlmpr` utility. |
| KAPL10653-W | The entered value is invalid. Please re-enter it [y/n]: | Details<br>A value other than `y` or `n` has been entered for a [y/n] request.<br>Action<br>Enter `y` or `n`. |

| | | |
|---|---|---|
| KAPL10665-I | The dlmpr utility completed. | Details<br>The `dlmpr` utility completed normally.<br>Action<br>None. For details on the `dlmpr` utility, see section 7.6. |
| KAPL10922-E | The version of the kernel supported by the installed HDLM does not match the currently running kernel version. | Details<br>The version of the kernel to which installed HDLM has adjusted and the kernel that is running now is not corresponding. There is a possibility of installing the package of the kernel after installing HDLM.<br>Action<br>Install HDLM that adjusts to the kernel that is running now. |
| KAPL10923-E | The HDLM file is invalid. File name = *aa...aa* | Details<br>A file necessary for processing HDLM cannot be found. Alternatively, the content of a file is invalid.<br>*aa...aa*: Name of file where error was detected.<br>Action<br>Reinstall an HDLM. |

## 8.13 Messages from HDLM remote access interface

| | | |
|---|---|---|
| KAPL11901-I | *aa...aa* has started. | Details<br>The operation has started on the host.<br>*aa...aa*: Operation (character string)<br>▪ Get Path Information<br>▪ Get Option Information<br>▪ Set Option Information<br>▪ Clear Data<br>▪ Get HDLM Manager Status<br>▪ Get HDLM Driver Status<br>▪ Get HDLM Alert Driver Status<br>Action<br>None. |
| KAPL11902-I | *aa...aa* has started. PathID = *bb...bb* | Details<br>The operation has started on the host.<br>*aa...aa*: Operation (character string)<br>▪ Online<br>▪ Offline<br>*bb...bb*: The Path ID of the target path for the operation (decimal number)<br>Action<br>None. |
| KAPL11903-I | *aa...aa* has completed normally. | Details<br>Operation has completed normally on the host.<br>*aa...aa*: Operation (character string)<br>▪ Get Path Information<br>▪ Get Option Information<br>▪ Set Option Information<br>▪ Clear Data<br>▪ Get HDLM Driver Status<br>▪ Get HDLM Manager Status<br>▪ Get HDLM Alert Driver Status<br>▪ Online<br>▪ Offline<br>▪ Get SNMP Trap Information<br>▪ Set SNMP Trap Information<br>▪ Set LU Load Balance<br>Action<br>None. |

| | | |
|---|---|---|
| KAPL11904-E | *aa...aa* has completed abnormally. Error status = *bb...bb* | Details<br>Operation has completed abnormally on the host.<br>*aa...aa*: Operation (character string)<br>▪ Get Path Information<br>▪ Get Option Information<br>▪ Set Option Information<br>▪ Clear Data<br>▪ Get HDLM Driver Status<br>▪ Get HDLM Manager Status<br>▪ Get HDLM Alert Driver Status<br>▪ Online<br>▪ Offline<br>▪ Get SNMP Trap Information<br>▪ Set SNMP Trap Information<br>▪ Set LU Load Balance<br>*bb...bb*: Error status returned from the API (character string)<br>Action<br>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 section 7.2. |
| KAPL11905-E | An unexpected error occurred. | Details<br>An exception occurred during processing in the host.<br>Action<br>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 section 7.2. |
| KAPL11906-I | GUI information - *aa...aa* | Details<br>This information is required for determining the cause of the problem (if any).<br>*aa...aa*: Trace information (character string)<br>Action<br>None. |
| KAPL11907-I | XML reception - *aa...aa* | Details<br>This information is required for determining the cause of the problem (if any).<br>*aa...aa*: XML information (character string)<br>Action<br>None. |

| | | |
|---|---|---|
| KAPL11908-I | XML transmission - *aa...aa* | Details<br>This information is required for determining the cause of the problem (if any).<br>*aa...aa*: XML information (character string)<br>Action<br>None. |

## 8.14 Messages from the dlmlfk Utility for Supporting LifeKeeper

| | | |
|---|---|---|
| KAPL12001-I | The utility for supporting LifeKeeper *aa...aa* started. | Details<br>*aa...aa*: Option when the utility for supporting LifeKeeper (`dlmlfk`) was executed<br>Action<br>None. |
| KAPL12002-I | The utility for supporting LifeKeeper *aa...aa* is successful. | Details<br>*aa...aa*: Option when the utility for supporting LifeKeeper (`dlmlfk`) was executed<br>Action<br>None. |
| KAPL12003-I | Support change already exist in *aa...aa* | Details<br>*aa...aa*: Name of the file that an attempt was made to change<br>Action<br>None. |
| KAPL12004-I | The utility for supporting LifeKeeper *aa...aa* failed. | Details<br>*aa...aa*: Option when the utility for supporting LifeKeeper (`dlmlfk`) was executed<br>Action<br>Check for the reason of failure. Error messages will be displayed. For details on the `dlmlfk` utility, see section 7.4. |
| KAPL12005-I | Directory does not exist. Directory = *aa...aa* | Details<br>*aa...aa*: Referenced directory<br>Action<br>None. |
| KAPL12006-I | HDLM mapping information could not be found in the Lifekeeper disk.mapping directory. | Action<br>None. |
| KAPL12007-I | Please use *aa...aa* to restore *bb...bb* | Details<br>*aa...aa*: Backup file name<br>*bb...bb*: Name of file to be restored<br>Action<br>From the backup directory, manually copy the file *aa...aa* to the file *bb...bb*. The paths are displayed in the message. |
| KAPL12008-E | The file does not contain default entries. File name = *aa...aa* | Details<br>*aa...aa*: Name of the device_pattern file for LifeKeeper<br>Action<br>Confirm the device_pattern file, and re-execute the utility for supporting LifeKeeper (`dlmlfk`). |

| | | |
|---|---|---|
| KAPL12009-E | Could not modify file. File name = *aa...aa*, Errno = *bb...bb* | Details<br>*aa...aa*: Name of file to be updated<br>*bb...bb*: Detailed error code (decimal number)<br>Action<br>Based on the error number, resolve the cause of the problem. |
| KAPL12010-I | Execution of the utility for supporting LifeKeeper *aa...aa* is aborted by user. | Details<br>*aa...aa*: Option when the utility for supporting LifeKeeper (`dlmlfk`) was executed<br>Action<br>None. |
| KAPL12011-E | Abort processing for the utility for supporting LifeKeeper *aa...aa* completed with errors. | Details<br>*aa...aa*: Option when the utility for supporting LifeKeeper (`dlmlfk`) was executed<br>Action<br>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 section 7.2. |
| KAPL12012-E | Invalid parameter passed to the utility for supporting LifeKeeper. Parameter = *aa...aa* | Details<br>*aa...aa*: Option when the utility for supporting LifeKeeper (`dlmlfk`) was executed<br>Action<br>Re-execute with the correct parameter. For details on the `dlmlfk` utility, see section 7.4. |
| KAPL12013-E | This script can only be executed with the name dlmlfk. | Action<br>Change the name of the script to `dlmlfk`. |
| KAPL12014-E | Necessary information could not be obtained. File name = *aa...aa* | Details<br>*aa...aa*: Referenced file name<br>Action<br>Check the *aa...aa* file, and re-execute. |
| KAPL12015-E | LifeKeeper is running. | Action<br>Stop LFK and re-execute `dlmlfk`. For details on the `dlmlfk` utility, see section 7.4. |
| KAPL12016-E | Necessary file or directory does not exist. File or directory name = *aa...aa* | Details<br>*aa...aa*: Referenced directory or file name<br>Action<br>Make sure the directory and file exist:<br>1. For `/etc/opt/DynamicLinkManager`, create it and reinstall HDLM.<br>2. For `device_pattern`, `lk_cspec` and `disk.mapping`, re-install LFK. |

| | | |
|---|---|---|
| KAPL12017-E | Failed to copy file. File name = *aa...aa*, Errno = *bb...bb* | Details<br>*aa...aa*: Backup file name<br>*bb...bb*: Detailed error code (decimal number)<br>Action<br>Check for the reason of cp the command failure (OS specific). |
| KAPL12018-E | Failed to create directory. Directory = *aa...aa*, Errno = *bb...bb* | Details<br>*aa...aa*: Directory that an attempt was made to create<br>*bb...bb*: Detailed error code (decimal number)<br>Action<br>Check for the reason of the mkdir command failure (OS specific). |
| KAPL12019-W | Could not modify file. File name = *aa...aa*, Errno = *bb...bb* | Details<br>*aa...aa*: Updated file name<br>*bb...bb*: Detailed error code (decimal number)<br>Action<br>From the backup directory, manually restore the file for which unconfiguration failed. |
| KAPL12020-E | LifeKeeper is not installed. | Action<br>Install LFK and then execute `dlmlfk`. For details on the `dlmlfk` utility, see section 7.4. |
| KAPL12021-E | Could not find HDLM character special file. File name = *aa...aa* | Details<br>*aa...aa*: HDLM character special file name<br>Action<br>Execute the HDLM-configuration definition utility (`dlmcfgmgr`) with `-r`, and then restart LifeKeeper. For details on the `dlmcfgmgr` utility, see section 7.3. |
| KAPL12022-E | File does not have executable permission. File name = *aa...aa* | Details<br>*aa...aa*: Execution file name<br>Action<br>Assign executable permissions to the script executer. Alternatively:<br>1. For `lktest`, re-install LFK.<br>2. For `dlmlfk`, re-install HDLM. |
| KAPL12023-E | Multiple instances of the utility for supporting LifeKeeper cannot be executed concurrently. | Action<br>Wait until execution of the utility for supporting LifeKeeper (`dlmlfk`) ends, and then re-execute it. For details on the `dlmlfk` utility, see section 7.4. |
| KAPL12024-W | Usage: /opt/DynamicLinkManager/bin/dlmlfk {-c\| -u} | Action<br>Make sure the parameter is correct. |

| | | |
|---|---|---|
| KAPL12025-W | Failed to remove directory. Directory name = *aa...aa*, Errno = *bb...bb* | Details<br>*aa...aa*: Deleted file name<br>*bb...bb*: Detailed error code (decimal number)<br>Action<br>Check whether the file exists. If it does, delete it. |
| KAPL12026-E | Execution of dlmlfk to configure disk.mapping files for HDLM support failed. | Action<br>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 section 7.2. |
| KAPL12035-E | The version of the kernel supported by the installed HDLM does not match the currently running kernel version. | Details<br>The version of the kernel to which installed HDLM has adjusted and the kernel that is running now is not corresponding. There is a possibility of installing the package of the kernel after installing HDLM.<br>Action<br>Install HDLM that adjusts to the kernel that is running now. |
| KAPL12036-E | The HDLM file is invalid. File name = *aa...aa* | Details<br>A file necessary for processing HDLM cannot be found. Alternatively, the content of a file is invalid.<br>*aa...aa*: Name of file where error was detected.<br>Action<br>Reinstall an HDLM. |
| KAPL12051-E | The version of the kernel supported by the installed HDLM does not match the currently running kernel version. | Details<br>The version of the kernel to which installed HDLM has adjusted and the kernel that is running now is not corresponding. There is a possibility of installing the package of the kernel after installing HDLM.<br>Action<br>Install HDLM that adjusts to the kernel that is running now. |
| KAPL12052-E | The HDLM file is invalid. File name = *aa...aa* | Details<br>A file necessary for processing HDLM cannot be found. Alternatively, the content of a file is invalid.<br>*aa...aa*: Name of file where error was detected.<br>Action<br>Reinstall an HDLM. |

## 8.15 Messages from the dlmmkinitrd Utility for Supporting Boot Disk

| | | |
|---|---|---|
| KAPL12301-E | This script can only be executed with the name dlmmkinitrd. | Action<br>Please rename the script to `dlmmkinitrd`. |
| KAPL12302-E | Multiple instances of the utility for supporting the boot disk cannot be executed concurrently. | Action<br>Wait until the `dlmmkinitrd` utility for supporting the boot disk ends, and then re-execute the `dlmmkinitrd` utility. For details on the `dlmmkinitrd` utility, see section 7.5. |
| KAPL12303-E | A file or directory does not exist. File name = *aa...aa* | Details<br>*aa...aa*: The name of the non-existent file<br>Action<br>If the file name displayed in the message is the file name specified in the fstab parameter of the Boot Disk Support Utility (`dlmmkinitrd`)<br>Check the specified file name, specify the file name of the existing fstab file, and then re-execute the utility.<br>If the file name displayed in the message is the HDLM file name shown below, HDLM may not be installed correctly. In this situation, re-install HDLM.<br>In the list shown below, kernel-version indicates the value displayed by the `uname -r` command.<br>▪ `/etc/opt/DynamicLinkManager/dlmmkinitrd.conf`<br>▪ `/lib/modules/kernel-version/kernel/drivers/DynamicLinkManager/sddlmfdrv.o`<br>▪ `/lib/modules/kernel-version/kernel/drivers/DynamicLinkManager/sddlmfdrv.ko`<br>▪ `/lib/modules/kernel-version/kernel/drivers/DynamicLinkManager/sddlmadrv.o`<br>▪ `/lib/modules/kernel-version/kernel/drivers/DynamicLinkManager/sddlmadrv.ko`<br>If the file name displayed in the message is another file name<br>The OS (kernel) or the HBA driver may not be installed correctly. Confirm sure the existence of the files. |
| KAPL12305-I | The utility for supporting the boot disk was stopped by the user. | Action<br>Retry the `dlmmkinitrd` utility for supporting the boot disk. |
| KAPL12306-I | The list of files required in the initial ramdisk image is being created. | Action<br>None. |
| KAPL12307-I | The initial ramdisk file system is being created. | Action<br>None. |
| KAPL12308-I | The initial ramdisk file system is being mounted. Mount point = *aa...aa* | Details<br>*aa...aa*: Mount point to be used<br>Action<br>None. |
| KAPL12309-I | Required files are being copied to the initial ramdisk. | Action<br>None. |
| KAPL12310-I | The linuxrc executable is being created. | Action<br>None. |

| | | |
|---|---|---|
| KAPL12311-I | The fstab file for the initial ramdisk is being created. | Action<br>None. |
| KAPL12312-I | The directories for the initial ramdisk are being created. | Action<br>None. |
| KAPL12313-I | Initial ramdisk is being unmounted. | Action<br>None. |
| KAPL12314-I | The compressed initial ramdisk image is being created. | Action<br>None. |
| KAPL12315-I | A compressed initial ramdisk image *aa...aa* was created with the ramdisk size = *bb...bb* KB. | Details<br>*aa...aa*: The created initial ramdisk image<br>*bb...bb*: The ramdisk size required for the created initrd image (decimal number)<br>Action<br>None. |
| KAPL12316-I | The SCSI root partition is *aa...aa* | Details<br>*aa...aa*: The SCSI device partition on which the root is mounted<br>Action<br>None. |
| KAPL12318-W | Usage: /opt/DynamicLinkManager/bin/dlmmkinitrd [-v] [-f] [-fstab fstab-name] initrd-image kernel-version | Action<br>Confirm the parameters specified when the `dlmmkinitrd` utility for supporting the boot disk was executed, and then retry. For details on the `dlmmkinitrd` utility, see section 7.5. |
| KAPL12319-E | The utility for supporting the boot disk has failed. | Action<br>This message is always preceded by some other specific error message. |
| KAPL12320-E | A file or directory already exists. File name = *aa...aa* | Details<br>*aa...aa*: The name of the already existing file<br>Action<br>To the file name of the initial RAM disk image file specified in the parameter of the Boot Disk Support Utility (`dlmmkinitrd`), do not specify an existing file name. However, if overwriting the existing file, specify the `-f` parameter.<br>Note that if you attempt to specify an existing directory name as the initial RAM disk image file, it is impossible to overwrite the file even if the `-f` parameter is specified. To overwrite the file, it is necessary to delete the directory first. |
| KAPL12321-E | An invalid root partition is specified in the file. File name = *aa...aa* | Details<br>*aa...aa*: The name of the file containing the invalid root partition entry<br>Action<br>Change the root partition specified in the `/etc/fstab` file to a SCSI or HDLM device, or specify settings so that HDLM manages the device that is listed next to the root partition specified in the `/etc/fstab` file. |
| KAPL12322-I | Execution of linuxrc started. | Action<br>None. |

| | | |
|---|---|---|
| KAPL12323-I | The insertion of the module was started. Module name = *aa...aa* | Details<br>*aa...aa*: The name of the module being inserted<br>Action<br>None. |
| KAPL12324-E | The module could not be inserted. Module name = *aa...aa* | Details<br>*aa...aa*: The name of the module whose insertion has failed<br>Action<br>Confirm that the system environments are not changed after creating the initial ramdisk image file.<br>Check whether there is a mistake in the initial ramdisk image file specified in the config file of the boot loader. |
| KAPL12325-I | The command started. Command name = *aa...aa* | Details<br>*aa...aa*: The started command<br>Action<br>None. |

| | | |
|---|---|---|
| KAPL12326-E | The command could not execute.<br>Command name = *aa...aa* | Details<br>*aa...aa*: The failed command<br>Action<br>Perform the following confirmation corresponding to the output command, and then retry.<br>mkdir:<br>▪ Confirm that the free space is sufficient for `/etc/opt/DynamicLinkManager`.<br>▪ Confirm that the free space is sufficient for `/opt/DynamicLinkManager`.<br>▪ Confirm that the free space is sufficient for the root directory.<br>dd:<br>Confirm that the free space is sufficient for `/etc/opt/DynamicLinkManager`.<br>mount:<br>Confirm that there is free space in the loop device.<br>umount:<br>A temporary file remains, so do the following:<br>1. Unmount `/etc/opt/DynamicLinkManager/tmp/initrdmnt.$$` ($$ is a process ID).<br>2. Remove the `/etc/opt/DynamicLinkManager/tmp/initrdmnt.$$` ($$ is a process ID) file.<br>gzip:<br>Confirm that that the free space is sufficient at the creation destination of the initial ramdisk image file.<br>dlmcfgmgr:<br>▪ Confirm that the version of the OS has not changed since you made the initial ramdisk image file.<br>▪ Take action according to the contents output in the log of the `dlmcfgmgr` utility.<br>dlnkmgr:<br>▪ Confirm that HDLM has been configured.<br>▪ Take action according to the contents output in the log of the `dlmcfgmgr` utility.<br>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 section 7.2. |
| KAPL12327-E | An invalid entry is registered in the file.<br>File name = *aa...aa* | Details<br>*aa...aa*: The name of the file containing the invalid entry<br>Action<br>Make sure that an HBA driver that can be applied when using an HDLM device as a boot disk has been installed. |
| KAPL12328-I | Execution of linuxrc completed. | Action<br>None. |
| KAPL12329-I | The utility for supporting the boot disk started. | Action<br>None. |

| | | |
|---|---|---|
| KAPL12330-I | The utility for supporting the boot disk completed. | Action<br>None. |
| KAPL12331-E | A directory with the specified kernel version does not exist. Directory name = *aa...aa* | Details<br>*aa...aa*: The directory requiring the specified kernel version<br>Action<br>Confirm that the kernel version specified in the parameters of the `dlmmkinitrd` utility for supporting the boot disk is installed, and then retry. |
| KAPL12332-E | The root partition could not be found in the file. File name = *aa...aa* | Details<br>*aa...aa*: The name of the file referenced by the `dlmmkinitrd` utility for supporting the boot disk<br>Action<br>Confirm that the settings of the root partition are in the `/etc/fstab` file.<br>Confirm that the root partition is mounted. |
| KAPL12333-E | An HBA that HDLM can support as a boot disk was not found in the file. File name = *aa...aa* | Details<br>*aa...aa*: The name of the file referenced by the `dlmmkinitrd` utility for supporting the boot disk<br>Action<br>Make sure that an HBA driver that can be used when the boot disk is an HDLM device is in the following file:<br>▪ When Red Hat Enterprise Linux AS3/ES3 is used:<br>`/etc/modules.conf`<br>▪ When Red Hat Enterprise Linux AS4/ES4 is used:<br>`/etc/modprobe.conf`<br>Also, make sure that the HBA driver has been installed. |
| KAPL12334-E | Multiple root partition entries exist in the file. File name = *aa...aa* | Details<br>*aa...aa*: The name of the file referenced by the `dlmmkinitrd` utility for supporting the boot disk<br>Action<br>Modify the `fstab` file specified in the `fstab` parameter of the boot disk support utility (`dlmmkinitrd`), and then retry the operation. Alternatively, modify the root partition setting in the `/etc/fstab` file and then retry the operation. |
| KAPL12335-E | A file system that HDLM can support as a boot disk was not found in the file. File name = *aa...aa* | Details<br>*aa...aa*: The name of the file referenced by the `dlmmkinitrd` utility for supporting the boot disk<br>Action<br>Modify the `fstab` file specified in the `fstab` parameter of the boot disk support utility (`dlmmkinitrd`), or the file system of the `/etc/fstab` file root partition, to a file system that can be applied when the HDLM device is used as a boot disk (ext2 or ext3). Also, make sure that the kernel package has been installed correctly. |
| KAPL12336-E | The user does not have permission to execute the utility for supporting the boot disk. | Action<br>Retry as a user with root permission. |
| KAPL12341-I | Execution of bootup script started. | Action<br>None. |

| | | |
|---|---|---|
| KAPL12342-I | Execution of bootup script completed. | Action<br>None. |
| KAPL12343-I | The bootup executable is being created. | Action<br>None. |
| KAPL12344-I | A compressed initial ramdisk image (*aa...aa*) was created. | Details<br>*aa...aa*: The created initial ramdisk image<br>Action<br>None. |
| KAPL12345-E | The version of the kernel supported by the installed HDLM does not match the currently running kernel version. | Details<br>The version of the kernel to which installed HDLM has adjusted and the kernel that is running now is not corresponding. There is a possibility of installing the package of the kernel after installing HDLM.<br>Action<br>Install HDLM that adjusts to the kernel that is running now. |
| KAPL12346-E | The HDLM file is invalid. File name = *aa...aa* | Details<br>A file necessary for processing HDLM cannot be found. Alternatively, the content of a file is invalid.<br>*aa...aa*: Name of file where error was detected.<br>Action<br>Reinstall an HDLM. |

## 8.16 Messages from the dlmsetopt Utility for Setting HDLM Driver Option

| | | |
|---|---|---|
| KAPL12551-E | The utility for setting HDLM driver option can be executed with the name dlmsetopt only. | Action<br>Change the utility name back to dlmsetopt, and then execute it. |
| KAPL12552-E | Multiple instances of the utility for setting HDLM driver option (dlmsetopt) cannot execute concurrently. | Action<br>Wait until the executing utility finishes, and then re-execute `dlmsetopt`.<br>If two or more utilities are not executing simultaneously, delete the following files, and then re-execute the utility:<br>`/etc/opt/DynamicLinkManager/.dlm_modules.conf.lock` |
| KAPL12553-W | Usage: /opt/DynamicLinkManager/bin/dlmsetopt {-r retrycount \| -inqt InquiryTimeout \| -inqr InquiryRetry \| -h} | Action<br>Refer to 7.7, and then specifying the correct argument. Then, re-execute the `dlmsetopt` utility. |
| KAPL12554-I | The utility for setting HDLM driver option has started. | Action<br>None. |
| KAPL12555-I | The utility for setting HDLM driver option completed normally. | Action<br>None. |
| KAPL12556-I | An attempt to execute the utility for setting HDLM driver option has failed. | Action<br>Check the action recommended in the error message output before this message, and take action accordingly. |
| KAPL12557-I | The user stopped the utility for setting HDLM driver option. | Details<br>Even if the message KAPL12555-I is output before this message, the contents set by the execution of the `dlmsetopt` utility will have no effect and the settings will return to the state before the execution.<br>Action<br>If necessary, re-execute the `dlmsetopt` utility. |
| KAPL12558-I | Please restart the computer so that the option settings take effect. | Action<br>Restart the computer. |
| KAPL12559-E | An invalid value was specified in the parameter. Option = *aa...aa*, Parameter = *bb...bb* | Details<br>*aa...aa*: Option<br>*bb...bb*: Specified parameter(Maximum 10 characters. If there are more than 10 characters, only 10 characters will be output followed by `...`)<br>Action<br>Refer to 7.7, and then specifying the correct argument. Then, re-execute the `dlmsetopt` utility. |
| KAPL12560-E | An option is not specified. | Action<br>Refer to 7.7, specify the correct argument, and then re-execute the `dlmsetopt` utility. |

| | | |
|---|---|---|
| KAPL12561-E | An invalid option was specified. Option = *aa...aa* | Details<br>*aa...aa*: Option<br>Action<br>Refer to 7.7, specify the correct argument, and then re-execute the `dlmsetopt` utility. |
| KAPL12562-E | A necessary file or directory was not found. File or directory name = *aa...aa* | Details<br>*aa...aa*: Referenced directory or file name<br>Action<br>Make sure that the directory or file indicated in the message has not been renamed. If the directory or file has been renamed, change the name back to its original name, and then re-execute the utility for setting HDLM driver option (`dlmsetopt`). |
| KAPL12563-E | An attempt to create a directory has failed. Directory = *aa...aa* | Details<br>*aa...aa*: Directory that an attempt was made to create<br>Action<br>Make sure there is enough free disk space, and then re-execute the utility for setting HDLM driver option (`dlmsetopt`).<br>Delete any unnecessary files in the file system, and then re-execute the `dlmsetopt` utility. |
| KAPL12564-E | An attempt to back up a file has failed. File name = *aa...aa* | Details<br>*aa...aa*: Backup file name<br>Action<br>Make sure there is enough free disk space, and then re-execute the utility for setting HDLM driver option (`dlmsetopt`). |
| KAPL12565-E | Could not modify file. File name = *aa...aa* | Details<br>*aa...aa*: Name of file to be updated<br>Action<br>Make sure there is enough free disk space, and then re-execute the utility for setting HDLM driver option (`dlmsetopt`). If this message still appears, stop any unnecessary processes, and then re-execute the `dlmsetopt` utility. |

| | | |
|---|---|---|
| KAPL12566-E | The utility for setting HDLM driver option was executed and ended abnormally last time. | Action<br>Compare the contents of the `/etc/modprobe.conf` file and the `/etc/opt/DynamicLinkManager/DLMSETOPT/modprobe.conf` file. According to the result, perform the following procedures:<br>The `/etc/opt/DynamicLinkManager/DLMSETOPT/modprobe.conf` file was copied for backup by HDLM before the utility for setting HDLM driver option (`dlmsetopt`) is to be executed. If Red Hat Enterprise Linux AS3/ES3 is being used, `modprobe.conf` must be read as `modules.conf`.<br>▪ When the file contents match:<br>Delete the `/etc/opt/DynamicLinkManager/DLMSETOPT/modprobe.conf` file. Then, re-execute the `dlmsetopt` utility to make sure that an error does not occur.<br>▪ When the file contents do not match:<br>Overwrite the `/etc/modprobe.conf` file with the `/etc/opt/DynamicLinkManager/DLMSETOPT/modprobe.conf` file, and then re-configure the content that was configured in the `/etc/modprobe.conf` file according to the user's environment. Then, delete the `/etc/opt/DynamicLinkManager/DLMSETOPT/modprobe.conf` file. Finally, re-execute the `dlmsetopt` utility to make sure that an error does not occur. |
| KAPL12567-E | The version of the kernel supported by the installed HDLM does not match the currently running kernel version. | Details<br>The version of the kernel to which installed HDLM has adjusted and the kernel that is running now is not corresponding. There is a possibility of installing the package of the kernel after installing HDLM.<br>Action<br>Install HDLM that adjusts to the kernel that is running now. |
| KAPL12568-E | The HDLM file is invalid. File name = *aa...aa* | Details<br>A file necessary for processing HDLM cannot be found. Alternatively, the content of a file is invalid.<br>*aa...aa*: Name of file where error was detected.<br>Action<br>Reinstall an HDLM. |
| KAPL12801-E | The version of the kernel supported by the installed HDLM does not match the currently running kernel version. | Details<br>The version of the kernel to which installed HDLM has adjusted and the kernel that is running now is not corresponding. There is a possibility of installing the package of the kernel after installing HDLM.<br>Action<br>Install HDLM that adjusts to the kernel that is running now. |
| KAPL12802-E | The HDLM file is invalid. File name = *aa...aa* | Details<br>A file necessary for processing HDLM cannot be found. Alternatively, the content of a file is invalid.<br>*aa...aa*: Name of file where error was detected.<br>Action<br>Reinstall an HDLM. |

| | | |
|---|---|---|
| KAPL12851-E | The version of the kernel supported by the installed HDLM does not match the currently running kernel version. | Details<br>The version of the kernel to which installed HDLM has adjusted and the kernel that is running now is not corresponding. There is a possibility of installing the package of the kernel after installing HDLM.<br>Action<br>Install HDLM that adjusts to the kernel that is running now. |
| KAPL12852-E | The HDLM file is invalid. File name = *aa...aa* | Details<br>A file necessary for processing HDLM cannot be found. Alternatively, the content of a file is invalid.<br>*aa...aa*: Name of file where error was detected.<br>Action<br>Reinstall an HDLM. |
| KAPL12901-E | The version of the kernel supported by the installed HDLM does not match the currently running kernel version. | Details<br>The version of the kernel to which installed HDLM has adjusted and the kernel that is running now is not corresponding. There is a possibility of installing the package of the kernel after installing HDLM.<br>Action<br>Install HDLM that adjusts to the kernel that is running now. |
| KAPL12902-E | The HDLM file is invalid. File name = *aa...aa* | Details<br>A file necessary for processing HDLM cannot be found. Alternatively, the content of a file is invalid.<br>*aa...aa*: Name of file where error was detected.<br>Action<br>Reinstall an HDLM. |

## 8.17 Messages from the installgetras Utility for Collecting HDLM Installation Error Information

| | | |
|---|---|---|
| KAPL13401-E | No parameter has been specified. | Details<br>No parameter (directory to which collected information is output) has been specified.<br>Action<br>Check the parameters of the utility for collecting HDLM install error information, and then retry. |
| KAPL13402-E | The first parameter has not been set to a directory. Value= *aa...aa* | Details<br>The first directory must be set to a directory to which collected information is output.<br>*aa...aa*: first parameter<br>Action<br>Check the parameters of the utility for collecting HDLM install error information, and then retry. |
| KAPL13403-E | You lack write permission for the specified directory. Value= *aa...aa* | Details<br>You do not have write permission for the specified directory, or the creation of a subdirectory of the specified directory failed.<br>*aa...aa*: first parameter<br>Action<br>Check the following.<br>1. Check whether you have access permission for the specified directory.<br>2. Check whether the specified directory name is correct.<br>3. Check that the disk has sufficient free space. |
| KAPL13404-W | The specified directory already exists. Do you want to overwrite it? [y/n]: | Details<br>The specified directory already exists. Select `y` to overwrite it or `n` to cancel.<br>Action<br>The specified directory already exists. Enter `y` to overwrite the existing file. Enter `n` or press any other key to terminate the utility for collecting HDLM install error information without executing it. |
| KAPL13405-E | The root directory has been specified in the first parameter. | Details<br>A root "`/`" cannot be specified in a directory to which collected information is output.<br>Action<br>Check the parameters of the utility for collecting HDLM install error information, and then retry. |
| KAPL13406-E | You lack privileges for executing the utility for collecting HDLM install error information. | Details<br>The utility for collecting HDLM install error information must be executed by a root user.<br>Action<br>Retry as a user with root privileges. |

| | | |
|---|---|---|
| KAPL13407-I | The file has been obtained successfully. File = *aa...aa*, Collection time = *bb...bb*(GMT:*cc...cc*) | Details<br>The file to be collected has been obtained.<br>*aa...aa*: Collected file name<br>*bb...bb*: the year of grace/month/day hour:minute:second<br>*cc..cc*: the year of grace/month/day hour:minute:second(GMT)<br>Action<br>None. |
| KAPL13408-E | Processing terminated before completion because a signal was received. | Details<br>The process has been terminated by an operation such as `Ctrl` + `c`.<br>Action<br>The utility for collecting HDLM install error information terminated before completion.<br>If the directory is unnecessary, delete directory. |
| KAPL13409-E | The utility for collecting HDLM install error information completed normally. | Details<br>The processing of the utility for collecting HDLM install error information has been terminated as `n` was sent to a confirmation.<br>Action<br>None. |
| KAPL13410-I | A user terminated the utility for collecting HDLM install error information. | Details<br>The processing of the utility for collecting HDLM install error information has been terminated as `n` was sent to a confirmation.<br>Action<br>None. |
| KAPL13411-W | The entered value is invalid. Continue operation? [y/n]: | Details<br>A value other than `y` or `n` has been entered for a [`y` /`n` ] request. Enter `y` or `n`.<br>Action<br>Enter `y` or `n`. |
| KAPL13412-E | The entered value is invalid. The utility for collecting HDLM install error information stops. | Details<br>The processing of the utility for collecting HDLM install error information has been terminated as an invalid response was sent three times consecutively to a request.<br>Action<br>Re-execute the utility for collecting HDLM install error information. |
| KAPL13413-W | The file does not exist. File = *aa...aa* | Details<br>The file subject to collecting does not exist.<br>*aa...aa*: Name of file to collect<br>Action<br>None. |

| | | |
|---|---|---|
| KAPL13414-E | The file could not be copied. File = *aa...aa*, Details = *bb...bb* | Details<br>Execution of the cq command failed.<br>*aa...aa*: file name you tried to copy<br>*bb...bb*: cq output message<br>Action<br>Check the system configuration. |
| KAPL13415-E | An attempt to archive the install error information failed. Details = *aa...aa* | Details<br>Execution of the tar command failed.<br>*aa...aa*: tar command output message<br>Action<br>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. |
| KAPL13416-E | An attempt to gzip the install error information failed. Details = *aa...aa* | Details<br>Execution of the gzip command failed.<br>*aa...aa*: gzip command output message<br>Action<br>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. |
| KAPL13417-E | The install error information does not exist. | Details<br>The install error information does not exist.<br>Action<br>After installing HDLM, re-execute the `installgetras` utility. |
| KAPL13418-E | Too many parameters have been specified. | Details<br>Two or more parameters have been specified.<br>Action<br>Check the parameters of the utility for collecting HDLM install error information, and then retry. |

## 8.18 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<br>There is no path on which the operation can be performed.<br>Action<br>Refresh the host information, check the path status, and then perform the operation again. |
| 1003 | Details<br>No path was detected.<br>Action<br>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<br>Memory required for HDLM internal processing could not be allocated.<br>Action<br>Terminate unnecessary applications to increase free memory, or restart the host. |
| 1006 | Details<br>The Offline path cannot be placed Online.<br>Action<br>Remove the error in the path, and then retry. |
| 1007 | Details<br>The target path of the offline operation is the last path of the device and cannot be placed in the offline status.<br>Action<br>Click Refresh to update the host information, check the path status, and then retry the offline operation. |
| 1015 | Details<br>The Offline path cannot be placed Online.<br>Action<br>Remove the error in the path, and then retry. |
| 1016 | Details<br>The target path(s) are already Online.<br>Action<br>Update the host information, and then check the path status. |
| 1017 | Details<br>The target path(s) are already Offline(C).<br>Action<br>Update the host information, and then check the path status. |

| | |
|---|---|
| 1019 | Details<br>An error occurred during HDLM internal processing.<br>Action<br>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 section 7.2. |
| 1020 | Details<br>An unexpected error occurred during HDLM internal processing.<br>Action<br>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 section 7.2. |
| 1025 | Details<br>A parameter value is invalid.<br>Action<br>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 section 7.2. |
| 1026 | Details<br>The getting Path Information has been stopped because the configuration of paths was changed during the processing of the getting Path Information.<br>Action<br>Refresh the host information, check the path status, and then perform the operation again. |
| 1027 | Details<br>The error monitoring interval and the number of times that the error is to occur conflict with the automatic failback checking interval.<br>Action<br>Set the monitoring interval to a value that is equal to or more than (auto failback checking interval x number of times). |
| 1033 | Details<br>An attempt to acquire the HDLM version information failed.<br>Action<br>Re-execute. 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 there is a maintenance contract for HDLM. For details on the `DLMgetras` utility, see section 7.2. |
| 1034 | Details<br>An attempt to acquire information about the HDLM version or SP version has failed.<br>Action<br>Re-execute. 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 there is a maintenance contract for HDLM. For details on the `DLMgetras` utility, see section 7.2. |

| | |
|---|---|
| 1035 | Details<br>An attempt to acquire information about the HDLM version or SP version has failed.<br>Action<br>Re-execute. 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 there is a maintenance contract for HDLM. For details on the `DLMgetras` utility, see section 7.2. |
| 1036 | Details<br>An attempt to acquire information about the HDLM version or SP version has failed.<br>Action<br>Re-execute. 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 there is a maintenance contract for HDLM. For details on the `DLMgetras` utility, see section 7.2. |
| 1037 | Details<br>A parameter value is invalid.<br>Action<br>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 section 7.2. |
| 1041 | Details<br>An attempt to communicate with the HDLM Manager has failed.<br>Action<br>Check whether the HDLM Manager is running on the host. If it is not running, start the HDLM Manager. |
| 1042 | Details<br>Information about the path configuration on the specified LU does not match the path configuration information held by HDLM.<br>Action<br>Refresh the host information, check the path status, and then perform the operation again. |
| 1045 | Details<br>A parameter value is invalid.<br>Action<br>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 section 7.2. |
| 1046 | Details<br>A parameter value is invalid.<br>Action<br>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 section 7.2. |

# Appendix A Notes on Linux Commands and Files

## A.1 Notes on Linux Commands and Files

This section gives notes on Linux commands and files.

## A.2 Notes on the /proc/partitions File

The HDLM device and each partition of the HDLM device that are recognized by Linux and can be used are recorded in the `/proc/partitions` file. The `/proc/partitions` file records the I/O statistics per device; however, it does not record the I/O statistics of the HDLM device. The I/O statistics are always 0.

If you are going to obtain statistics about an HDLM device, calculate from the value recorded for the SCSI device that corresponds to the HDLM device. The following is an example of output statistics.

```
# cat /proc/partitions
major minor  #blocks  name     rio rmerge rsect ruse wio wmerge wsect wuse running use aveq

 253     0    2403360 sddlmaa 0 0 0 0 0 0 0 0 0 0 0
 253     1    2401686 sddlmaa1 0 0 0 0 0 0 0 0 0 0 0
 253    16    2403360 sddlmab 0 0 0 0 0 0 0 0 0 0 0
 253    17     803218 sddlmab1 0 0 0 0 0 0 0 0 0 0 0
 253    18     803250 sddlmab2 0 0 0 0 0 0 0 0 0 0 0
 253    19     795217 sddlmab3 0 0 0 0 0 0 0 0 0 0 0
                  •
                  •
                  •
   8     0    2403360 sda 6 0 12 80 0 0 0 0 0 80 80
   8     1    2401686 sda1 4 0 8 0 0 0 0 0 0 0 0
   8    16    2403360 sdb 12 6 36 100 1 0 2 0 0 100 100
   8    17     803218 sdb1 0 0 0 0 0 0 0 0 0 0 0
   8    18     803250 sdb2 0 0 0 0 0 0 0 0 0 0 0
   8    19     795217 sdb3 0 0 0 0 0 0 0 0 0 0 0
```

**Figure A.0.1 Example of a Displayed /proc/partitions File**

## A.3 Notes on Linux Commands

When the host and the storage subsystem are connected by multiple paths, HDLM uses all of these paths as the HDLM devices. Consequently, with `Open()` and `Close()` system calls to the HDLM device, open or close processing occurs on each of the multiple paths. Linux commands that issue successive `Open()` or `Close()` system calls to multiple HDLM devices, such as the `lvmdiskscan` or `pvscan` commands, will take longer to execute after HDLM is installed.

Table A.0.1 shows an example of how long the execution of these Linux commands can be expected to take before and after HDLM installation. The execution times shown in Table A.0.1were measured by using LVM1 commands. LVM2 commands show little difference before and after HDLM installation.

**Table A.0.1 Examples of Linux Command Execution Times**

| Linux commands | Time | | | |
|---|---|---|---|---|
| | Example 1#1 | | Example 2#2 | |
| | Before HDLM installation | After HDLM installation | Before HDLM installation | After HDLM installation |
| lvmdiskscan | 6.1 sec | 19.6 sec | 3.8 sec | 10.6 sec |
| pvscan | 2.6 sec | 12.6 sec | 1.5 sec | 5.2 sec |

# 1

CPU: Pentium III 1.4 GHz x 2

Memory: 1 GB

Connection between host and storage subsystems: 64 LU x 2 paths (total 128 paths)

# 2

CPU: Xeon 2.0GHz x 2

Memory: 2 GB

Connection between host and storage subsystems: 64 LU x 2 paths (total 128 paths)

## A.4 Notes on the iostat Command

If an HDLM device is not specified in the `iostat` command, the HDLM device will not be displayed. Information of an HDLM device is displayed only when the HDLM device is specified in the `iostat` command (in such case as `iostat -x /dev/sddlmaa`). The information (CPU or I/O statistics) displayed for the HDLM device is always `0`, and the actual value is displayed for each SCSI device that conforms to the HDLM device.

If you are going to obtain information of an HDLM device, calculate from the value of a SCSI device that conforms to the HDLM device. The following is an example of executing the `iostat` command.

```
# iostat -x /dev/sddlmaa /dev/sda /dev/sdq
Linux 2.4.21-37.ELsmp (localhost)      05/20/06

avg-cpu:  %user   %nice    %sys %iowait   %idle
           0.09    0.00    0.10    0.02   99.78

Device: rrqm/s wrqm/s r/s  w/s rsec/s wsec/s rkB/s wkB/s avgrq-sz avgqu-sz await svctm %util
sddlmaa   0.00   0.00 0.00 0.00  0.00   0.00  0.00  0.00     0.00     0.00  0.00  0.00  0.00
sdq       0.00   0.00 0.00 0.00  0.00   0.00  0.00  0.00     8.00     0.00  2.86  2.86  0.00
sda       0.18   1.37 0.09 0.20  2.03  12.65  1.01  6.32    50.61     0.02  55.41 1.11  0.03
```

**Figure A.0.2 Example of Executing the iostat Command**

## A.5 Notes on the mkfs Command

If you are going to create a file system for an HDLM device, specify the name of a logical device file of the HDLM device in the `mkfs` command. However, if a file system is created for a SCSI device that is already in use, the applicable HDLM device will be created with a file system even without creating a new file system by using the `mkfs` command.

Furthermore, if the `mkfs` command is executed to a SCSI device, a message confirming whether to execute the command will appear, however, when an HDLM device is specified, the message will not appear.

## A.6 Notes on the fdisk Command

You can perform operations on a partition of an HDLM device by executing the `fdisk` command, with the name of the logical device file of the HDLM device specified. Up to 15 partitions can be created just as for a SCSI device. In addition, partitions can be set in a corresponding SCSI device at the same time by setting partitions in the HDLM device.

However, partitions for a device managed by HDLM will not be automatically created at the same time even when partitions are created for a corresponding SCSI device that HDLM manages. In this case, restart the host to make the partitions the same as those for the corresponding SCSI device.

If HDLM is installed for the first time, an HDLM device that inherits the partitions of the corresponding SCSI device will be created automatically.

Also, if you execute the `fdisk` command, which is in an IPF or EM64T/AMD64 environment, on the HDLM device while running SUSE LINUX Enterprise Server 9, SUSE LINUX Enterprise Server 10 or Red Hat Enterprise Linux AS4/ES4, an error message (KAPL05023-E) is output to the syslog file. However, HDLM operations are not affected.

## A.7 Notes on the sar Command

Executing the `sar` command with the `-d` or the `-A` option does not display HDLM devices. To obtain statistics about an HDLM device, calculate from the value of the SCSI device that conforms to the HDLM device. The following is an example of executing the `sar` command.

```
# sar -d
Linux 2.4.21-37.ELsmp  (localhost)        05/20/06

00:00:00          DEV       tps  rd_sec/s  wr_sec/s
00:10:00      dev65-0      0.00      0.00      0.00
00:10:00     dev65-16      0.00      0.00      0.00
00:10:00     dev65-32      0.00      0.00      0.00
                 :
Average:      dev65-0      0.00      0.00      0.00
Average:     dev65-16      0.00      0.00      0.00
Average:     dev65-32      0.00      0.00      0.00
                 :
```

**Figure A.0.3 Example of Executing the sar Command**

## A.8 Notes on the fdisk and parted Commands

When creating partitions on an HDLM device, only use the `fdisk` command or the `parted` command.

If partitions have been created on an HDLM device using the `fdisk` or `parted` command, some of the partitions created by the `fdisk` command might not be visible to the `parted` command. Likewise, some of the partitions created by the `parted` command might not be visible to the `fdisk` command. For this reason, if partitions that are created using the `fdisk` command are modified using the `parted` command, the partition information might be corrupt.

## A.9 Notes on the parted Command

If the `parted` command is used to create partitions on an HDLM device, after issuing the `parted` command, execute the following command:

```
# blockdev --rereadpt HDLM device SCSI device 1 ... SCSI device n
```

Where:

*HDLM device*: The HDLM device name on which the partition was created (for example, /dev/sddlmad)

*SCSI device* 1..*n*: The SCSI device names which are being managed by the above HDLM device (for example, /dev/sda)

## A.10 Notes on the vgrename and lvrename Commands

If a logical volume is created on an md device, you can rename a volume group or a logic volume by using the `vgrename` or `lvrename` commands. After renaming, if the `vgdisplay -v -D` command, `vgscan`, or `lvscan` commands are executed, the information from the previous volume group or logic volume, might be displayed. To display information from after the change, deactivate the volume group, and then activate the volume group after restarting the md device.

# Appendix B Differences in Functionality between HDLM versions

This section gives precautions on differences in functionality between HDLM versions.

## B.1 Differences in Functionality Between Versions 5.9.1 and Versions Earlier than 5.9.1

- In versions earlier than 5.9.1, the timeout value of `SCSI INQUIRY` is fixed. Therefore, depending on the connected storage subsystem, a timeout occurs when the response for `SCSI INQUIRY` is delayed. In versions 5.9.1 and later, you can use the `dlmsetopt` utility to customize the timeout value according to the storage subsystem in use.
- In versions earlier than 5.9.1, the message below will *not* be output when you install HDLM. In versions 5.9.1 and later, the following message will be output when you install HDLM.
  - `Preparing packages for installation...`
  - `HDLM-`*x.xx.x.xxx.-xx*

The following shows an execution example on Red Hat Enterprise Linux AS 4:

```
# /media/cdrom/installhdlm
KAPL09093-I HDLM xx-xx will be installed. Is this OK ? [y/n]: y
Preparing packages for installation...
KAPL09076-I The permanent license was installed.
HDLM-x.xx.x.xxx-xx
KAPL09043-I The installation of HDLM-x.xx.x.xxx-xx completed successfully.
#
```

## B.2 Differences in Functionality Between Versions 5.9 and Versions Earlier than 5.9

The function for operating HDLM from HDLM GUI or HDLM Web GUI is no longer supported.

## B.3 Differences in Functionality Between Versions 5.8 and Versions Earlier than 5.8

Table B.0.1 describes the version of JRE or JDK required to use the HDLM GUI on a host.

**Table B.0.1 JRE/JDK Version Required for HDLM GUI**

| OS | Required version of JRE or JDK | |
|---|---|---|
| | HDLM 5.7.1 | HDLM 5.8 |
| Red Hat Enterprise Linux AS 3 (IA32)<br>Red Hat Enterprise Linux ES 3 (IA32)<br>Red Hat Enterprise Linux AS 4 (IA32)<br>Red Hat Enterprise Linux ES 4 (IA32)<br>SUSE LINUX Enterprise Server 9 (IA32) | JRE v1.4.2_08 | JRE v1.4.2_09 |

| | | |
|---|---|---|
| Red Hat Enterprise Linux AS 3 (IPF)<br>Red Hat Enterprise Linux ES 3 (IPF)<br>Red Hat Enterprise Linux AS 4 (IPF)<br>Red Hat Enterprise Linux ES 4 (IPF) | JDK v1.4.2_08 | JDK v1.4.2_09 |
| Red Hat Enterprise Linux AS 3 (EM64T/AMD64)<br>Red Hat Enterprise Linux ES 3 (EM64T/AMD64)<br>Red Hat Enterprise Linux AS 4 (EM64T/AMD64)<br>Red Hat Enterprise Linux ES 4 (EM64T/AMD64)<br>SUSE LINUX Enterprise Server 9 (EM64T/AMD64) | JRE v1.5.0_03<br>(JRE 64 bit version is supported) | JRE v1.5.0_05<br>(JRE 32 bit version is supported) |

## B.4 Differences in Functionality Between Versions 5.7.1 or Later and Versions Earlier than 5.7.1

In versions earlier than 5.7.1, if the HDLM driver fails to allocate memory for an I/O operation when reading or writing from the HDLM device, it returns an error instead of trying again.

In versions 5.7.1 or later, the system retries to allocate the above memory and an I/O operation is executed successfully. The number of retries can be set by the user. By default, the system retries the operation until memory is allocated.

## B.5 Differences in Functionality Between Versions 5.7.0-01 or Later and Versions Earlier than 5.7.0-01

In versions earlier than 5.7.0-01, when an I/O request is sent to a path in the Online(E) status despite all paths being disconnected, the status of the path is changed to Offline(E) and another path is placed in Online(E) status. In versions 5.7.0-01 or later, even if an I/O request is sent to a path in the Online(E), the status of the path does not change.

In the event of path recovery after all paths have failed, the paths must be placed in Online status by automatic failback or by command before you can access the data.

## B.6 Differences in Functionality Between Versions 5.7 or Later and Versions Earlier than 5.7

- The following distributions are no longer supported:
  - Red Hat Enterprise Linux AS2.1 for IA32
  - Red Hat Enterprise Linux AS2.1 for IPF
  - SUSE LINUX Enterprise Server 8 for IA32
  - SUSE LINUX Enterprise Server 8 for IPF
- In versions 5.7 or later, if there are no files under the `/etc/opt/DynamicLinkManager` directory when HDLM is uninstalled, the directory is deleted.

## B.7 Differences in Functionality Between Versions 5.6.3 or Later and Versions Earlier than 5.6.3

Table B.0.2 describes the version of JRE or JDK required to use the HDLM GUI on a host.

**Table B.0.2 JRE/JDK Version Required for HDLM GUI**

| OS | Required version of JRE or JDK | | |
|---|---|---|---|
| | Earlier than HDLM 5.4.2 | HDLM 5.6 - 5.6.2 | HDLM 5.6.3 |
| Red Hat Enterprise Linux AS 2.1 (IA32)<br>SUSE LINUX Enterprise Server 8 (IA32) | JRE v1.3.1_08 | JRE v1.3.1_15 | JRE v1.4.2_08 |
| Red Hat Enterprise Linux AS 3 (IA32)<br>Red Hat Enterprise Linux ES 3 (IA32) | JRE v1.4.2_03 | JRE v1.4.2_07 | JRE v1.4.2_08 |
| Red Hat Enterprise Linux AS 2.1 (IPF)<br>Red Hat Enterprise Linux AS 3 (IPF)<br>Red Hat Enterprise Linux ES 3 (IPF)<br>SUSE LINUX Enterprise Server 8 (IPF) | JRE v1.4.2_03 | JRE v1.4.2_07 | JRE v1.4.2_08 |

## B.8 Differences in Functionality Between Versions 5.4 or Later and Versions Earlier than 5.4

- The permission for the error log file (`/var/DynamicLinkManager/log/dlmmgrX.log`) has been changed to 600.
- The permission for the trace log directory (`/var/opt/hitachi/HNTRLib2/spool`) has been changed to 755, and the permission for the log file (`/var/opt/hitachi/HNTRLib2/spool/hntr2X.log`) has been changed to 644.
- License keys that worked with versions 04-00, 05-02, and 05-03 are no longer valid with version 5.4.
- HDLM no longer supports LVMs that use HDLM devices on Red Hat Enterprise Linux AS2.1.

# Acronyms and Abbreviations

| | |
|---|---|
| AMS | Automatic Storage Management |
| CFQ | Complete Fair Queuing |
| CHA | channel adapter |
| CPU | Central Processing Unit |
| CRS | Cluster Ready Service |
| DRBD | Distributed Replicated Block Device |
| EVMS | Enterprise Volume Manager |
| GUI | Graphical User Interface |
| HBA | Host Bus Adapters |
| HDLM | HiCommand Dynamic Link Manager |
| I/O | input/output |
| LDEV | logical device |
| LU | logical unit |
| LUN | logical unit number |
| LVM | Logical Volume Manager |
| OS | operating system |
| SAN | storage area network |
| USP | Universal Storage Platform |
| VCS | VERITAS Cluster Server |

# Glossary

| | |
|---|---|
| automatic failback | Functionality for checking the status of paths at regular intervals, and automatically placing the status of a path recovered from an error into Online status.<br>If a path with Offline(E) or Online(E) status recovers from an error, automatic failback places the path Online.<br>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 section 6.4. |
| AutoPATH_ID | An ID that HDLM assigns to a path at system startup. Every path has a unique AutoPATH_ID.<br>(See also: *path*) |
| boot disk environment | An environment in which the startup disk is in a storage subsystem instead of in the host. |
| 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. |
| Dev(Device) | A logical division of an LU that HDLM controls and operates.<br>A Dev is equivalent to a *partition* in Linux.<br>In Linux, each LU can be divided into more than one Dev.<br><br>Each Dev has a *Dev number*.<br>Depending on the access method, Devs are classified into two types: block devices and character devices.<br>(See also: *Dev number*) |
| Dev number | A Dev number (DNum column) in the configuration list in HDLM.<br>This number is assigned to each Dev within an LU. In Linux, a Dev number is called a *partition number*.<br>The Dev number is set to `0` because HDLM manages Devs in units of LUs.<br>(See also: *Dev*) |
| emulation type | An LU type accessible from a host. Since 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.<br>For details on emulation types supported by a storage subsystem, see the maintenance manual for that storage subsystem. |

| | |
|---|---|
| failback | Functionality for placing the status of a path recovered from an error into Online status, and switching the access path to the recovered path. |
| failover | Functionality for switching to another normal path if there is an error in a path, thereby enabling the system to continue to operate. |
| HBA (Host Bus Adapter) | Device that functions as an interface connecting a host with external devices.<br>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 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. |
| host | A generic term for both servers and clients. |
| host device | An area within a host LU.<br>(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. |
| host LU | An LU that the host recognizes. The actual HDev entity is a Dev in the` storage subsystem.<br>Each host LU has a *host LU number*.<br>(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.<br>(See also: *host LU*, *path name*) |
| intermittent error | An error that occurs irregularly due to reasons such as a loose cable connection. |
| 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. HDLM uses this value to identify an LU. |
| load balancing | Functionality for distributing the load across the paths that access each area within an LU. To distribute loads, load balancing uses multiple paths to perform I/O operations.<br>Two types of algorithms are available: round-robin and extended round-robin. |
| logical device file | A device file in the `/dev` directory.<br>***Note***: the term *logical device* here is different from *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. |
| node | A server of cluster members. |

| | |
|---|---|
| non-owner controller | A controller other than an owner controller.<br>Non-owner controllers exist only when the Thunder 9200, Thunder 9500V Series, or TagmaStore AMS/WMS series is in use.<br>(See also: *owner controller*, *non-owner path*) |
| non-owner path | A path that passes through a controller other than an owner controller.<br>Non-owner paths exist only when the Thunder 9200, Thunder 9500V Series, or TagmaStore AMS/WMS series is in use.<br>(See also: *owner path*, *non-owner controller*) |
| owner controller | A controller that has been set as an owner controller for an LU on the storage subsystem when using the Thunder 9200, Thunder 9500V Series, or TagmaStore AMS/WMS series. A controller other than an owner controller is called a *non-owner controller*.<br>When using Lightning 9900 Series, Lightning 9900V Series, TagmaStore USP or Universal Storage Platform V, all controllers are owner controllers.<br>(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 using the Thunder 9200, Thunder 9500V Series, or TagmaStore AMS/WMS series. A path that passes through a *controller* other than an owner controller is called a *non-owner path*.<br>When using the Lightning 9900 Series, Lightning 9900V Series, TagmaStore USP or Universal Storage Platform V, all paths are owner paths.<br>(See also: *owner controller*, *non-owner path*) |
| path | An access path from a host to 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`.<br>(See also: *AutoPATH_ID*) |
| path health checking | Functionality for checking the status of paths at regular intervals.<br>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.<br>Path health checking checks only those paths that have Online status. |
| path name | The name consisting of the following four elements, separated by periods:<br>-Host port number (hexadecimal number)<br>-Bus number (hexadecimal number)<br>-Target ID (hexadecimal number)<br>-Host LU number (hexadecimal number)<br>A path name is used to identify a physical path.<br>(See also: *host LU number*) |
| persistent reservation | Similar to *reservation*, persistent reservation 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.<br>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.<br>(See also: *reservation*) |

| | |
|---|---|
| physical path | An access path from a host to an LU. A path name is used to identify a physical path. |
| reservation | The reservation functionality 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.<br>(See also: *persistent reservation*) |
| 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) | A function, supported by the TagmaStore USP and Universal Storage Platform V, for logically splitting a storage subsystem. This function splits the resources (ports, CLPR, and volumes) in the storage subsystem, so that the user can manage each resource independently. |

# Index

## U



## V



## W