HITACHI
DATA SYSTEMS

Hitachi Freedom Storage™ Thunder 9500™ V Series

Command Control Interface (CCI) User and Reference Guide

Trademarks

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-92DF609-P	September 2002	Preliminary Release
MK-92DF609-0	November 2002	Initial Release
MK-92DF609-1	January 2003	Revision 1, supersedes and replaces MK-92DF609-0
MK-92DF609-2	February 2003	Revision 2, supersedes and replaces MK-92DF609-1
MK-92DF609-3	April 2003	Revision 3, supersedes and replaces MK-92DF609-2
MK-92DF609-4	July 2003	Revision 4, supersedes and replaces MK-92DF609-3

Source Document Revision Level

The following source document was used to produce this 9500V user document:

- *Hitachi Freedom Storage™ Thunder 9500™ V Series Command Control Interface (CCI) User and Reference Guide* (RSD-92DF609-3)

Changes in this Revision:

- Added a note to section 3.1 stating that Windows® 2000 dynamic disks are not supported

Referenced Documents

- *Hitachi Freedom Storage™ Thunder 9500™ V Series ShadowImage 9500V User's Guide* (MK-92DF607)
- *Hitachi Freedom Storage™ Thunder 9500™ V Series Synchronous TrueCopy 9500V User's Guide* (MK-92DF608)

Preface

The *Hitachi Freedom Storage™ Thunder 9500™ V Series Command Control Interface (CCI) User and Reference Guide* describes and provides instructions for performing Command Control Interface (CCI) software operations on the Thunder 9500V array subsystem. The CCI software enables the user to issue Synchronous Remote Copy (hereafter called TrueCopy) and/or Hitachi ShadowImage 9500V commands to the Thunder 9500V array subsystem from the open-systems host (UNIX® -based or PC server).

This user's guide assumes the following:

- The user has a background in data processing and understands RAID storage subsystems and their basic functions.
- The user is familiar with the Thunder 9500V array subsystem.
- The user is familiar with the functionality of the TrueCopy/ShadowImage 9500V features.
- The user has read and understands the *Hitachi Freedom Storage™ Thunder 9500™ V Series ShadowImage 9500V User's Guide* (MK-92DF607).
- The user has read and understands the *Hitachi Freedom Storage™ Thunder 9500™ V Series Synchronous TrueCopy 9500V User's Guide* (MK-92DF608). **Please note that the 9530 disk array subsystem does not use Synchronous TrueCopy 9500V.**

Note: The term "9500V" refers to the entire Hitachi 9500V subsystem family, unless otherwise noted. Please refer to the *Hitachi Freedom Storage™ Thunder 9500™ V Series User and Reference Guide* (MK-92DF601) for further information on the 9500V disk array subsystems, or contact your Hitachi Data Systems account team.

Note: The use of the CCI software and all other Hitachi Data Systems products is governed by the terms of your license agreement(s) with Hitachi Data Systems.

Applicable Platforms

This document applies to the following platforms:

Vendor	Operating System	Cluster (Failover) S/W	Volume Manager	I/O Interface
HP	HP-UX 10.20, 11.0	MC/Service Guard	LVM	Fibre
SUN	Solaris 2.6, 7, 8	VCS	VxVM	Fibre
IBM	AIX 4.3.3, 5.1.0	HACMP	LVM	Fibre
COMPAQ	Tru64 UNIX 5.1	—	LSM	Fibre
—	Windows NT 4.0/2000	—	—	Fibre
—	Red Hat Linux 7.0	—	—	Fibre
SGI	IRIX 6.5.x	—	LVM	Fibre

Terminology

Please note the following:

- The term DF600 refers to the Thunder 9500™ V Series subsystem.

COMMENTS

Please send us your comments on this document: doc.comments@hds.com.

Make sure to include the document title, number, and revision.
Please refer to specific page(s) and paragraph(s) whenever possible.

(All comments become the property of Hitachi Data Systems Corporation.)

Thank you!

Contents

List of Figures

List of Tables

Chapter 1 Overview of Hitachi Copy Solutions

1.1 Hitachi Command Control Interface (CCI)

The licensed Hitachi Command Control Interface (CCI) software product enables you to perform Synchronous TrueCopy 9500V and ShadowImage 9500V operations on the Hitachi Thunder 9500™ V subsystem by issuing commands from the system hosts to the Thunder 9500™ V subsystem. The CCI software interfaces with the system software and high-availability (HA) software on the system hosts as well as with the Synchronous TrueCopy/ShadowImage 9500V software on the Thunder 9500™ V subsystem. For additional information on Synchronous TrueCopy 9500V and ShadowImage 9500V, see the *Hitachi Freedom Storage™ Thunder 9500™ Synchronous TrueCopy 9500V User's Guide* (MK-92DF608) and the *Hitachi Freedom Storage™ Thunder 9500™ V Series ShadowImage 9500V User's Guide* (MK-92DF607).

Note: The 9530 disk array subsystem *does not* use Synchronous TrueCopy 9500V.

CCI provides failover and operation commands which support mutual hot standby in conjunction with industry-standard failover products. CCI also supports a scripting function for defining multiple Synchronous TrueCopy 9500V/ShadowImage 9500V operations in a script (or text) file. Using CCI scripting, you can set up and execute a large number of Synchronous TrueCopy 9500V and/or ShadowImage 9500V commands in a short period of time while integrating host-based high-availability control over remote copy operations.

Synchronous TrueCopy 9500V operations (see section 1.2) involve the primary (main) subsystems and the secondary (remote) subsystems. The primary subsystems contain the Synchronous TrueCopy 9500V primary volumes (P-VOLs), which are the original data volumes. The secondary subsystems contain the Synchronous TrueCopy 9500V secondary volumes (S-VOLs). When Synchronous TrueCopy 9500V is performed using CCI, you need to reserve and configure one volume on each subsystem as the CCI command device.

ShadowImage 9500V operations (see section 1.3) involve primary and secondary volumes within one subsystem. The ShadowImage 9500V P-VOLs contain the original data, and the S-VOLs are the internal duplicate volumes. ShadowImage 9500V allows one S-VOL to be created for each P-VOL. Each S-VOL is paired with the P-VOL independently, allowing each S-VOL to be maintained as an independent copy set.

Note: The use of the CCI software and all other Hitachi Data Systems products is governed by the terms of your license agreement(s) with Hitachi Data Systems.

1.2 Synchronous TrueCopy 9500V

Synchronous TrueCopy 9500V is an optional function and can be installed on the disk subsystem. The Synchronous TrueCopy 9500V enables you to create and maintain remote copies of the data stored on the 9500Vsubsystem for data backup and disaster recovery purposes.

Synchronous TrueCopy 9500V operations can be performed using the Hitachi Command Control Interface (CCI) software on the UNIX®/PC server host. The CCI software on the UNIX®/PC server displays Synchronous TrueCopy 9500V information and allows you to perform Synchronous TrueCopy 9500V operations from the UNIX® command line or via a script file. The CCI software interfaces with the 9500V subsystem through a dedicated LU called a command device.

Read the following notices and follow them; otherwise, a path failure will occur.

- When turning on the subsystem where a path has already been set, turn on the remote subsystem, and turn on the local subsystem after the subsystem become READY. When turning off the subsystem where a path has already been set, turn off the local subsystem and turn off the remote subsystem.
- A path blockage that occurred while using the Synchronous TrueCopy 9500V function, even if the remote subsystem was off, also occurs. The remote subsystem is turned on, and automatically recovered when the remote subsystem is READY (about 4 minutes).

  If a path blockage is not recovered, regardless of being READY, call the Hitachi maintenance personnel.

1.3 ShadowImage 9500V

The ShadowImage 9500V data duplication feature enables you to set up and maintain multiple copies of logical volumes within the same Thunder 9500™ V subsystem. ShadowImage 9500V operations for UNIX®/PC server-based data can be performed using Hitachi Multiple RAID Coupling Feature – Lite (ShadowImage 9500V) software on the host where the Hitachi Command Control Interface (CCI) software is installed.

The Hitachi CCI software on the UNIX®/PC server displays ShadowImage 9500V information and allows you to perform ShadowImage 9500V operations by issuing commands from the UNIX® command line or by executing a script file. The CCI software interfaces with the Thunder 9500™ V subsystem through a dedicated LU called a command device.

1.4 Cooperation between the User and Hitachi Personnel

If a Synchronous TrueCopy 9500V/ShadowImage 9500V suspended-error occurs, the cause is usually due to a failure in the hardware (or when the user forcibly suspends the pair). To recover from a suspended status (PSUE), the hardware error and data must be recovered. To accomplish this task, cooperation between the user and Hitachi Data Systems maintenance personnel is necessary.

To prepare for a PSUE error, make certain that the user (system engineer) monitors the CCI system log. Confirm that the PSUE error is not caused by user operation, then contact Hitachi Data Systems maintenance personnel.

Chapter 2 Overview of CCI Operations

CCI allows you to perform Synchronous TrueCopy 9500V and ShadowImage 9500V operations by issuing Synchronous TrueCopy 9500V/ShadowImage 9500V commands from the UNIX®/PC server host to the Thunder 9500™ V subsystem. ShadowImage 9500V operations are non-disruptive and allow the primary volume of each volume pair to remain online to all hosts for both read and write operations (except when an error occurred during reverse—resync).

Note: Windows® 2000 dynamic disks are not supported.

2.1 Features of Paired Volumes

The logical volumes, which have been handled independently by server machines, can be combined or separated in a pair being handled uniformly by the Synchronous TrueCopy 9500V and/or ShadowImage 9500V pairing function. Synchronous TrueCopy 9500V and ShadowImage 9500V regard those two volumes to be combined or separated as a unique paired logical volume used by the servers. It is possible to handle paired volumes as groups by selecting them in units of server software or in units of the database and its attributes.

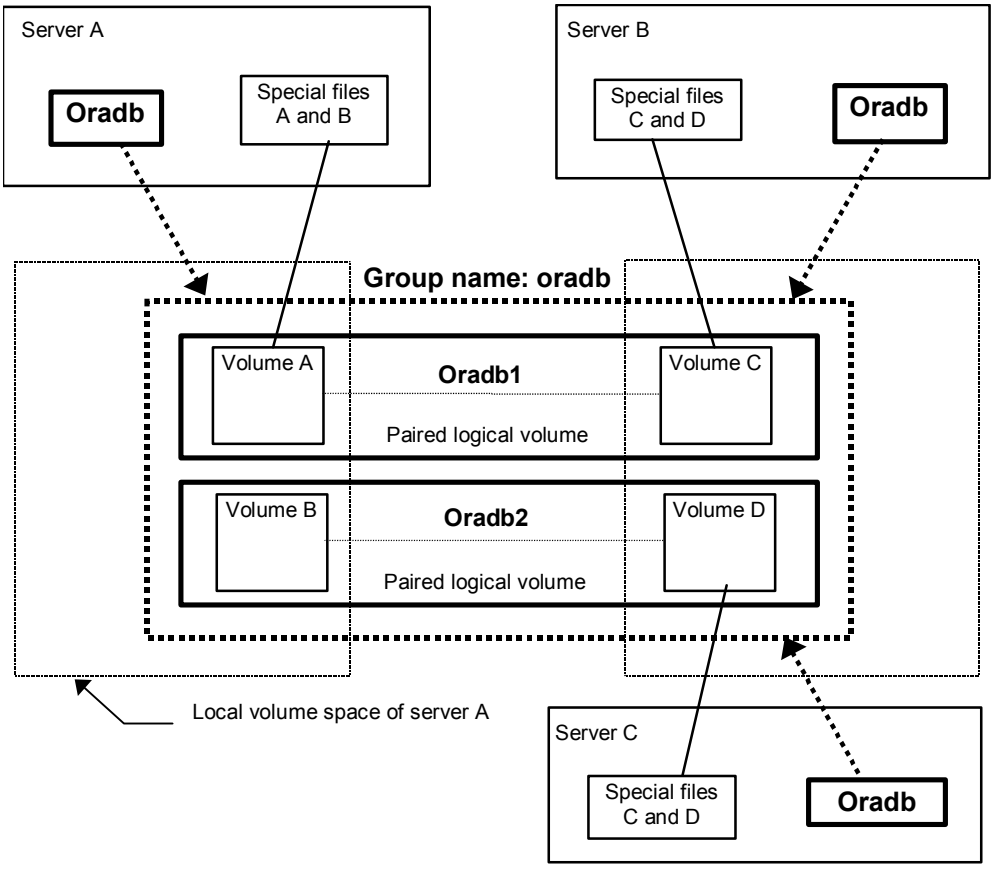

Figure 2.1 Paired Volumes

Addressing paired logical volumes: The correspondence between paired logical volumes and physical volumes is defined by describing the intended paired logical volume names and group names in the configuration definition file of each server. It is possible to define a server for the paired logical volumes in units of the group name. Each paired logical volume must belong to a group in order to determine the corresponding server.

Specification of volumes by commands: Volume names specified by the CCI commands must be given using the paired logical volume names or the group names.

2.1.1 ShadowImage 9500V Mirroring

Mirroring of a single primary volume is possible when the ShadowImage 9500V feature is used. Therefore, the mirror descriptor (MU#) is fixed (MU#0).

2.2 Overview of CCI Operations

This section includes the following:

- Synchronous TrueCopy 9500V Takeover Commands
- Synchronous TrueCopy 9500V Remote Commands
- Synchronous TrueCopy 9500V Operations
- ShadowImage Operations

2.2.1 Synchronous TrueCopy 9500V Takeover Commands

CCI Synchronous TrueCopy 9500V operates in conjunction with the software on the UNIX®/PC servers and the 9500V subsystem Synchronous TrueCopy 9500V functions. The CCI software provides failover and other functions such as backup commands to allow mutual hot standby in cooperation with the failover product on the UNIX®/PC server (e.g., MC/ServiceGuard, FirstWatch®, HACMP). For the proper maintenance of Synchronous TrueCopy 9500V operations, it is important to find failures in paired volumes, recover the volumes from the failure as soon as possible, and continue operation in the original system.

Figure 2.2 illustrates the server failover system configuration. When a server software error or a node error is detected, the operation of the failover software causes the Cluster Manager (CM) to monitor server programs, and causes the CM of the standby node to automatically activate the HA control script of the corresponding server program. The HA control script usually contains the database recovery procedures, server program activation procedures, and other procedures. The takeover commands provided by Synchronous TrueCopy 9500V are activated by the control HA script and execute the control needed for failover of the server.

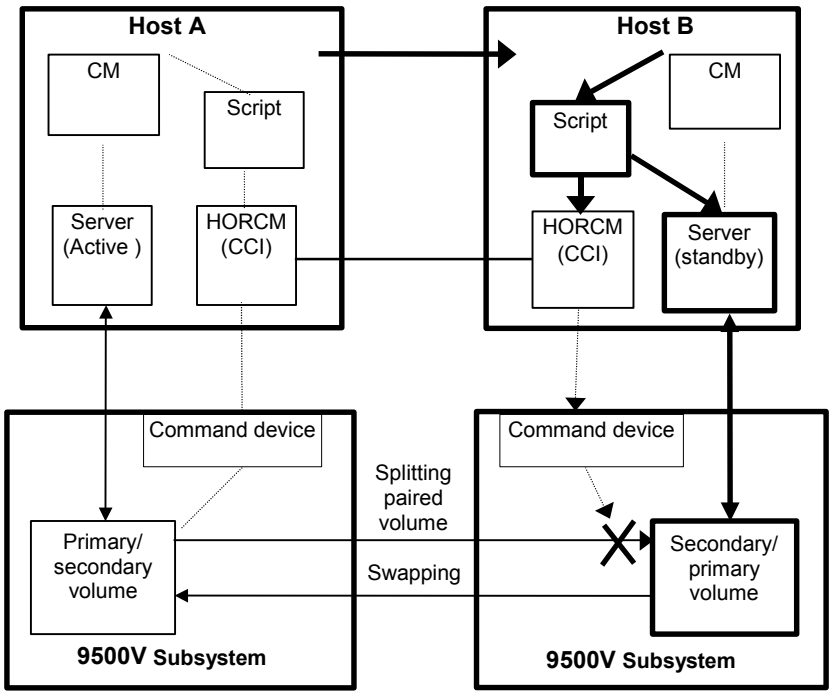

Figure 2.2 Server Failover System Configuration

In a high availability (HA) environment, a package is a group of applications that are scripted to run on the secondary host in the event of a primary host failure. When using the HA software (e.g., MC/ServiceGuard), the package can be transferred to the standby node as an operation executed by the system administrator (see Figure 2.3). However, if the operation is performed in an environment in which Synchronous TrueCopy 9500V is used, the volume is switched from primary to secondary as if an error has occurred, even though data consistency is assured. When returning the package to the current node, it is necessary to copy the secondary volume data into the primary volume, and this operation can take as much time as the initial copy operation for the pair. In actual operation, no package can be transferred when Synchronous TrueCopy 9500V is used. The secondary package is switched to the primary package, and vice versa, when the primary volume is switched to the secondary volume. Therefore, the primary and secondary Synchronous TrueCopy 9500V volumes should be switched depending on the package state.

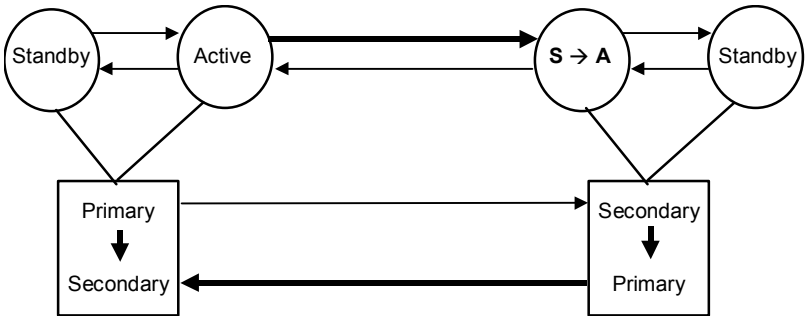

Figure 2.3 Package Transfer on High Availability (HA) Software

2.2.2 Synchronous TrueCopy 9500V Remote Commands

Figure 2.4 illustrates a Synchronous TrueCopy 9500V remote configuration. The Synchronous TrueCopy 9500V remote commands support a function which links the system operation for the purpose of volume backup among UNIX® servers with the operation management of the server system. The Synchronous TrueCopy 9500V remote pair commands are also used to copy volumes in the failover configuration of the servers and to recover the volumes after the takeover.

- **Pair creation command:** Creates a new volume pair. Volume pairs can be created in units of volume or group.
- **Pair splitting command:** Splits a volume pair and allows read and write access to the secondary volume.
- **Pair resynchronization command:** Resynchronizes a split volume pair based on the primary volume. The primary volume remains accessible during resynchronization.
  - **Swaps(p) option (Synchronous TrueCopy 9500V only).** Swaps volume from the S-VOL(P-VOL) to P-VOL(S-VOL) at suspending state on the S-VOL(P-VOL) side and resynchronizes the NEW_S-VOL based on the NEW_P-VOL. As a result of this operation, the volume attributes of own host (local host) become the attributes for the NEW_P-VOL(S-VOL).
- **Event waiting command:** Used to wait for completion of volume pair creation or resynchronization and to check the pair status.
- **Pair status display and configuration confirmation command:** Displays the pair status and configuration of the volume pairs, used for checking the completion of pair creation or pair resynchronization.

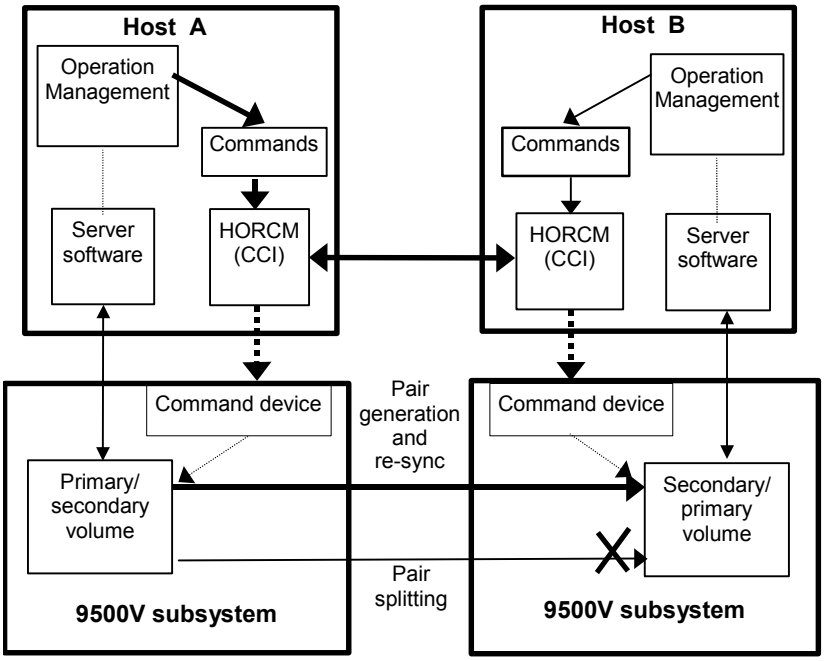

Figure 2.4 Synchronous TrueCopy 9500V Remote System Configuration

2.2.3 ShadowImage Operations

Figure 2.5 illustrates the ShadowImage 9500V configuration. The ShadowImage 9500V commands support a function, which links the system operation with the server system operation management to create a volume backup among UNIX® servers. For detailed information on the operational requirements for ShadowImage 9500V, refer to the *Hitachi Freedom Storage™ Thunder 9500™ ShadowImage 9500V User's Guide* (MK-92DF607).

- **Pair creation command:** Creates a new volume pair. Volume pairs can be created in units of volume.
- **Pair splitting command:** Splits a volume pair and allows read and write access to the secondary volume.
- **Pair resynchronization command:** Resynchronizes a split volume pair based on the primary volume. The primary volume remains accessible during resynchronization.
  - **Restore option:** Resynchronizes a split pair based on the secondary volume (reverse resync). The primary volume is not accessible during resync with restore option.
- **Event waiting command:** Used to wait for completion of volume pair creation or resynchronization and to check the pair status.
- **Pair status display and configuration confirmation command:** Displays the pair status and configuration of the volume pairs, used for checking the completion of pair creation or pair resynchronization.

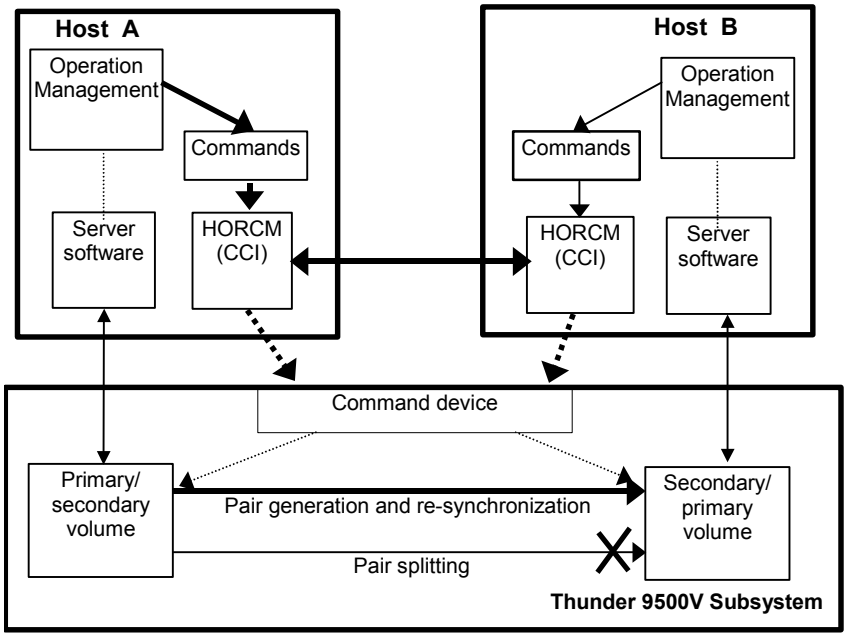

Figure 2.5 ShadowImage 9500V System Configuration

2.3 Volume Pairs

ShadowImage 9500V paired volumes can be created between the same volume sizes. Synchronous TrueCopy 9500V and ShadowImage 9500V commands allow you to create a Synchronous TrueCopy 9500V/ShadowImage 9500V volume pair consisting of one P-VOL and one S-VOL, and up to 31 volume pairs. The Synchronous TrueCopy 9500V P-VOL and S-VOL can be in different subsystems. ShadowImage 9500V pairs are contained within the same subsystem and are maintained by re-synchronizing from P-VOL to S-VOL.

Each volume pair that you create must be registered in the CCI configuration file. ShadowImage 9500V volume pairs must include an MU (mirrored unit) number assigned to the S-VOL. The MU number indicates that the pair is a ShadowImage 9500V pair. Once the correspondence between the paired logical volumes has been defined in the HORCM_DEV section of the configuration file, you can use the configuration file to group the paired volumes into volume groups that can be managed by the host operating system's LVM (logical volume manager).

The host's LVM allows you to manage the Synchronous TrueCopy 9500V/ShadowImage 9500V volumes as individual volumes. Synchronous TrueCopy 9500V/ShadowImage 9500V commands can specify individual logical volume names.

2.3.1 Synchronous TrueCopy 9500V Volume Status

Each Synchronous TrueCopy 9500V pair consists of one P-VOL and one S-VOL.

Table 2.1lists and lists and describes the Synchronous TrueCopy 9500V pair status terms. The P-VOL controls the pair status for the primary and secondary volumes. The major pair statuses are SMPL, PAIR, PSUS/PSUE, and COPY. Read and write requests from the host are accepted or rejected depending on the pair status of the volume.

The pair status changes when a Synchronous TrueCopy 9500V command is executed. The validity of the specified operation is checked according to the status of the volume (primary volume). The user must keep track of the issued commands and the changes of status so that an appropriate response can be made if an error occurs.

Table 2.1 Synchronous TrueCopy 9500V Pair Status

Status	Synchronous TrueCopy 9500V Pair Status	Primary	Secondary
SMPL	Unpaired volume	R/W enabled	R/W enabled
PAIR	Paired volume. Initial copy is complete. Updates are processed synchronously.	R/W enabled	R enabled
COPY	In paired state, but initial copy, pairsplit, or resync operation is not complete. Includes COPY(PD), and COPY(SP) status.	R/W enabled	R enabled
PSUS SSUS (split)	In paired state, but updates to the S-VOL data are suspended due to user-requested pairsplit. The 9500V keeps track of P-VOL and S-VOL updates while the pair is split.	R/W enabled	R/W enabled when using write enable pairsplit option.
PSUE (error)	In paired state, but updates to the S-VOL volume data are suspended due to an error condition. When a PSUE pair is resynced, the 9500V copies the entire P-VOL to the S-VOL (same as initial copy).	R/W enabled if no error has occurred and when the fence level is set to **data** in the primary volume.	R enabled

Table 2.2 shows the relationship between pair status and Synchronous TrueCopy 9500V command acceptance.

Table 2.2 Pair Status Versus Synchronous TrueCopy 9500V Commands

	Synchronous TrueCopy 9500V Command					
	Paircreate		Pairsplit			Pairresync
Pair Status	Copy	Nocopy	Suspend –r, -rw	Suspend –P	Simplex	Resync
① SMPL	Accepted ②	Accepted ③	Rejected	Rejected	Acceptable	Rejected
② COPY ↓	Acceptable	Acceptable	Accepted ④	Rejected	Accepted ①	Acceptable
③ PAIR	Acceptable	Acceptable	Accepted ④	Accepted ④	Accepted ①	Acceptable
④ PSUS	Rejected	Rejected	Acceptable	Acceptable	Accepted ①	Accepted ② (see note 1)
⑤ PSUE	Rejected	Rejected	Rejected	Rejected	Accepted ①	Accepted ② (see note 1)

Note 1: For the **SSWS** state after **SVOL-SSUS-takeover**, the pairresync command (from PVOL to SVOL) is rejected because the delta data for S-VOL becomes dominant, and its state expects to use -swaps(p) option of pairresync. If the pairresync command (from P-VOL to S-VOL) is rejected, verify this special state using the -fc option of the pairdisplay command.

2.3.2 ShadowImage 9500V Volume Status

Each ShadowImage 9500V pair consists of one P-VOL and one S-VOL. Table 2.3 lists and describes the ShadowImage 9500V pair status terms. The P-VOL controls the pair status for the primary and secondary volumes. The major pair statuses are SMPL, PAIR, PSUS/PSUE, and COPY/RCPY. Read and write requests from the host are accepted or rejected depending on the pair status of the volume.

The pair status changes when a ShadowImage 9500V command is executed. The validity of the specified operation is checked according to the status of the volume (primary volume). The user must note the issued commands and the changes of the status so that appropriate action can be taken if an error occurs.

Table 2.3 ShadowImage 9500V Pair Status

Status	ShadowImage 9500V Pair Status	Primary	Secondary
SMPL	Unpaired volume	R/W enabled	R/W enabled
PAIR	Paired volume. Initial copy is complete. Updates are processed asynchronously.	R/W enabled	R enabled*
COPY	In paired state, but initial copy, pairsplit, or resync operation is not complete. Includes COPY (PD), COPY (SP), and COPY (RS) status.	R/W enabled	R enabled*
RCPY	In paired state, but reverse resync operation is not complete. Includes COPY (RS-R) status.	R/W enabled*	R enabled
PSUS SSUS (split)	In paired state, but updates to the S-VOL data are suspended due to user-requested pairsplit. The 9500V keeps track of P-VOL and S-VOL updates while the pair is split.	R/W enabled	R/W enabled when using write enable pairsplit option.
PSUE (error)	In paired state, but updates to the S-VOL volume data are suspended due to an error condition. When a PSUE pair is resynced, the 9500V copies the entire P-VOL to the S-VOL (same as initial copy).	R/W enabled if no error has occurred in the primary volume. (If the status transits from RCPY, all access is disabled.)	R enabled* (If the status transits from RCPY, all access is disabled.)

* Read is disabled when the **-m noread** option of the paircreate command is specified.

Table 2.4 shows the relationship between pair status and ShadowImage 9500V command acceptance.

Table 2.4 Pair Status versus ShadowImage 9500V Commands

	ShadowImage 9500V Command					
	Paircreate		Pairsplit			Pairresync
Pair Status	No -split	-split	-E option	-C option	-S option	Resync
① SMPL	Accepted ②	Accepted ②→④	Rejected	Rejected	Acceptable	Rejected
② COPY RCPY	Acceptable	Accepted* ②→④	Accepted ⑤	Accepted* ②→④	Accepted ①	Acceptable
③ PAIR	Acceptable	Accepted ④	Accepted ⑤	Accepted ④	Accepted ①	Acceptable
④ PSUS	Rejected	Acceptable	Accepted ⑤	Acceptable	Accepted ①	Accepted ②
⑤ PSUE	Rejected	Rejected	Acceptable	Rejected	Accepted ①	Accepted ②

Note 1: If the P-VOL does not have Write in the PAIR state, then data identical with an S-VOL is guaranteed. Therefore, when using the S-VOL with the SMPL state, after stopping Write to the P-VOL, generate a paired volume, and split the paired volume after confirming that the paired volume has PAIR status. In the PSUE state, ShadowImage 9500V does not manage differential data at the P-VOL or S-VOL. Therefore, pairresync issued to a pair in the PSUE state is all copy performance, but the copy progress rate returned by the **-fc** option of the pairdisplay command indicates "0%".

Note 2: The state change (②→④) is effective COPY state only that is changed without specification of -split for paircreate command.

Example:

```
# pairsplit -g oradb
# pairdisplay -g oradb –fc
Group   PairVol(L/R) (Port#,TID,LU-M), Seq#, LDEV#.P/S, Status, % , P-LDEV# M
oradb   oradev3(L)   (CL2-A , 3,  4-0) 8071   28..P-VOL PSUS,  100     29   W
oradb   oradev3(R)   (CL2-A , 3,  5-0) 8071   29..S-VOL COPY,   97     28   -
```

2.3.3 Synchronous TrueCopy 9500V Fence-Level Settings

Synchronous TrueCopy 9500V volume pairs are assigned a fence level for write I/Os to ensure the mirroring consistency of critical volumes. Accordingly, when the secondary volume takes over from the primary volume, the takeover action is determined according to the pair status and fence level of the corresponding secondary volume. Table 2.5 shows the relationship between Synchronous TrueCopy 9500V pair status and fence level.

Table 2.5 Relationship between Synchronous TrueCopy 9500V Pair Status and Fence Level

Synchronous TrueCopy 9500V Pair Status of Volume	Fence Level and Write Response	
	Data [1]	Never [2]
Write response 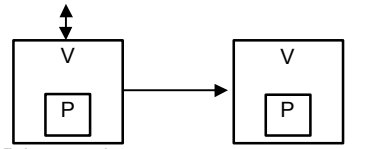 	OK	OK
	Mirroring consistency assured	Mirroring consistency assured
Write response 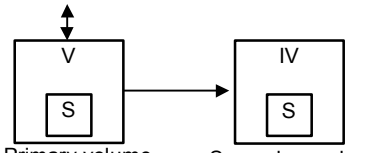 	ERROR	OK
	Mirroring consistency assured	Mirroring consistency not assured
Write response 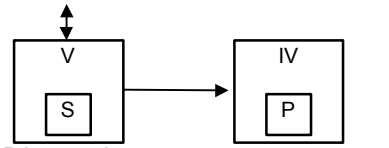 	ERROR	OK
	Mirroring consistency assured	Mirroring consistency not assured

Note 1: When fence level is **data**, the P-VOL returns a write error whenever data consistency is lost, so mirroring consistency is assured. The S-VOL can continue operation, regardless of its status.

Note 2: When fence level is **never**, writing is enabled whenever data consistency is lost, whether or not the S-VOL status is updated. Thus, the currency of the S-VOL can be evaluated as follows:
S: The secondary volume is dubious.
P: The secondary volume is substantially dubious, since it can continue operation and is also dubious. The P-VOL status must be checked to confirm the mirroring consistency.

2.3.3.1 Setting the Fence Level

Figure 2.6 shows the relationship between redo log files (journal) and data files. If the S-VOL takes over from the P-VOL in the status shown in Figure 2.6(where two errors have occurred), the secondary host leaves data (V) unprocessed in the rollback processing and cannot be recovered completely. Therefore, the fence level of a redo log file must be defined as **data**. Once the fence level is set to **data**, the P-VOL returns an error if data is inconsistent when a write request is issued by the host. Since the writing into the data file has not been executed due to a write error of the redo log file, the log file stays consistent with the data file. However, when the fence level is set to **data**, a write I/O error occurs even when operation is suspended due to an error in the S-VOL. Accordingly, the duplication becomes meaningless when the S-VOL takes over. Thus, applications using paired volumes with the **data** fence level should be able to handle write I/O errors properly. For example, Oracle creates multiple redo log files by itself (three by default). The fence level can be set to **data** when disk errors are permissible by creating multiple files.

Since most UNIX®-based file systems (excluding JFS and VxFS) have no journal files, the fence level should be defined as **Never**. When a takeover by the S-VOL occurs, **fsck** is executed on the volume and the file system is cleaned up, even if the S-VOL is undefined at the secondary host. The data that will be lost depends on how much differential data is contained in the P-VOL when the S-VOL is suspended. During operation, error recovery should be performed when the suspended status (PSUE) is detected (when one error occurs).

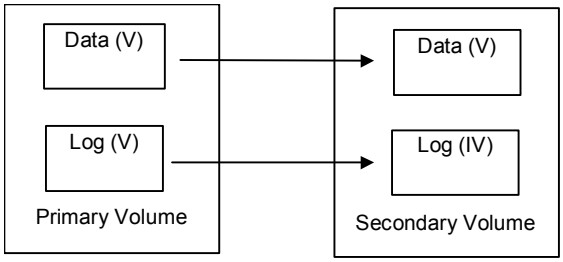

Figure 2.6 Relation Between Logs and Data in Paired Status

2.4 Applications of CCI Commands

This section provides examples of tasks that can be performed using CCI commands.

2.4.1 Backing Up Secondary Volume in Paired Status

2.4.1.1 Synchronous TrueCopy 9500V

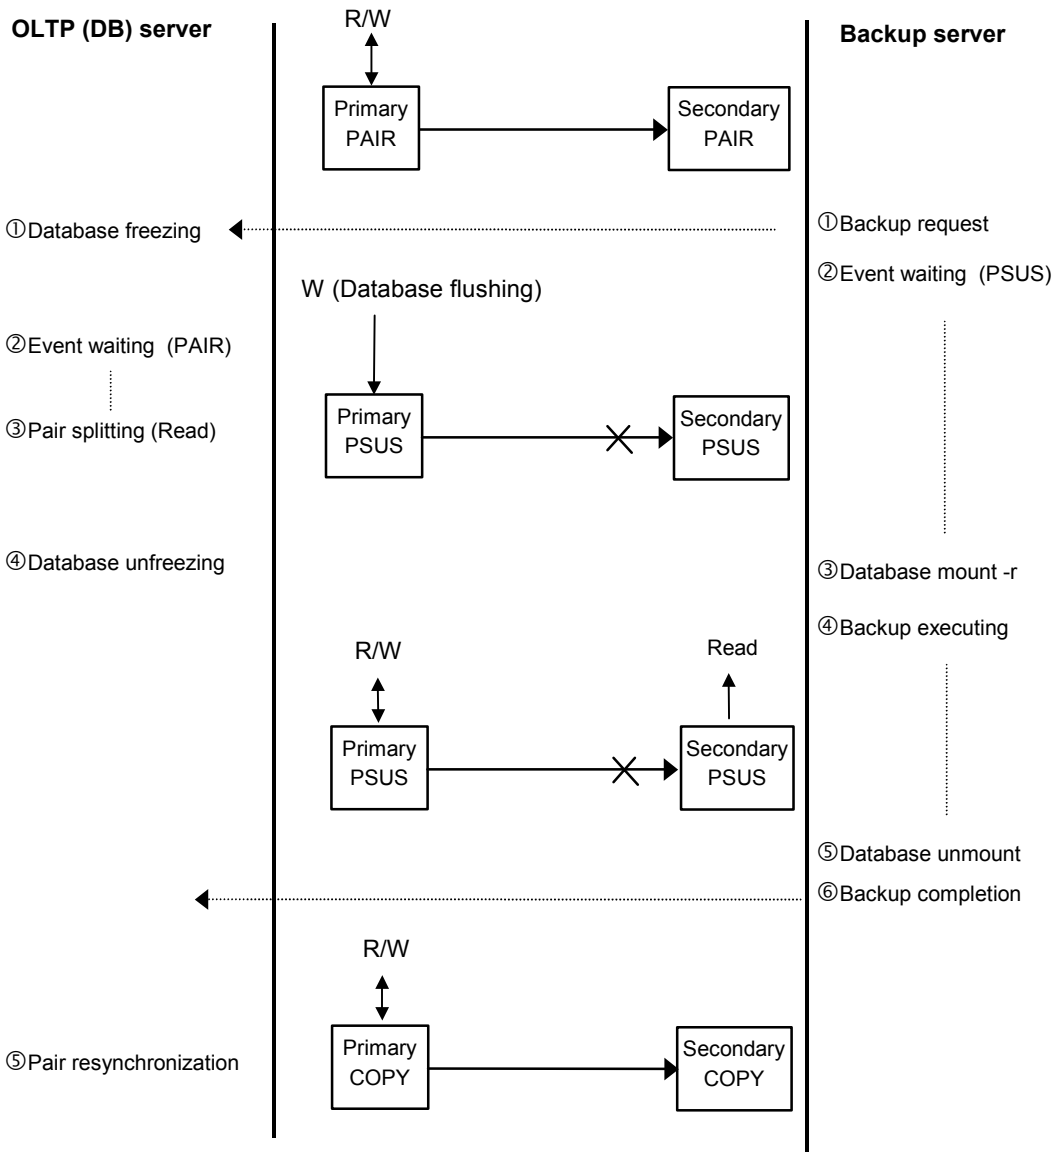

Figure 2.7 Backing Up S-Vol in Paired Status Using Synchronous TrueCopy 9500V

2.4.1.2 ShadowImage 9500V

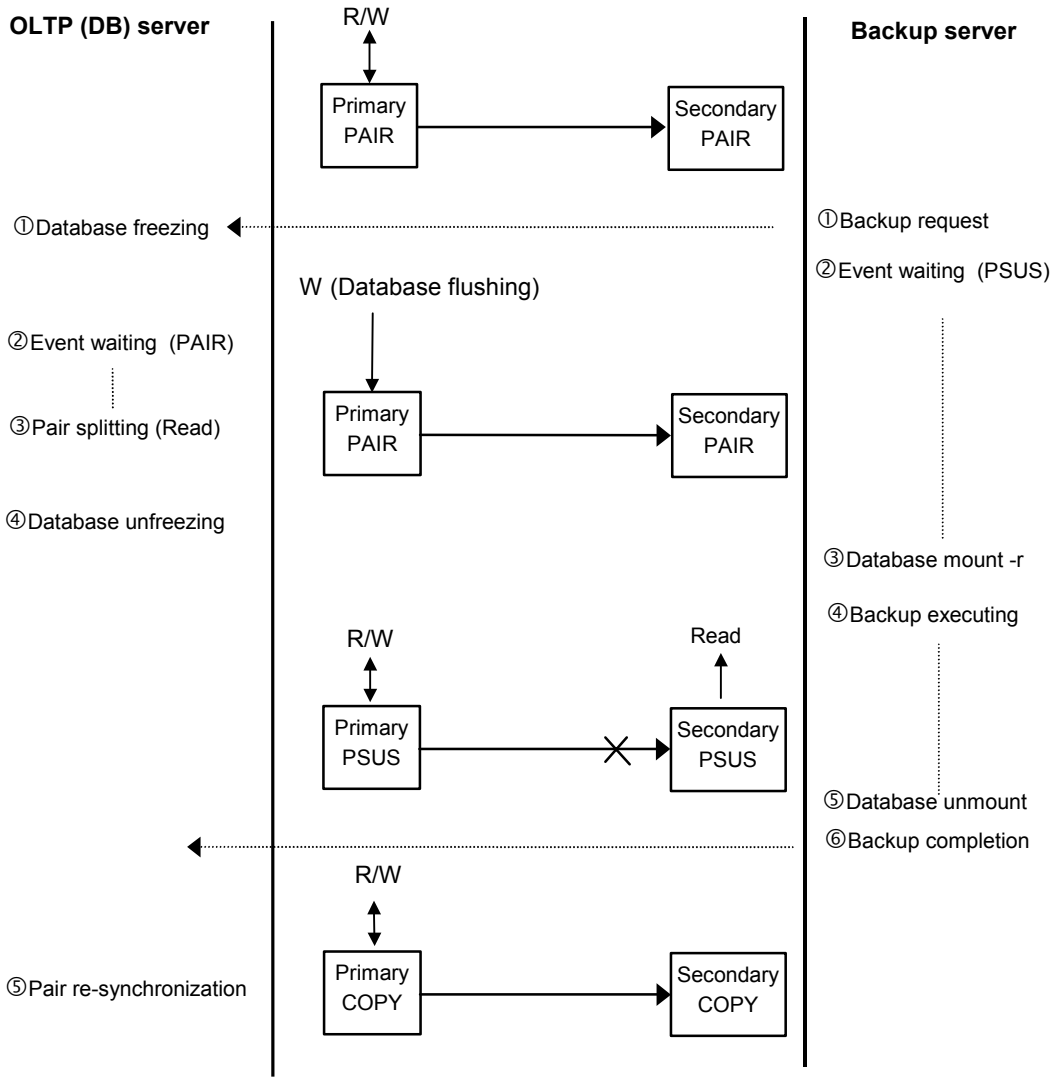

Figure 2.8 Backing Up S-VOL in Paired Status Using ShadowImage 9500V

2.4.2 Restoring Secondary Volume to Primary Volume in Split Status

2.4.2.1 Synchronous TrueCopy 9500V

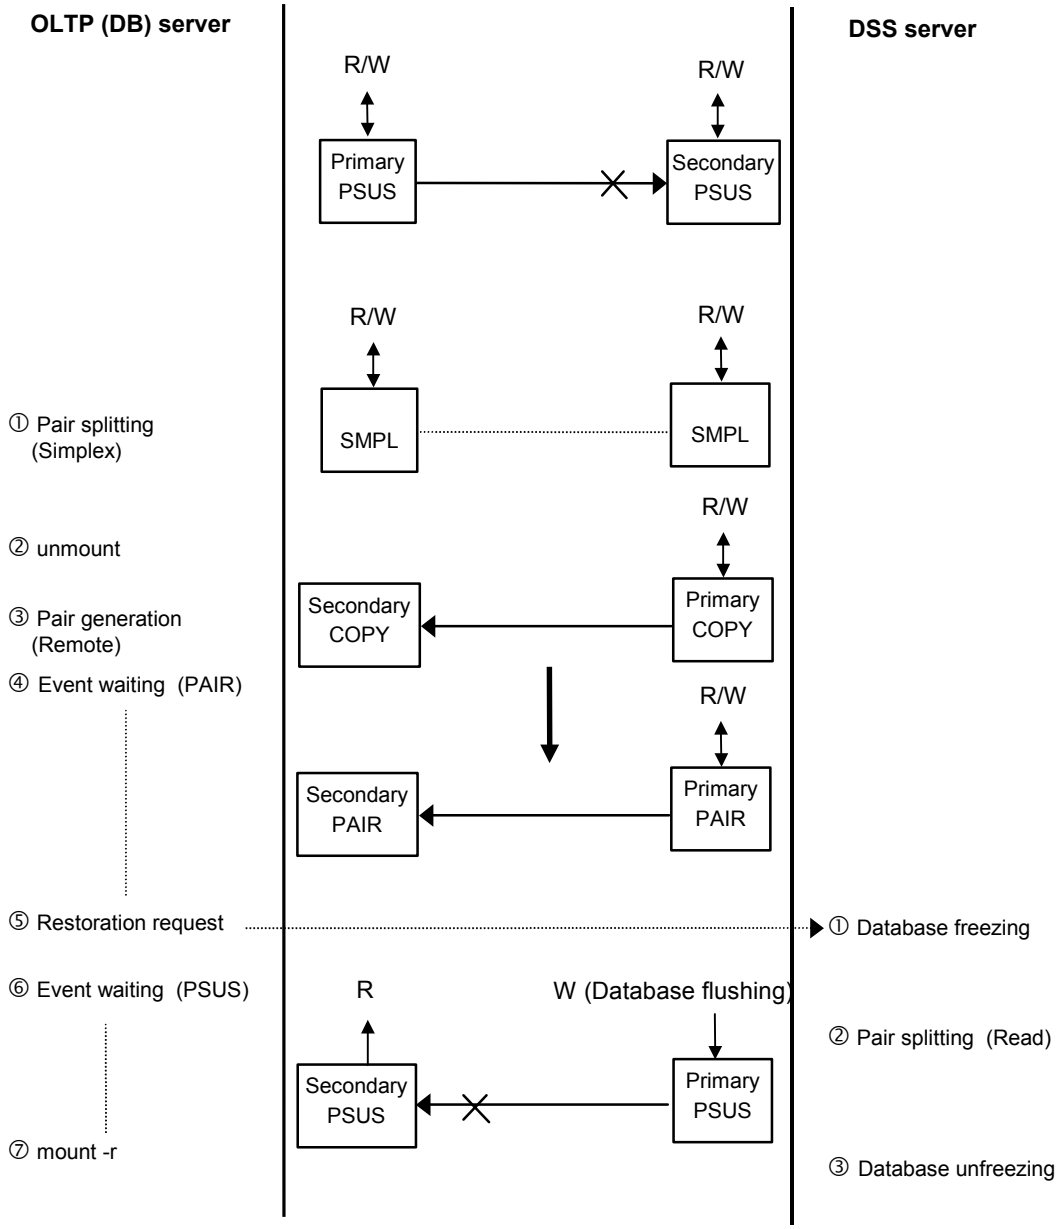

Figure 2.9 Restoring S-VOL to P-VOL in Split Status Using Synchronous TrueCopy 9500V

2.4.2.2 ShadowImage 9500V

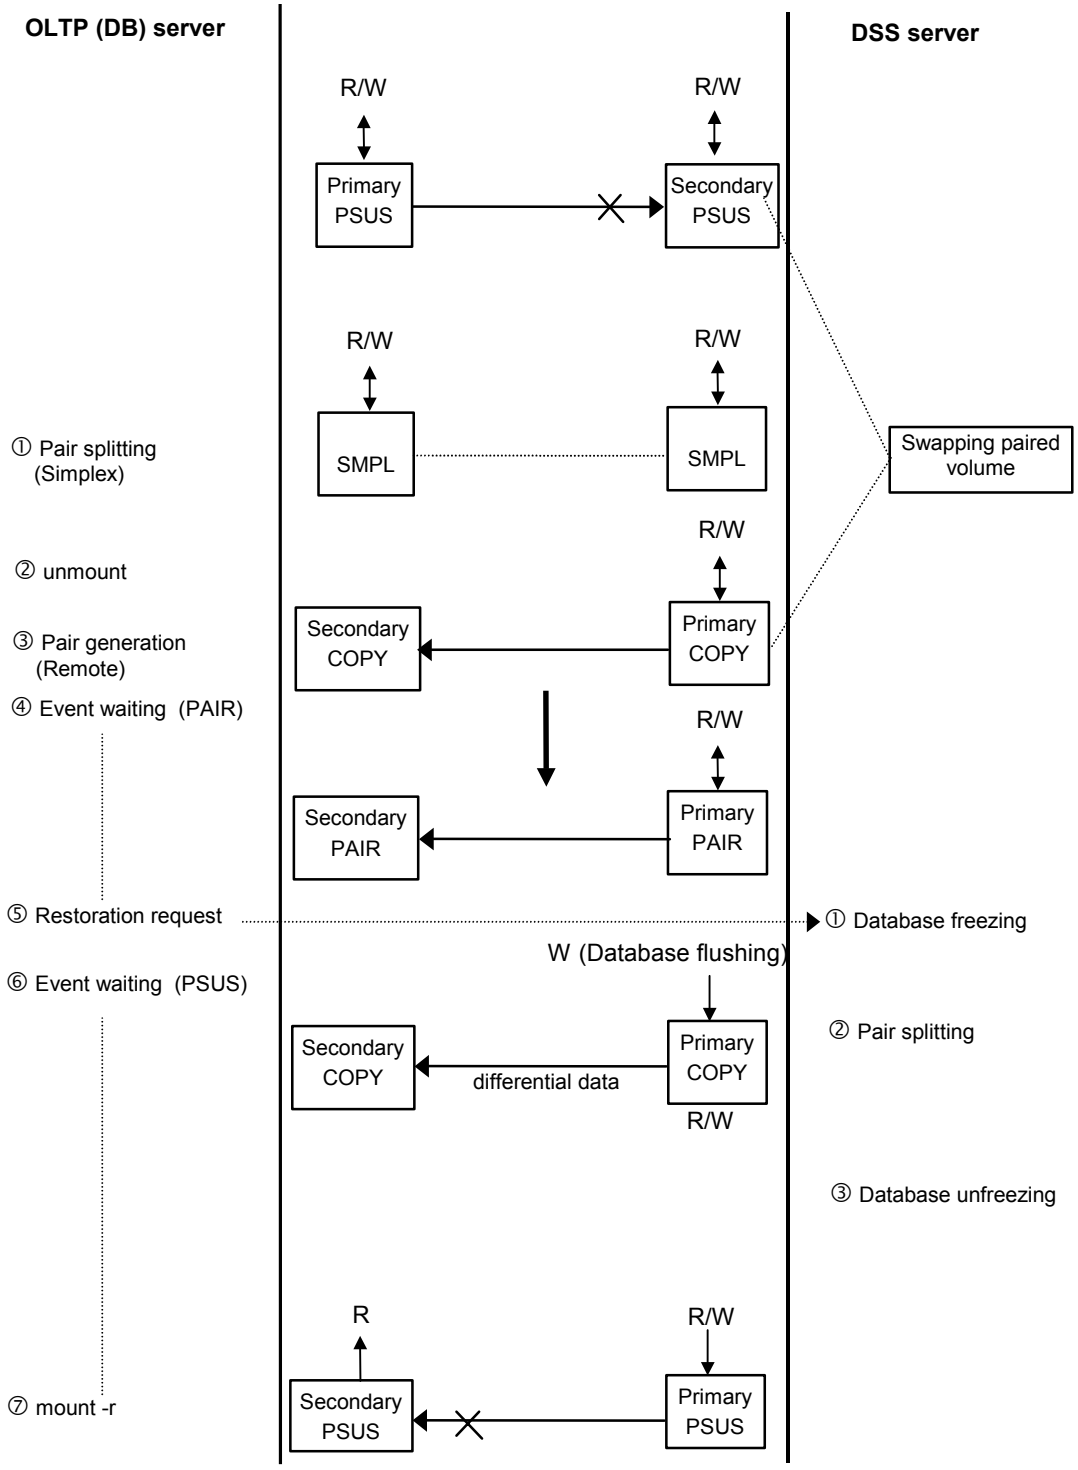

Figure 2.10 Restoring S-VOL to P-VOL in Split Status Using ShadowImage 9500V

2.4.3 Swapping Paired Volume for Duplex Operation

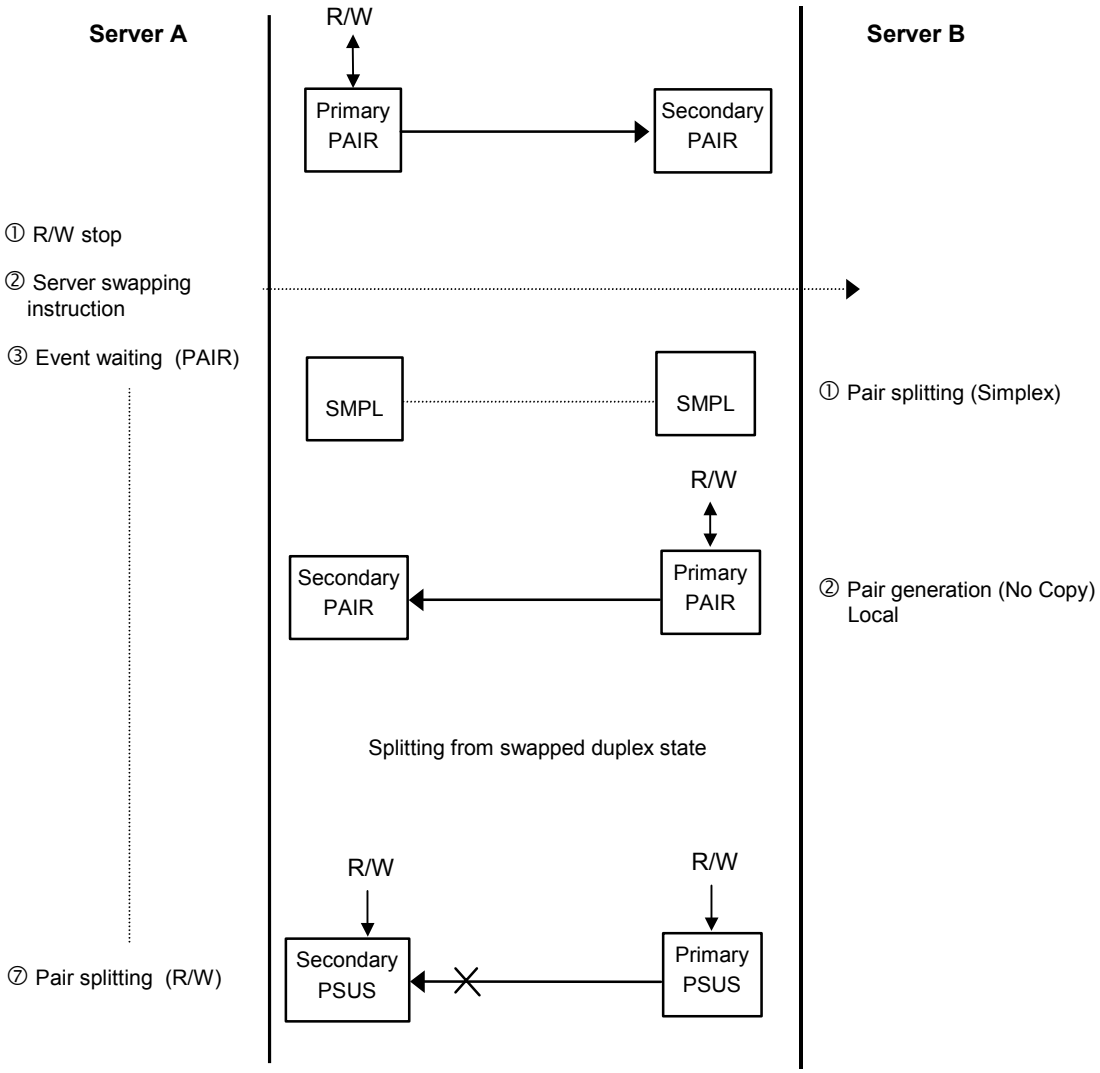

Figure 2.11 Swapping Paired Volume for Duplex Operation (Synchronous TrueCopy 9500V Only)

2.4.4 Restoring S-VOL for Duplex Operation

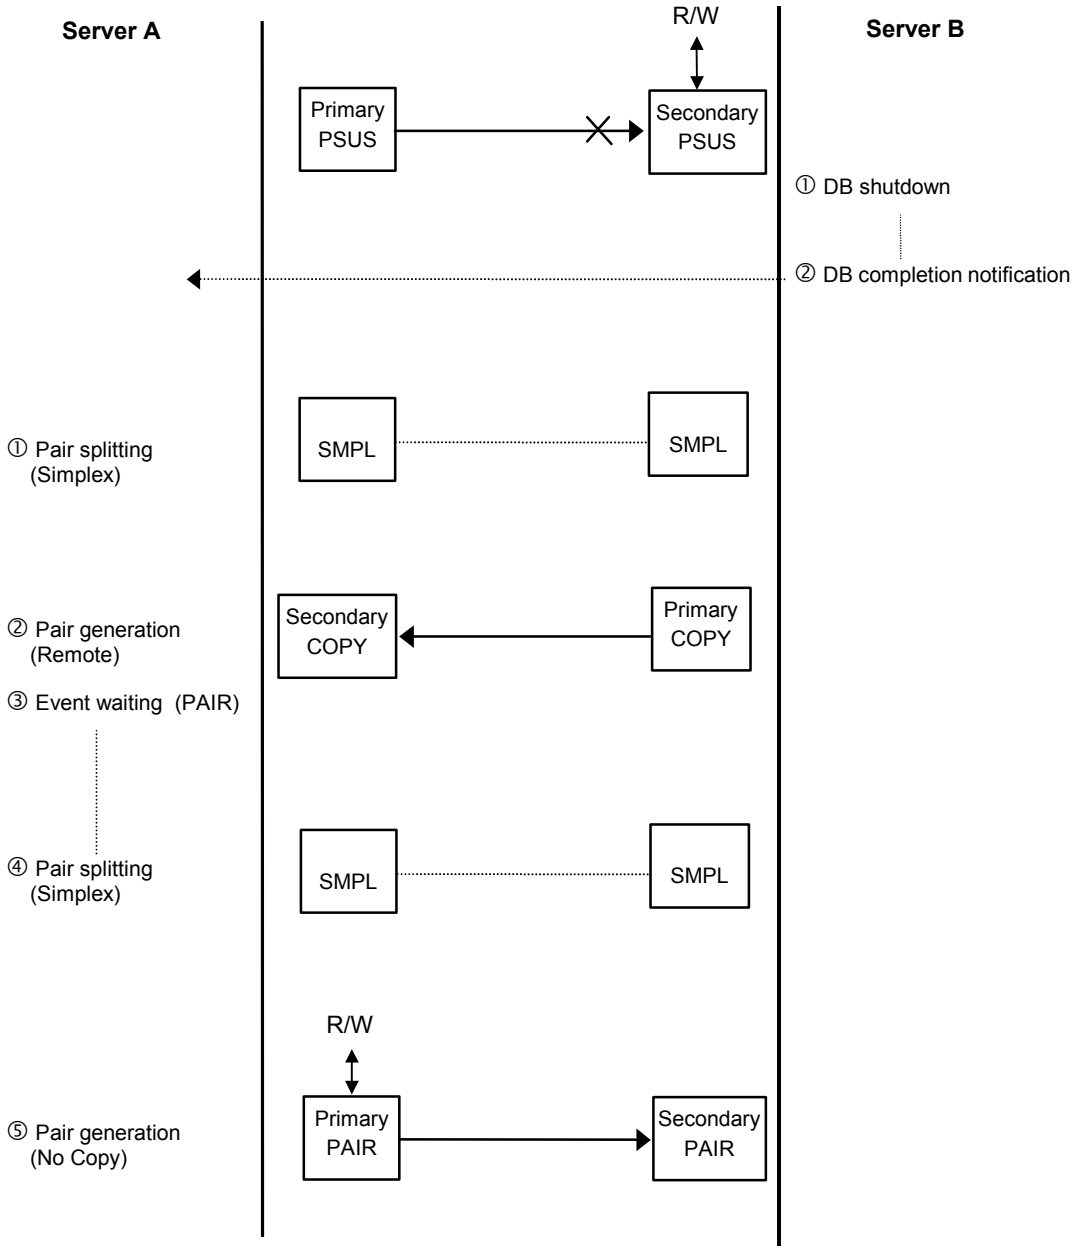

Figure 2.12 Restoring S-VOL for Duplex Operation (Synchronous TrueCopy 9500V Only)

2.5 CCI Software Structure

Figure 2.13 illustrates the CCI software structure: the CCI components on the Thunder 9500™ V subsystem, and the CCI instance on the UNIX®/PC server. The CCI components on the Thunder 9500™ V subsystem include the command device(s) and the Synchronous TrueCopy 9500V/ShadowImage 9500V volumes. Each CCI instance on a UNIX®/PC server includes:

- HORC Manager (HORCM):
  - Log and trace files
  - A command server
  - Error monitoring and event reporting files
  - A configuration management feature
- Configuration definition file (defined by the user)
- The Synchronous TrueCopy 9500V and/or ShadowImage 9500V user execution environments, which contain the Synchronous TrueCopy 9500V/ShadowImage 9500V commands, a command log, and a monitoring function.

2.5.1 HORCM Operational Environment

The HORCM operates as a daemon process on the host server and is activated automatically when the server machine starts up or manually by the start-up script. HORCM refers to the definitions in the configuration file when it is activated. The environmental variable HORCM_CONF is used to define the configuration file to which it is referred.

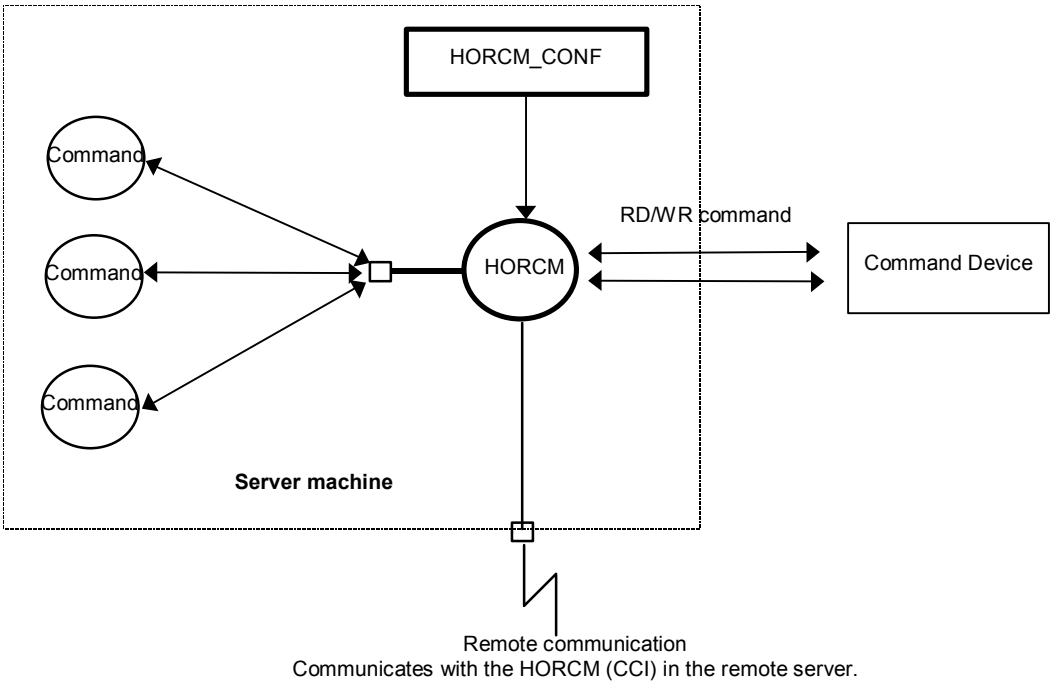

Figure 2.13 HORCM Operational Environment

2.5.2 CCI Instance Configurations

The basic unit of the CCI software structure is the CCI instance. Each copy of CCI on a server is a CCI instance. Each instance uses a defined configuration file to manage volume relationships while maintaining awareness of the other CCI instances. Each CCI instance normally resides on one server (one node). If two or more nodes are run on a single server (e.g., for test operations), it is possible to activate two or more instances using instance numbers.

Note: The default command execution environment for CCI is Synchronous TrueCopy 9500V (without specification of HORCC_MRCF). Therefore, in order to use CCI command for ShadowImage 9500V, the user must specify the environment variable HORCC_MRCF in the configuration definition file (HORCM_CONF).

The CCI instance shown in Figure 2.14 has a remote execution link and a connection to the Thunder 9500™ V subsystem. The remote execution link is a network connection to another PC to allow you to execute CCI functions remotely. The connection between the CCI instance and the 9500V illustrates the connection between the CCI software on the host and the command device. The command device accepts both Synchronous TrueCopy 9500V and ShadowImage 9500V CCI commands and communicates read and write I/Os between the host and the volumes on the 9500V. The host does not communicate Synchronous TrueCopy 9500V or ShadowImage 9500V commands directly to the volumes on the 9500V. The CCI commands are always sent through the 9500V command device.

Note: The Thunder 9500V command device must be defined using the Resource Manager 9500. For details on setting the command device, refer to the *Hitachi Freedom Storage™ Thunder 9500™ Synchronous TrueCopy 9500V User's Guide* (MK-92DF608) and the *Hitachi Freedom Storage™ Thunder 9500™ ShadowImage 9500V User's Guide* (MK-92DF607).

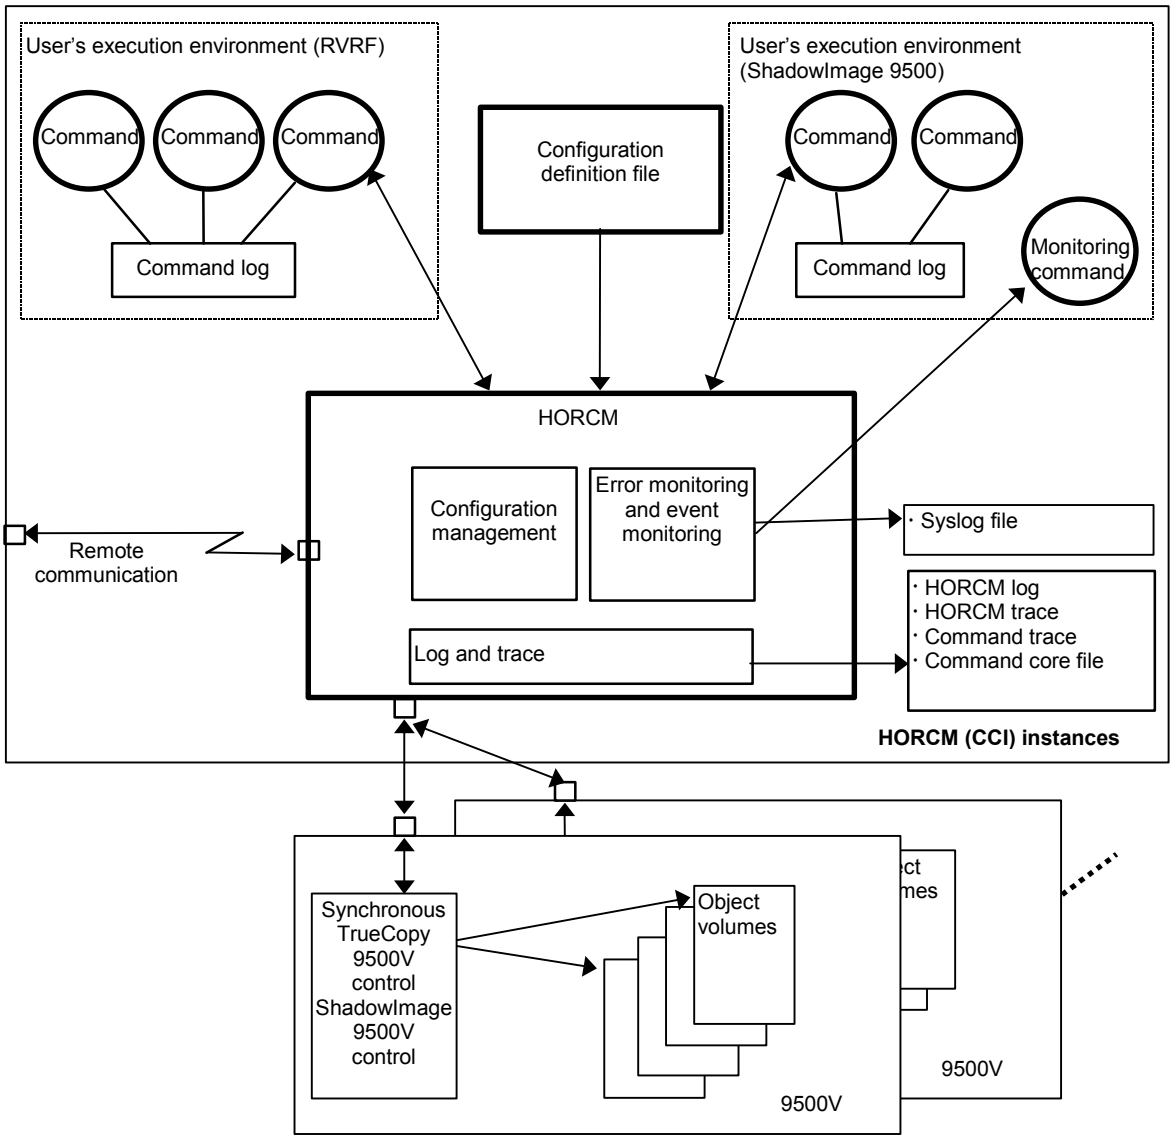

Figure 2.14 CCI Software Structure

2.5.3 Configuration Definition File

The CCI configuration definition file (HORCM_CONF) is the text file, which defines connected hosts and the volumes and groups known to the CCI instance. Physical volumes (special files) used independently by the servers are combined when paired logical volume names and group names are given to them. The configuration definition file describes the correspondence between the physical volumes used by the servers and the paired logical volumes and the names of the remote servers connected to the volumes. See section 2.6 for sample CCI configurations and their configuration definition file(s).

Figure 2.15 illustrates the configuration definition of paired volumes. Figure 2.16 shows a sample configuration file for a UNIX®-based operating system. Figure 2.17 shows a sample configuration file for the Windows NT® or Windows® 2000 operating system.

The CCI provides a sample configuration definition file (HORCM_CONF), so that the system administrator can copy this file to set necessary parameters and locate it in the specified directory. For details, see the HORCM_MON described later in this section.

The configuration definition file can be created automatically using the mkconf command tool (the user must customize the contents depending on the management). The value for the poll(10ms) must be specified manually. For details on the mkconf command tool, see section 4.15.2.

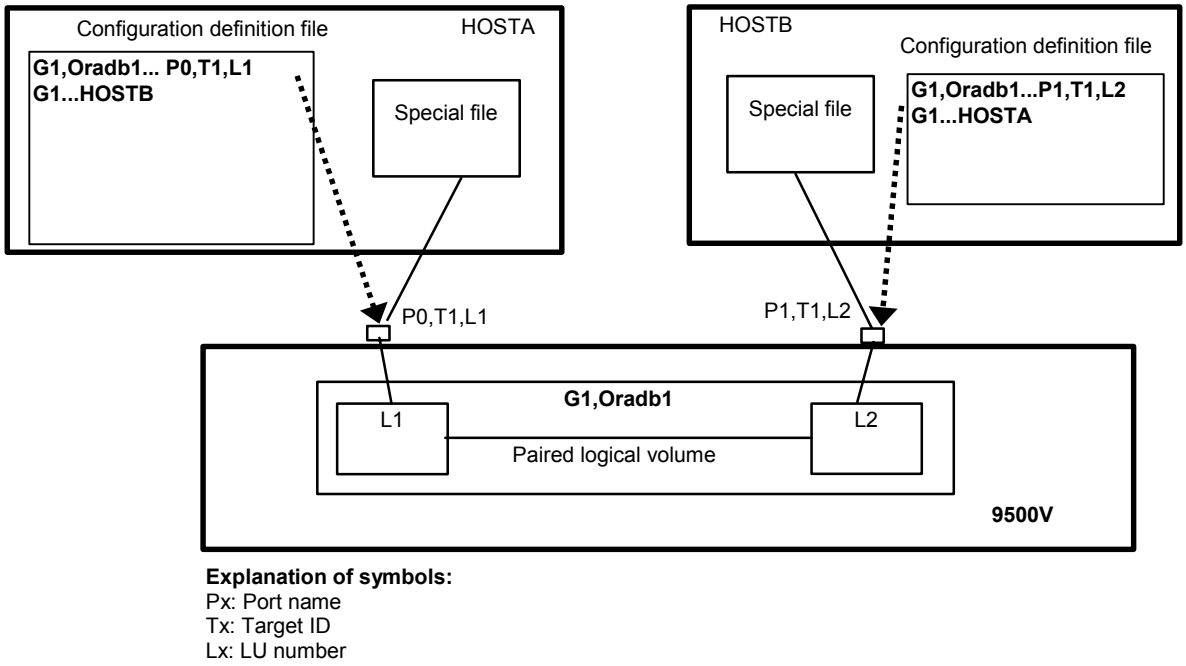

Figure 2.15 Configuration Definition of Paired Volumes

The following figures display an example of HORCM_CONF:

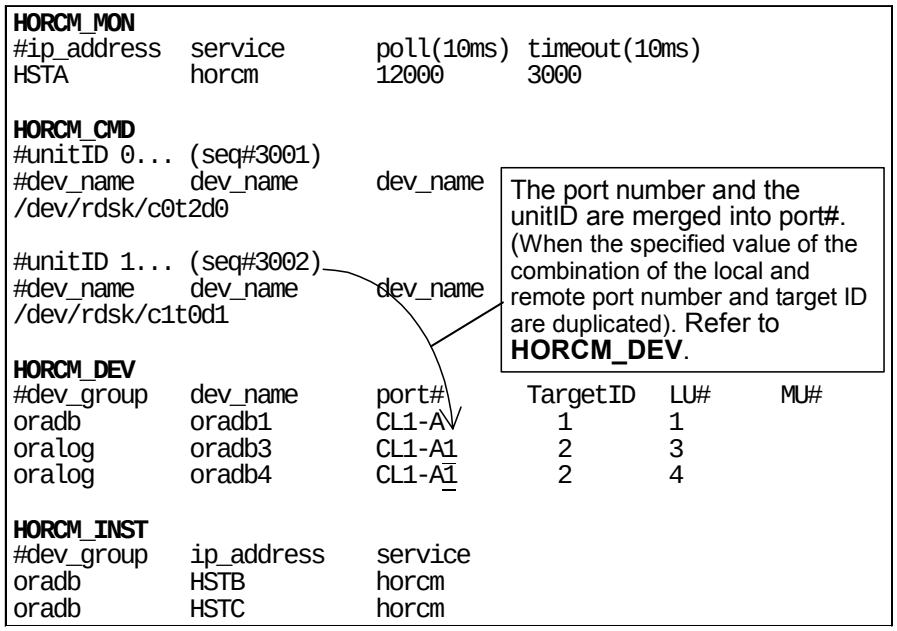

Figure 2.16 Configuration File Example Using Synchronous TrueCopy 9500V -- UNIX®-Based Servers

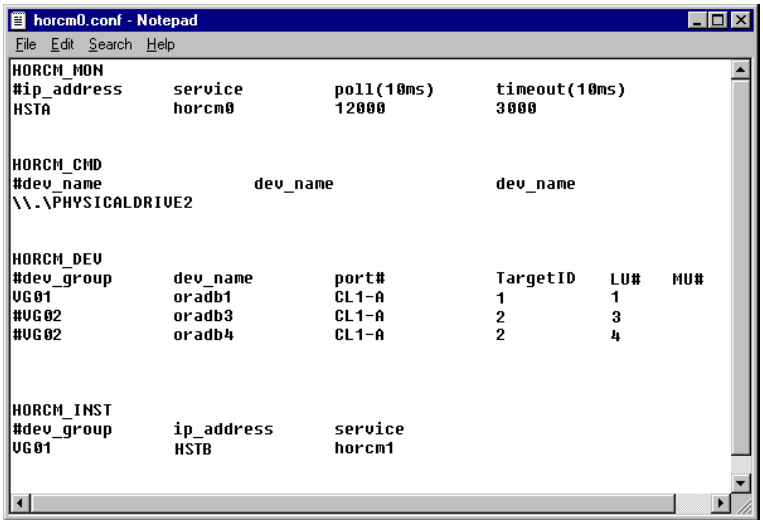

Figure 2.17 Configuration File Example Using Synchronous TrueCopy 9500V -- Windows NT®/2000 Servers

The procedure for creating and editing the configuration definition file manually is as follows:

1. Shutdown the HORCM (horcmshutdown).
2. Open the configuration definition file (HORCM_CONF) using the text editor. Set the parameters for **HORCM_MON** and **HORCM_CMD**. Save the file. ***Note:*** For details on configuration parameters, see section 3.4.
3. Start the HORCM (horcmstart).
4. Execute the raidscan command and write down the target IDs displayed in the execution result.
5. Shutdown the HORCM (horcmshutdown).
6. Open the configuration definition file (HORCM_CONF). Using the text editor, set the target ID based on the memo you took in step 4 above for the **HORCM_DEV** parameter.
7. Set the **HORCM_INST** parameter, and then save the configuration definition file.

The following items are in the configuration definition file:

HORCM_MON. The monitor parameter (HORCM_MON) defines the following values:

- **Ip_address:** The IP address of the local host. When HORCM has two or more network addresses on different subnets for communication, this item must be NONE.
- **Service:** The port name assigned to the CCI service (registered in the **/etc/services** file). The service parameter defines the CCI instance that runs on the local host. If a port number is specified instead of a port name, the port number will be used.
- **Poll(10ms):** The interval for monitoring paired volumes. To reduce the HORCM daemon load, make this interval longer. For Synchronous TrueCopy 9500V and/or ShadowImage 9500V operations, you must always set a value more than or equal to 6000. To calculate the poll(10ms) value, see the following equation and the example. Setting the value incorrectly may cause an internal conflict between CCI and the subsystem; the internal processing of the subsystem suspends temporarily. Processing may not proceed. If the interval is set to -1, the paired volumes are not monitored. The value of -1 is specified when two or more CCI instances run on a single machine.

  Calculating the value for poll(10ms):
  6000 × the number of all CCI instances that controls the subsystem, which its host is connected to the subsystem.

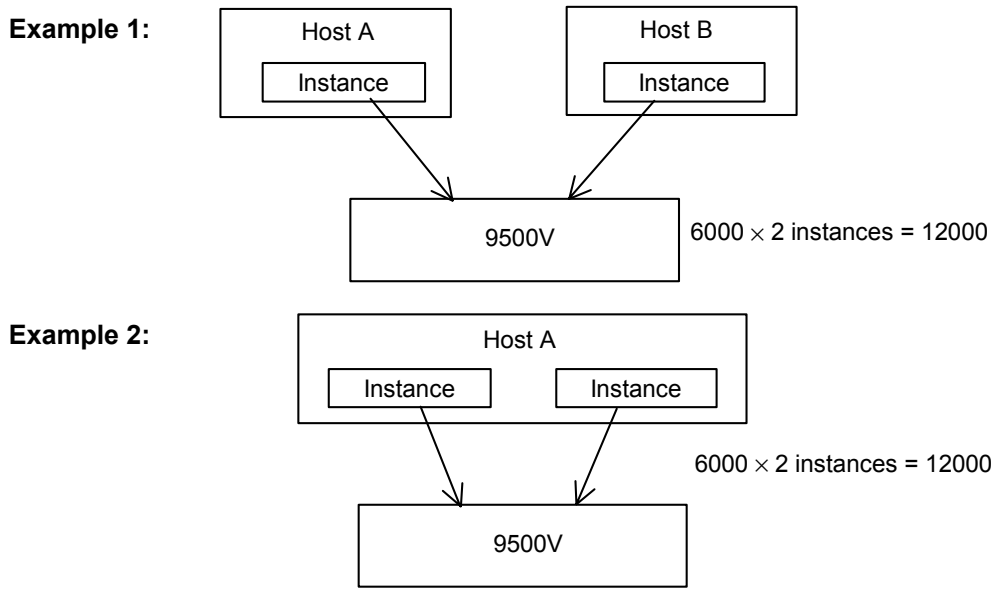

Figure 2.18 Setting the poll(10ms) value

- **Timeout(10ms):** The time-out period of communication with the remote server.

HORCM_CMD. The command parameter (HORCM_CMD) defines the UNIX® device path or Windows® physical device number of the 9500V command device. The command device must be mapped to fibre using the Resource Manager 9500. You can define two command devices to provide failover when the original command device becomes unavailable. For details on command devices, see section 2.5.4. For details on the alternate command device function, see section 2.5.5.

When a server is connected to two or more Thunder 9500™ V subsystems, the HORCM identifies each 9500V using the unit ID (see Figure 2.19). The unit ID is assigned sequentially in the order described in this section of the configuration definition file. If more than one command device (maximum of two command devices) is specified in a disk subsystem, the second command device has to be described side by side with the already described command device in a line. The server must be able to verify that the unit ID is the same Serial# (Serial ID) among servers when the 9500V is shared by two or more servers. This can be verified using the **raidqry** command.

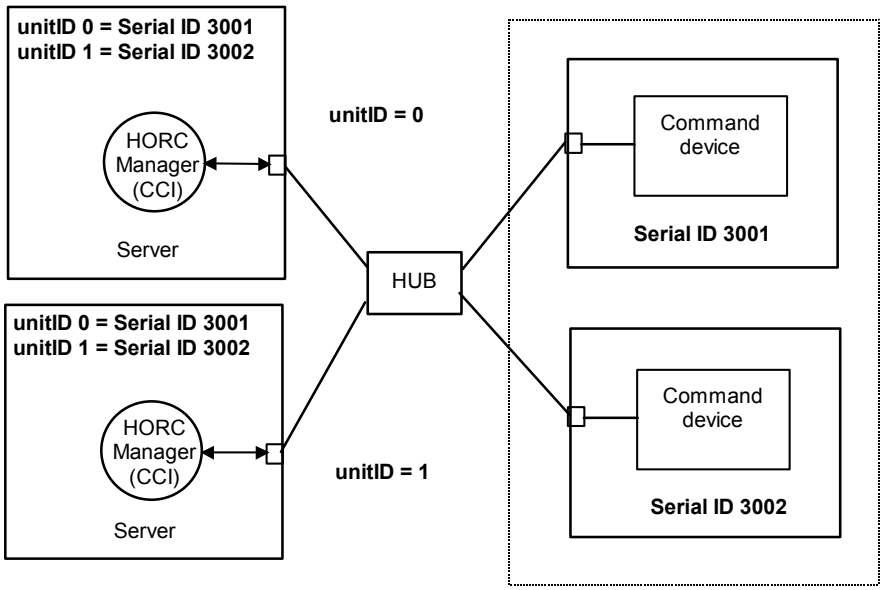

Figure 2.19 Configuration and Unit IDs for Multiple Thunder 9500™ V Subsystems

HORCM_DEV. The device parameter (HORCM_DEV) defines the 9500V device addresses for the paired logical volume names. When the server is connected to two or more 9500V subsystems, the 9500V unit ID is defined by the port# extension (when the specified value of the combination of the local and remote port number and target ID are duplicated. See Figure 2.16). Each group name is a unique name discriminated by a server that uses the volumes, the attributes of the volumes (such as database data, redo log file, UNIX® file), recovery level, etc. The group and paired logical volume names described in this item must reside in the remote server. The hardware SCSI/fibre bus, target ID, and LUN as hardware components need not be the same.

The following values are defined in the HORCM_DEV parameter:

- **dev_group:** Names a group of paired logical volumes. A command is executed for all corresponding volumes according to this group name.
- **dev_name:** Names the paired logical volume within a group (i.e., name of the special file or unique logical volume). The name of a paired logical volume must be different than the "dev name" on another group.
- **Port #:** Defines the 9500V port number of the volume that corresponds with the dev_name volume.

Table 2.6 Port Number of Volume Corresponding to dev_name Volume

Displayed by CCI	Port Name used in the Thunder 9500™ V Subsystem
CL1-A	Controller# 0, Port# A
CL1-B	Controller# 0. Port# B
CL2-A	Controller# 1. Port# A
CL2-B	Controller# 1. Port# B

- **Target ID:** Defines the fibre target ID number of the physical volume on the specified port. See Appendix D for further information on fibre address conversion.

  Note: The raidscan command cannot be executed while editing the configuration definition file. Therefore, in order to execute the raidscan command, edit the configuration definition file to the item HORCM_CMD, save the file, and then execute the raidscan command to obtain (write down) the target ID.

  Note: The conversion table for Windows NT/2000 is based on the Emulex driver. If a different fibre-channel adapter is used, the target ID indicated by the raidscan command may be different than the target ID indicated by the Windows NT/2000 system. In such case, for the configuration definition file, use the target ID that is displayed (obtained) by the raidscan -find command. For details on the conversion table, see Appendix D.

- **LU #:** Defines the fibre logical unit number (LU#) of the physical volume on the specified target ID and port.

  Note: For Host groups, when using fibre channel, if the target ID and LU# displayed on the system are different than the target ID on the fibre address conversion table, you must use the target ID and LU# indicated by the raidscan command in the CCI configuration file. Specify the 9500V LUN designation.

- **MU #:** Defines the remote copy number (zero fixed) of ShadowImage 9500V volumes. The MU#0 of a mirror descriptor is used for connection of the S-VOL.

HORCM_INST. The instance parameter (HORCM_INST) defines the network address (IP address) of the remote server (active or standby). It is used to refer to or change the status of the paired volume in the remote server (active or standby). When the primary volume is shared by two or more servers, there are two or more remote servers using the secondary volume. Thus, it is necessary to describe the addresses of all of these servers.

The following values are defined in the HORCM_INST parameter:

- **dev_group:** The server name described in dev_group of HORCM_DEV
- **ip_address:** The network address of the specified remote server
- **service:** The port name assigned to the HORCM communication path (registered in the /etc/services file). If a port number is specified instead of a port name, the port number will be used.

When HORCM has two or more network addresses on different subnets for communication, the ip_address of HORCM_MON must be NONE. This configuration for multiple networks can be found using **raidqry -r <group>** command option on each host. The current network address of HORCM can be changed using **horcctl -NC <group>** on each host.

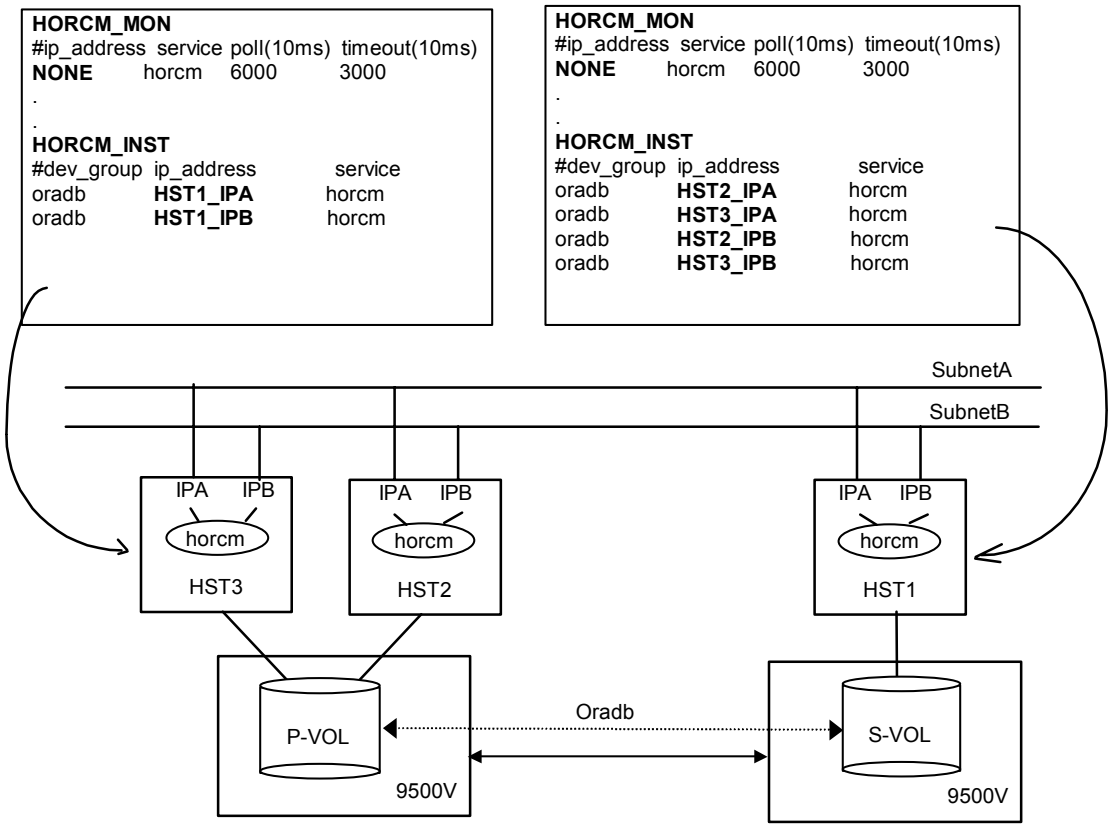

Figure 2.20 Configuration for Multiple Networks

2.5.4 Command Device

The Synchronous TrueCopy 9500V/ShadowImage 9500V commands are issued by the HORC Manager (HORCM) to the 9500V command device. The command device is a user-selected, dedicated logical volume on the Thunder 9500™ V subsystem which functions as the interface to the CCI software on the UNIX®/PC host. The command device is dedicated to CCI communications and cannot be used by any other applications. The command device accepts Synchronous TrueCopy 9500V and ShadowImage 9500V read and write commands that are executed by CCI.

The volume designated as the command device is used only by the 9500V and is blocked from the user. Set more than or equal to 65,538 blocks (1 block = 512 bytes) (33 MBs) for the command device LU.

WARNING: Do not write any user data on the volume that is to be selected as the command device.

The CCI software on the host issues reads and writes commands to the 9500V command device. When CCI receives an error notification in reply to a read or write request to the 9500V, the CCI software will activate an alternate command device, if one is defined. If a command device is blocked, you can activate an alternate command device manually. If no alternate command device is defined or available, all Synchronous TrueCopy 9500V and ShadowImage 9500V commands will terminate abnormally, and the host will not be able to issue commands to the subsystem. The user must set two command devices and use the alternate command device facility if a path error occurs or if the command device is blocked (the maximum is two command devices).

When you use the Synchronous TrueCopy 9500V function, the command devices must be set on both the local and remote disk subsystems.

Each command device must be set using the Resource Manager 9500. Each command device must also be defined in the HORCM_CMD section of the configuration file for the CCI instance on the attached host. If an alternate command device is not defined in the configuration file, the CCI software may not be able to use the device. See sections 2.5.3 and 2.6 for details.

2.5.5 Alternate Command Device Function

The CCI software issues commands to the 9500V command device via the UNIX®/PC raw I/O interface. If the command device fails in any way, all Synchronous TrueCopy 9500V and ShadowImage 9500V commands are terminated abnormally, and the user cannot use any commands. Because the use of alternate I/O pathing depends on the platform, restrictions are placed upon it. For example, on HP-UX® systems only devices subject to the LVM can use the alternate path PV-LINK. To avoid command device failure, CCI supports an alternate command device function.

- **Definition of alternate command devices.** To use an alternate command device, you must define two command devices for the HORCM_CMD item in the configuration definition file (see sections 2.5.3 and 2.6). When two command devices are defined, they are recognized as alternate command devices.
- **Timing of alternate command devices.** When the HORCM receives an error notification in reply from the operating system via the raw I/O interface, the command device is alternated. It is possible to alternate the command device forcibly by issuing an alternating command (**horcctl -C**).
- **Operation of alternating command.** If the command device will be blocked due to online maintenance (e.g., microcode replacement), the alternating command should be issued in advance. When the alternating command is issued again after completion of the online maintenance, the previous command device is activated again.
- **Multiple command devices on HORCM startup.** If at least one command device is available when one or more command devices were described in the configuration definition file, HORCM will be able to start with a warning message to the startup log by using an available command device. The user needs to confirm that all command devices can be changed by using the **horcctl -C** command option, or HORCM has been started without a warning message to the HORCM start up log.

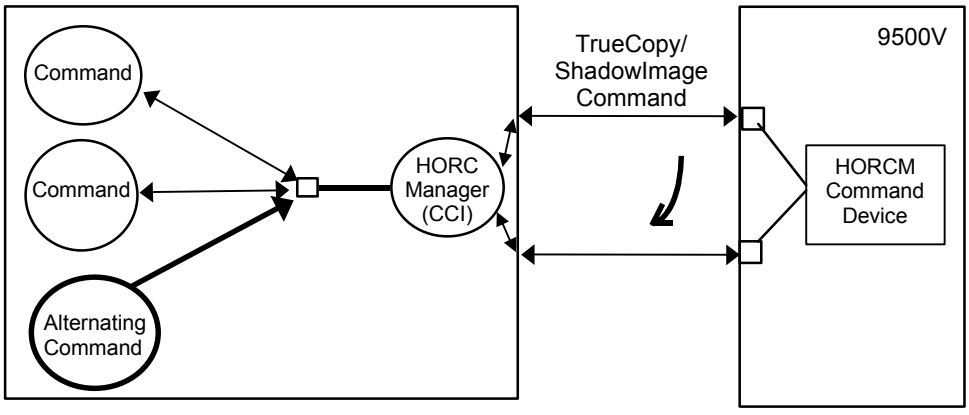

Figure 2.21 Alternate Command Device Function

For an example of setting two command devices, see section 2.6.3.

2.5.6 Protection Function

The CCI protection function protects a volume that cannot be recognized by the hosts from being operated (such as improper pair operation). The CCI identifies the existence of the environment variables when identifying the command device. If the current command device is specified to use the protection function, the protection function starts activating. Figure 2.22 shows the definition of the protected volumes.

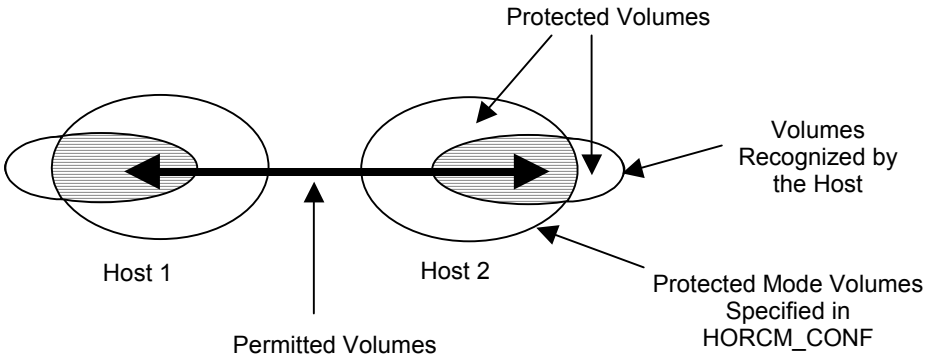

Figure 2.22 Definition of Protected Volumes

This protection function becomes effective by setting the environment variable $HORCMPROMOD. For details on setting the environment variable, see section 4.2.

Note: When the host cannot recognize an LU due to an error caused while the protection function is ON, operation to the pair of the unrecognized LU cannot be performed. In this case, set the LU so that it can be recognized by the host, or disable the protection function (OFF) once, remove the error cause, and then set the LU again to be recognized by the host so that the pair operation can be performed.

- **Protection mode and mirror descriptor.** Volumes described in the configuration definition file (HORCM_CONF) are targets for the protection function and are managed by each mirror descriptor (MU#0). The protection mode enables CCI to check whether the volumes described in the configuration definition file match all volumes recognized from the host at the time when CCI is activated. Permitted volumes are then registered in HORCM. Permitted volumes are volumes that are recognized from the host AND the mirror descriptors that are registered in the configuration definition file.

Table 2.7 Permitted Volumes and Mirror Descriptor

Volumes described in HORCM_CONF	Mirror descriptors described in HORCM_CONF			
	for TrueCopy 9500V		for ShadowImage 9500V	
			MU#0	
	E	none	E	none
Unknown				
/dev/rdsk/c0t0d0				
Unknown				

Permitted Volume

E — Mirror descriptor registered in HORCM_CONF

Unknown — Volumes that are described in HORCM_CONF but are not recognized from the host.

Notes:

- For the Fibre Switch environment, volumes that are recognized from the host must be set using the LUN Security feature.
- Volume matching test is performed when operating pairs. Inhibited volumes are rejected as EX_ENPERM.

- **Examples of a protected volume configuration.**
  - **Example of one host:** Group Ora1 and Ora2 in protection mode cannot recognize volume Grp2 and Grp4 from Host 1. The pair operation is inhibited. If the protection mode is OFF, the pair operation is permitted.

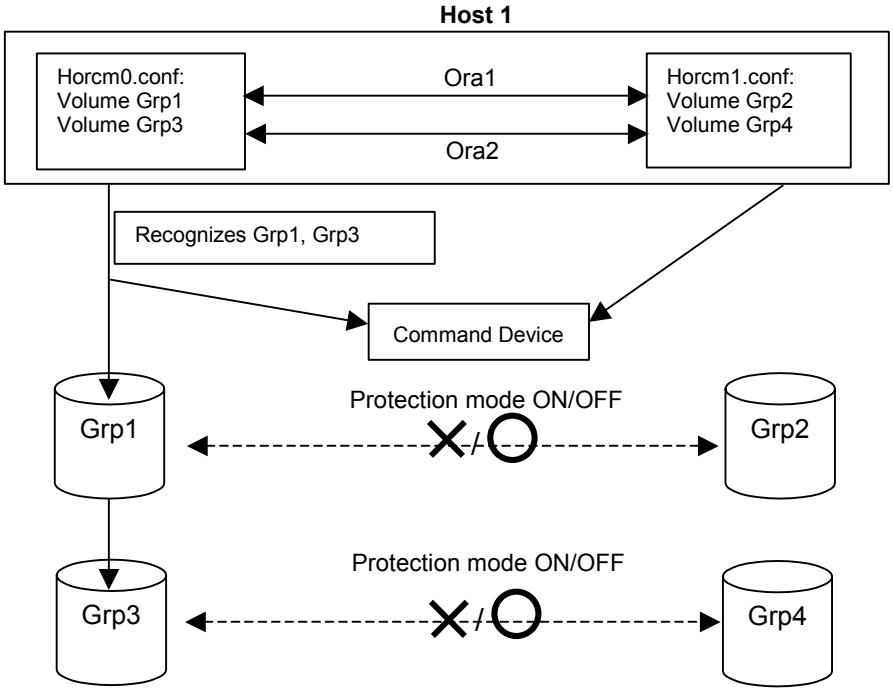

Figure 2.23 Protected Volume Configuration (one CCI host) (1)

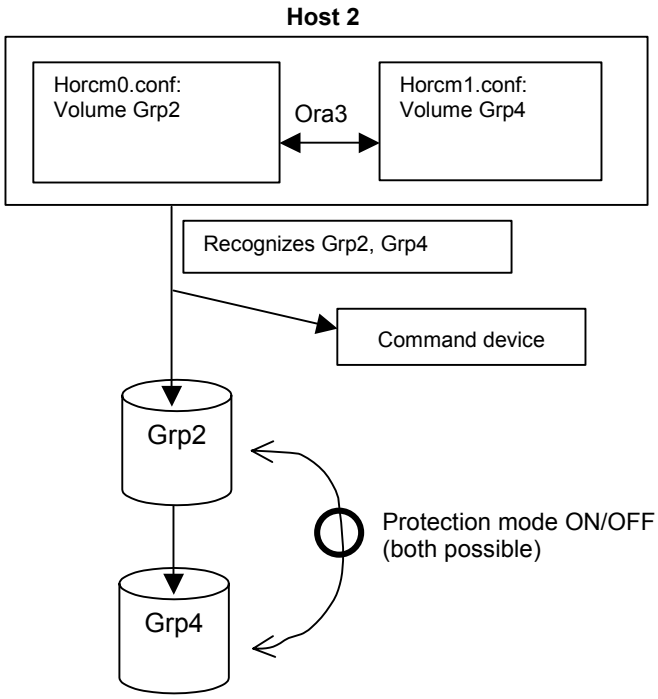

Figure 2.24 Protected Volume Configuration (one CCI host) (2)

— **Example of two CCI hosts:** Group Ora2 in protection mode cannot recognize volume Grp4 from Host 2. The pair operation is inhibited.

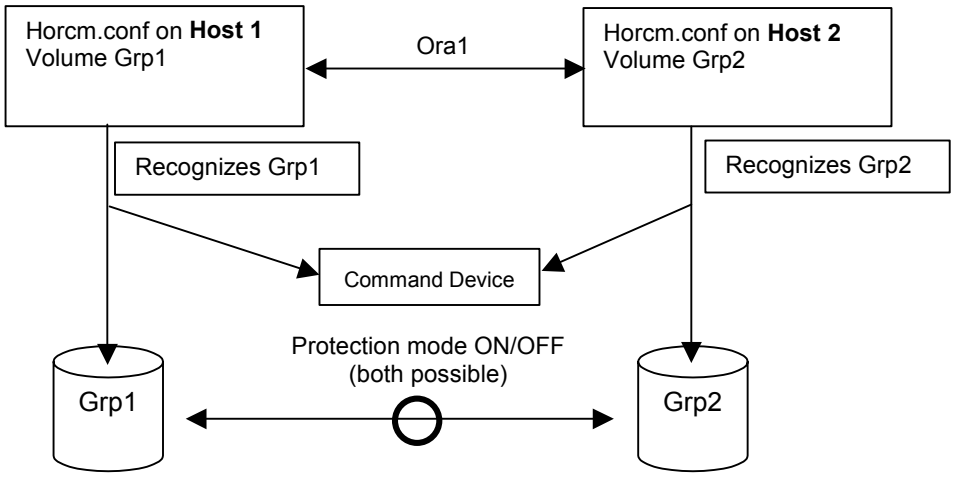

Figure 2.25 Protected Volume Configuration (two CCI hosts) (1)

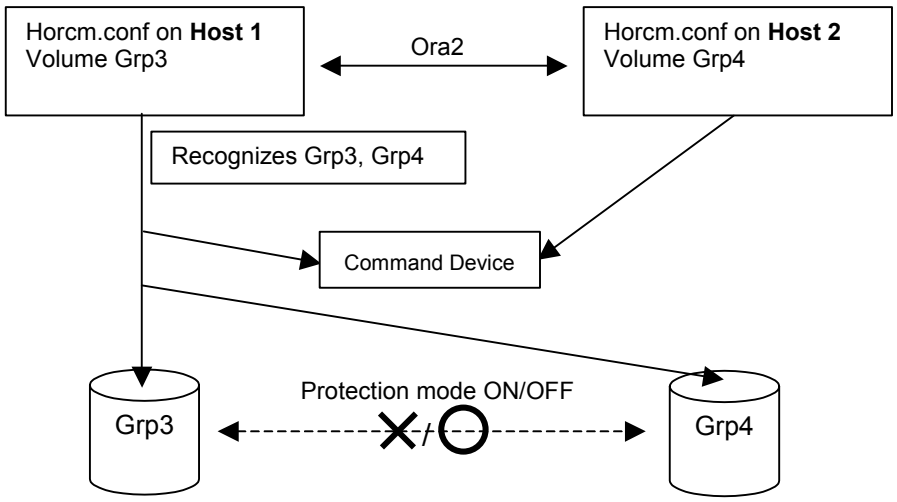

Figure 2.26 Protected Volume Configuration (two CCI hosts) (2)

- **Usage of commands.** Since volumes described in the configuration definition file (HORCM_CONF) are targets for the protection function, the following commands are also targets for the protection function (with the exception of pairdisplay).
  - The horctakeover, paircurchk, paircreate, pairsplit, pairresync, pairvolchk, and pairevtwait commands are the target commands. Command operation during protection mode rejects volumes (EX_ENPERM) that are not permitted.
  - The pairdisplay command displays volumes without LDEV attached information, and displays LDEV# "****" (for –CLI, "-").

Example:

```
# pairdisplay -g oradb
Group   PairVol(L/R) (Port#,TID,LU-M),Seq#, LDEV#.P/S,Status, Seq#,P-LDEV# M
oradb   oradev1(L)  (CL1-A , 3,  0-0) 3501 ****..----  ----,-----   ----  -
oradb   oradev1(R)  (CL1-A , 3,  1-0) 3501 ****..----  ----,-----   ----  -
```

  - The raidscan command is not affected by the HORCM_DEV and HORCM_INST section in the configuration definition file. The raidscan command is not the target of the protection function.

    You can see the correspondence of the permitted volumes and the DEVICE_FILE by using the raidscan –find option. In the configuration definition file, describing UID, PORT, TARG, and LUN of the volume specified by the –find option usually does not reject volumes as EX_PERM.

Example (for HP-UX):

```
# ioscan  -fun | grep rdsk | raidscan –find
DEVICE_FILE         UID  S/F PORT   TARG  LUN    SERIAL  LDEV  PRODUCT_ID
/dev/rdsk/c0t3d0     0    F  CL1-A     3     0     3501    17  DF600F
/dev/rdsk/c0t3d1     0    F  CL1-A     3     1     3501    18  DF600F
```

- **Registering permitted volumes.** The protection mode enables CCI to check whether the volumes described in the configuration definition file matches all the volumes recognized from the host when CCI is activated. Permitted volumes are then registered in HORCM. The permitted volumes are registered in /etc/horcmgr in the following steps.

  If the permitted file ($HORCMPERM) exists: CCI executes the following to the target volumes described in this file (e.g. If you want to permit only the LVM within your host, describe LVM volume in $HORCMPERM).

  — **For HP-UX®:** The default file name is /etc/horcmperm.conf or/etc/horcmperm*.conf (* as an instance number). CCI automatically executes cat $HORCMPERM | /HORCM/usr/bin/raidscan –find inst

Example:

```
# The following are an example to permit the LVM Volume groups.
# For MU# 0
vg00  /dev/rdsk/c0t3d0  /dev/rdsk/c0t3d1
vg00  /dev/rdsk/c0t3d2  /dev/rdsk/c0t3d3
```

Confirming vg01 groups: The following is an example for confirming whether the LVM volume group vg01 is correctly mapped to the group described in the configuration definition file.

Example:

```
# export HORCC_MRCF=1
# cat /etc/horcmperm.conf | grep vg01 | raidscan -find verify 1 -fd
```

OR

```
# vgdisplay -v /dev/vg01|grep dsk|sed 's/\/*\/dsk\//\/rdsk\//g'|raidscan -find verify 1 –fd
DEVICE_FILE        Group    PairVol     Device_File      M   SERIAL  LDEV
/dev/rdsk/c0t3d0   oradb1   oradev1     c0t3d0           1    3501    17
/dev/rdsk/c0t3d1   oradb1   oradev2     c0t3d1           1    3501    18
/dev/rdsk/c0t3d2   oradb    oradev3     c0t3d2           1    3501    19
/dev/rdsk/c0t3d3   -        -           -                1    3501    20
```

As shown in the example, the device file /dev/rdsk/c0t3d2 is mapped to the other group, and the /dev/rdsk/c0t3d3 is not described in the configuration definition file.

— **For Windows NT®/2000:** The default file name is \WINNT\horcmperm.conf or \WINNT\horcmperm*.conf (* as an instance number). CCI automatically executes type $HORCMPERM | x:\HORCM\etc\raidscan.exe –find inst

Example:

```
# The following are an example to permit the DB Volumes.
# Note: a numerical value is interpreted as Harddisk#.
# DB0 For MU# 0
Hd0-10
harddisk12  harddisk13  harddisk17
```

Example (Confirming DB1 groups):

```
set HORCC_MRCF=1
echo hd20-23 | raidscan -find verify 1 –fd
DEVICE_FILE          Group    PairVol     Device_File      M   SERIAL  LDEV
Harddisk20           oradb1   oradev1     Harddisk20       1    3501    17
Harddisk21           oradb1   oradev2     Harddisk21       1    3501    18
Harddisk22           oradb    oradev3     Harddisk22       1    3501    19
Harddisk23           -        -           -                1    3501    20
```

As shown in the example, Harddisk22 is mapped to the other group, and Harddisk23 is not described in the configuration definition file.

If the permitted file ($HORCMPERM) does not exist: CCI targets the volumes of your own host, and executes the following:

For HP-UX®:

```
‘ioscan  -fun | grep rdsk | /HORCM/usr/bin/raidscan -find inst’
```

For Solaris®:

```
‘ls /dev/rdsk/* | /HORCM/usr/bin/raidscan -find inst’
```

For AIX®:

```
‘ls –C –c disk | grep hdisk | /HORCM/usr/bin/raidscan -find inst’
```

For Linux®:

```
‘ls /dev/sd* | /HORCM/usr/bin/raidscan -find inst’
```

For Tru64 UNIX®:

```
‘ls /dev/rdisk/dsk* | /HORCM/usr/bin/raidscan -find inst’
```

For IRIX®:

```
‘ls /dev/rdsk/*vol /dev/rdsk/*/*vol/* | /HORCM/usr/bin/raidscan -find inst ‘
```

For Windows NT®/2000:

```
‘echo hd0-999 | x:\HORCM\etc\raidscan.exe -find inst’
```

Note: The default target number for scanning is 1,000.

Additional Notes: Since CCI automatically registers device files when CCI is activated, starting up may take time. To fasten start up by using the non-protection mode, create an empty HORCMPERM file (with no size) as a dummy to prohibit scanning. The file name displayed by the –fd option will be "Unknown". To display the file name using the –fd option, execute raidscan –find inst manually.

- **Confirming the protection mode command devices used by HORCM.** To display the protection mode command device used by HORCM, use the horcctl –D option. The device file name is displayed with "*" added.

```
# horcctl –D
Current control device = /dev/rdsl/c0t0d0*
```

Figure 2.27 Confirming the Protection Mode Command Device (HP-UX®)

2.5.7 CCI Software Files

The CCI software product consists of files supplied to the user, log files created internally, and files created by the user. These files are stored on the local disk in the server machine. Table 2.8 lists the CCI files which are provided for UNIX®-based systems. Table 2.9 lists the CCI files which are provided for Windows®-based systems.

Table 2.8 CCI Files for UNIX®-based Systems

No.	Title	File name	Command name	Mode	User*	Group
01	HORCM	/etc/horcmgr	horcmd*	0544	root	sys
02	HORCM_CONF	/HORCM/etc/horcm.conf	–	0444	root	sys
03	Takeover	/usr/bin/horctakeover	horctakeover*	0544	root	sys
04	Accessibility check	/usr/bin/paircurchk	paircurchk*	0544	root	sys
05	Pair generation	/usr/bin/paircreate	paircreate	0544	root	sys
06	Pair splitting	/usr/bin/pairsplit	pairsplit	0544	root	sys
07	Pair resynchronization	/usr/bin/pairresync	pairresync	0544	root	sys
08	Event waiting	/usr/bin/pairevtwait	pairevtwait	0544	root	sys
09	Error notification	/usr/bin/pairmon	pairmon	0544	root	sys
10	Volume check	/usr/bin/pairvolchk	pairvolchk	0544	root	sys
11	Pair configuration confirmation	/usr/bin/pairdisplay	pairdisplay	0544	root	sys
12	RAID scanning	/usr/bin/raidscan	raidscan	0544	root	sys
13	RAID activity reporting	/usr/bin/raidar	raidar	0544	root	sys
14	Connection confirming	/usr/bin/raidqry	raidqry	0544	root	sys
15	Trace control	/usr/bin/horcctl	horcctl	0544	root	sys
16	HORCM activation script	/usr/bin/horcmstart.sh	horcmstart.sh	0544	root	sys
17	HORCM shutdown script	/usr/bin/horcmshutdown.sh	horcmshutdown.sh	0544	root	sys
18	Connection confirming ***Note:*** Provided only for HP-UX®, and Solaris® systems.	/HORCM/usr/bin/inqraid	inqraid	0544	root	sys
19	Configuration file creating	/HORCM/usr/bin/mkconf.sh	mkconf	0544	root	sys
20	Synchronous waiting	/usr/bin/pairsyncwait	pairsyncwait*	0544	root	sys
21	Oracle Validation setting	/usr/bin/raidvchkset	raidvchkset	0544	root	sys
22	Oracle Validation displaying	/usr/bin/raidvchkdsp	raidvchkdsp	0544	root	sys
23	Oracle Validation scanning	/usr/bin/raidvchkscan	raidvchkscan	0544	root	sys

Note: For information and instructions on changing the UNIX® user for the CCI software, please see section 3.3.2.

Note*: ShadowImage 9500V/Synchronous TrueCopy 9500V does not support the horcmd, horctakeover, paircurchk, pairsyncwait, raidvchkset, raidvchkdsp, and the raidvchkscan command. However, Synchronous TrueCopy 9500V supports the horctakeover and the paircurchk command.

Table 2.9 CCI Files for Windows®-based Systems (Continued on next page)

No.	Title	File name	Command name
01	HORCM	\HORCM\etc\horcmgr.exe	horcmd*
02	HORCM_CONF	\HORCM\etc\horcm.conf	--
03	Takeover	\HORCM\etc\horctakeover.exe	horctakeover*
04	Accessibility check	\HORCM\etc\paircurchk.exe	paircurchk*
05	Pair generation	\HORCM\etc\paircreate.exe	paircreate
06	Pair splitting	\HORCM\etc\pairsplit.exe	pairsplit
07	Pair resynchronization	\HORCM\etc\pairresync.exe	pairresync
08	Event waiting	\HORCM\etc\pairevtwait.exe	pairevtwait
09	Error notification	\HORCM\etc\pairmon.exe	pairmon
10	Volume check	\HORCM\etc\pairvolchk.exe	pairvolchk
11	Pair configuration confirmation	\HORCM\etc\pairdisplay.exe	pairdisplay
12	RAID scanning	\HORCM\etc\raidscan.exe	raidscan
13	RAID activity reporting	\HORCM\etc\raidar.exe	raidar
14	Connection confirmation	\HORCM\etc\raidqry.exe	raidqry
15	Trace control	\HORCM\etc\horcctl.exe	horcctl
16	HORCM activation script	\HORCM\etc\horcmstart.exe	horcmstart
17	HORCM shutdown script	\HORCM\etc\horcmshutdown.exe	horcmshutdown
18	Synchronous waiting	\HORCM\etc\pairsyncwait.exe	pairsyncwait*
19	Connection confirming ***Note:*** Provided only for HP-UX®, and Solaris® systems.	\HORCM\usr\bin\inqraid	inqraid
20	Configuration file creating	\HORC\Tool\mkconf.sh	mkconf
21	Oracle Validation setting	\HORCM\usr\raidvchkset	raidvchkset*
22	Oracle Validation displaying	\HORCM\usr\raidvchkdsp	raidvchkdsp*
23	Oracle Validation scanning	\HORCM\usr\raidvchkscan	raidvchkscan*
24	Takeover	\HORCM\usr\bin\horctakeover.exe	horctakeover*
25	Accessibility check	\HORCM\usr\bin\paircurchk.exe	paircurchk*
26	Pair generation	\HORCM\usr\bin\paircreate.exe	paircreate
27	Pair splitting	\HORCM\usr\bin\pairsplit.exe	pairsplit
28	Pair resynchronization	\HORCM\usr\bin\pairresync.exe	pairresync
29	Event waiting	\HORCM\usr\bin\pairevtwait.exe	pairevtwait
30	Volume check	\HORCM\usr\bin\pairvolchk.exe	pairvolchk
31	Pair configuration confirmation	\HORCM\usr\bin\pairdisplay.exe	pairdisplay
32	RAID scanning	\HORCM\usr\bin\raidscan.exe	raidscan

Table 2.9 CCI Files for Windows®-based Systems (Continued)

No.	Title	File name	Command name
33	RAID connection confirmation	\HORCM\usr\bin\raidqry.exe	raidqry
34	Synchronous waiting	\HORCM\usr\bin\pairsyncwait.exe	pairsyncwait*
35	Oracle Validation setting	\HORCM\usr\bin\raidvchkset	raidvchkset*
36	Oracle Validation displaying	\HORCM\usr\bin\raidvchkdsp	raidvchkdsp*
37	Oracle Validation scanning	\HORCM\usr\bin\raidvchkscan	raidvchkscan*

Note: The commands in \HORCM\etc\ are used when you execute from the console window. If these commands are executed without an argument, the interactive mode will start up. The commands in \HORCM\usr\bin have no console window, and can therefore be used when you execute from the user application. The commands in \HORC\user\bin\ cannot mount directories.

Note*: ShadowImage 9500V/Synchronous TrueCopy 9500V does not support the horcmd, horctakeover, paircurchk, pairsyncwait, raidvchkset, raidvchkdsp, and the raidvchkscan command. However, Synchronous TrueCopy 9500V supports the horctakeover and the paircurchk command.

2.5.8 Log and Trace Files

The CCI software (HORCM) and Synchronous TrueCopy 9500V/ShadowImage 9500V commands maintain start-up log files, execution log files, and trace files which can be used to identify the cause of errors and maintain status transition history records of the paired volumes. Please refer to Appendix B for a complete description of the CCI log and trace files.

2.5.9 User-Created Files

Script Files. CCI supports scripting to provide automated and unattended copy operations. A CCI script contains a list of CCI commands, which describes a series of Synchronous TrueCopy 9500V and/or ShadowImage 9500V operations. The scripted commands for UNIX®-based platforms are defined in a shell script file. The scripted commands for Windows®-based platforms are defined in a text file. The host reads the script file and sends the commands to the 9500V command device to execute the Synchronous TrueCopy 9500V/ShadowImage 9500V operations automatically. The CCI scripts are:

- **HORCM startup script** (horcmstart.sh, horcmstart.exe): A script that starts HORCM (/etc/horcmgr), sets environmental variables as needed (e.g., HORCM_CONF, HORCM_LOG, HORCM_LOGS), and starts HORCM.
- **HORCM shutdown script** (horcmshutdown.sh, horcmshutdown.exe): A script for stopping the HORCM (/etc/horcmgr).
- **HA control script:** A script for executing takeover processing automatically when the cluster manager (CM) detects a server error.

When constructing the HORCM environment, the system administrator should make a copy of the **HORCM_CONF** file. The copied file should be set according to the system environment and registered as the following file (* is the instance number):

- UNIX®-based systems: **/etc/horcm.conf** or **/etc/horcm*.conf**
- Windows®-based systems: **\WINNT\horcm.conf** or **\WINNT\horcm*.conf**

2.6 Examples of Configuration Definition File

Figure 2.28 through Figure 2.29 show examples of CCI configurations, the configuration definition file(s) for each configuration, and examples of CCI command use for each configuration.

2.6.1 Two Hosts and Two Instances

The command device is defined using the system raw device name (character-type device file name). For example, the command devices for Figure 2.28 and Figure 2.29 would be:

- HP-UX®:

  HORCM_CMD of HOSTA = /dev/rdsk/c0t0d0
  HORCM_CMD of HOSTB = /dev/rdsk/c1t0d0

- Solaris®:

  HORCM_CMD of HOSTA = /dev/rdsk/c0t0d0s2
  HORCM_CMD of HOSTB = /dev/rdsk/c1t0d0s2

  The command device can be used without a label in the format command.

- AIX®:

  HORCM_CMD of HOSTA = /dev/rhdisk*X*
  HORCM_CMD of HOSTB = /dev/rhdisk*X*

  where *X* = device number created automatically by AIX®

- Tru64 UNIX®:

  HORCM_CMD of HOSTA = /dev/rdisk/dsk*X*c
  HORCM_CMD of HOSTB = /dev/rdisk/dsk*X*c

  where *X* = device number assigned by Tru64 UNIX®

- Linux®:

  HORCM_CMD of HOSTA = /dev/sd*X*
  HORCM_CMD of HOSTB = /dev/sd*X*

  where *X* = device number assigned by Linux®

- IRIX®:

  HORCM_CMD of HOSTA = /dev/rdsk/dks*X*d*X*l*X*vol
  OR
  HORCM_CMD of HOSTA = /dev/rdsk/node_wwn/lun*X*vol/c*X*p*X*

  HORCM_CMD of HOSTB = /dev/rdsk/dks*X*d*X*l*X*vol
  OR
  HORCM_CMD of HOSTB = /dev/rdsk/node_wwn/lun*X*vol/c*X*p*X*

  where *X* = device number assigned by IRIX®

- Windows NT®/2000:

  HORCM_CMD of HOSTA = \\.\PhysicalDrive*X*
  OR
  HORCM_CMD of HOSTA = \\.\Volume{guid} (Windows 2000 only)

  HORCM_CMD of HOSTB = \\.\PhysicalDrive*X*
  OR
  HORCM_CMD of HOSTB = \\.\Volume{guid} (Windows 2000 only)

  where *X* = device number assigned by Windows® NT/2000

  The PhysicalDrive number may change at every reboot. If the number changes, use Volume{guid} for which the same name is kept.

  If "\\.\Volume{guid}" is specified, CCI changes it to "\\.\PhysicalDrive?" to be corresponded. Volume{guid} is created when you make a partition by using the Windows' Disk Management. When the **inqraid $Volume -CLI -fv** or **raidscan -x findcmddev0.?** commands are used, you can find Volume{guid}.

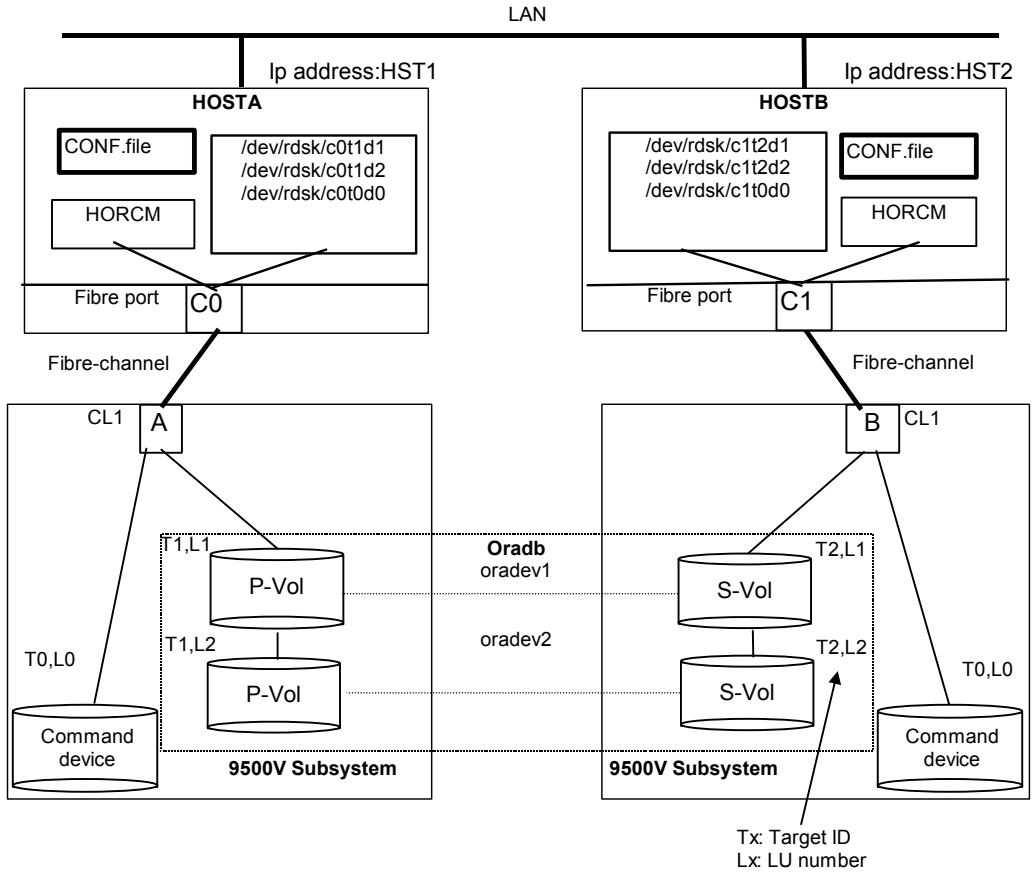

Configuration file for HOSTA (/etc/horcm.conf)

```
HORCM_MON
#ip_address  service   poll(10ms)   timeout(10ms)
HST1         horcm     12000 [Note 1]     3000

HORCM_CMD
#dev_name
/dev/xxx  [Note 2]

HORCM_DEV
#dev_group  dev_name  port#  TargetID  LU#  MU#
 Oradb      oradev1   CL1-A      1      1
 Oradb      oradev2   CL1-A      1      2

HORCM_INST
#dev_group     ip_address      service
 Oradb         HST2            horcm
```

Configuration file for HOSTB (/etc/horcm.conf)

```
HORCM_MON
#ip_address  service   poll(10ms)   timeout(10ms)
HST2         horcm     12000 [Note 1]     3000

HORCM_CMD
#dev_name
 /dev/yyy  [Note 2]

HORCM_DEV
#dev_group  dev_name  port#  TargetID  LU#  MU#
 Oradb      oradev1   CL1-B      2      1
 Oradb      oradev2   CL1-B      2      2

HORCM_INST
#dev_group     ip_address      service
 Oradb         HST1            horcm
```

Note 1: To calculate the value for poll(10ms), see section 2.5.3.
Note 2: The command device is dedicated to CCI communications and cannot be used by any other applications (neither the user). Command devices must be set using the Resource Manager 9500 program. If setting two command devices, define two command devices in the HORCM_CMD section.

Figure 2.28 Two Hosts and Two Instances Example Synchronous TrueCopy 9500V

Example of CCI commands with HOSTA (group Oradb): Synchronous TrueCopy 9500V

- Designate a group name (Oradb) and a local host P-VOL a case.
  # paircreate -g Oradb -vl
  This command creates pairs for all LUs assigned to group **Oradb** in the configuration definition file (two pairs for the configuration in Figure 2.28).

- Designate a volume name (oradev1) and a local host P-VOL a case.
  # paircreate -g Oradb -d oradev1 -vl
  This command creates pairs for all LUs designated as **oradev1** in the configuration definition file (CL1-A, T1, L1 and CL1-B, T2, L1 for the configuration in Figure 2.28).

- Designate a group name and display pair status.
  # pairdisplay -g Oradb

```
Group PairVol(L/R) (Port#,TID,LU),Seq#,LDEV#.P/S,Status,Fence, Seq#,P-LDEV# M
oradb  oradev1(L)   (CL1-A  ,  1,  1)  3005      18..P-VOL COPY NEVER , 3006   20  -
oradb  oradev1(R)   (CL1-B  ,  2,  1)  3006      20..S-VOL COPY NEVER , -----    18  -
oradb  oradev2(L)   (CL1-A  ,  1,  2)  3005      19..P-VOL COPY NEVER , 3006   21  -
oradb  oradev2(R)   (CL1-B  ,  2,  2)  3006      21..S-VOL COPY NEVER , -----    19  -
```

Example of CCI commands with HOSTB (group Oradb): Synchronous TrueCopy 9500V

- Designate a group name and a remote host P-VOL a case.
  # paircreate -g Oradb -f never -vr
  This command creates pairs for all LUs assigned to group **Oradb** in the configuration definition file (two pairs for the configuration in Figure 2.28).

- Designate a volume name (oradev1) and a remote host P-VOL a case.
  # paircreate -g Oradb -d oradev1 -f never -vr
  This command creates pairs for all LUs designated as **oradev1** in the configuration definition file (CL1-A, T1, L1 and CL1-B, T2, L1 for the configuration in Figure 2.28).

- Designate a group name and display pair status.
  # pairdisplay -g Oradb

```
Group PairVol(L/R) (Port#,TID,LU),Seq#,LDEV#.P/S,Status,Fence, Seq#,P-LDEV# M
oradb  oradev1(L)   (CL1-B  ,  2,  1)  3006      20..S-VOL COPY NEVER , -----    18  -
oradb  oradev1(R)   (CL1-A  ,  1,  1)  3005      18..P-VOL COPY NEVER , 3006   20  -
oradb  oradev2(L)   (CL1-B  ,  2,  2)  3006      21..S-VOL COPY NEVER , -----    19  -
oradb  oradev2(R)   (CL1-A  ,  1,  2)  3005      19..P-VOL COPY NEVER , 3006   21  -
```

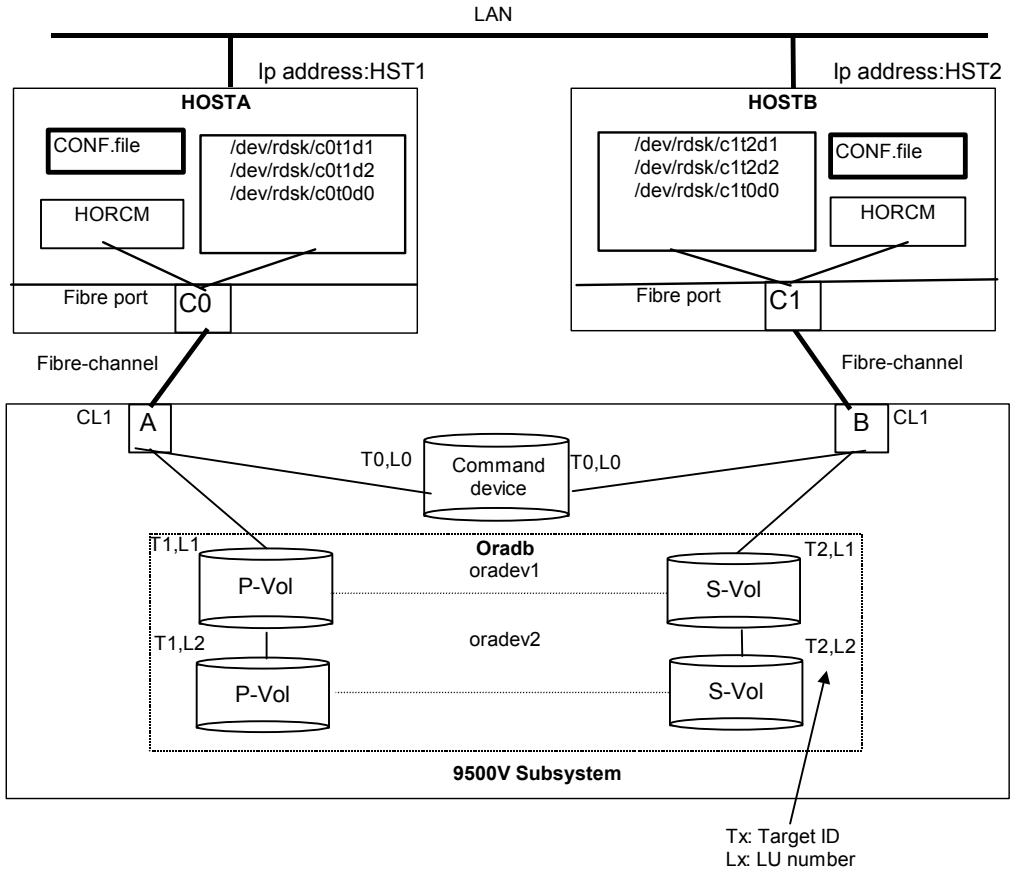

Configuration file for HOSTA (/etc/horcm.conf)

```
HORCM_MON
#ip_address  service   poll(10ms)   timeout(10ms)
HST1         horcm     12000 [Note 1]     3000

HORCM_CMD
#dev_name
/dev/xxx  [Note 2]

HORCM_DEV
#dev_group  dev_name  port#  TargetID  LU#  MU#
 Oradb      oradev1   CL1-A      1      1     0
 Oradb      oradev2   CL1-A      1      2     0

HORCM_INST
#dev_group      ip_address      service
 Oradb          HST2            horcm
```

Configuration file for HOSTB (/etc/horcm.conf)

```
HORCM_MON
#ip_address  service   poll(10ms)   timeout(10ms)
HST2         horcm     12000 [Note 1]     3000

HORCM_CMD
#dev_name
 /dev/yyy  [Note 2]

HORCM_DEV
#dev_group  dev_name  port#  TargetID  LU#  MU#
 Oradb      oradev1   CL1-B      2      1     0
 Oradb      oradev2   CL1-B      2      2     0

HORCM_INST
#dev_group      ip_address      service
 Oradb          HST1            horcm
```

Note 1: To calculate the value for poll(10ms), see section 2.5.3.
Note 2: The command device is dedicated to CCI communications and cannot be used by any other applications (neither the user). Command devices must be set using the Resource Manager 9500V program. If setting two command devices, define two command devices in the HORCM_CMD section.

Figure 2.29 Two Hosts and Two Instances Example ShadowImage 9500V

Example of CCI commands with HOSTA (group Oradb): ShadowImage 9500V

- When the command execution environment is not set, set HORCC_MRCF to the environment variable.

  For # setenv HORCC_MRCF 1
  Windows®: set HORCC_MRCF=1

- Designate a group name (Oradb) and a local host P-VOL a case.
  # paircreate -g Oradb -vl

  This command creates pairs for all LUs assigned to group **Oradb** in the configuration definition file (two pairs for the configuration in Figure 2.29).

- Designate a volume name (oradev1) and a local host P-VOL a case.
  # paircreate -g Oradb -d oradev1 -vl

  This command creates pairs for all LUs designated as **oradev1** in the configuration definition file (CL1-A, T1, L1 and CL1-B, T2, L1 for the configuration in Figure 2.29 Two Hosts and Two Instances Example ShadowImage 9500V

- Designate a group name and display pair status.
  # pairdisplay -g Oradb

```
Group PairVol(L/R) (Port#,TID,LU-M),Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M
oradb  oradev1(L)   (CL1-A,  1,  1 - 0)    3005      18..P-VOL COPY  3005      20  -
oradb  oradev1(R)   (CL1-B,  2,  1 - 0)    3005      20..S-VOL COPY  -----     18  -
oradb  oradev2(L)   (CL1-A,  1,  2 - 0)    3005      19..P-VOL COPY  3005      21  -
oradb  oradev2(R)   (CL1-B,  2,  2 - 0)    3005      21..S-VOL COPY  -----     19  -
```

Example of CCI commands with HOSTB (group Oradb): ShadowImage 9500V

- When the command execution environment is not set, set HORCC_MRCF to the environment variable.

  For # setenv HORCC_MRCF 1
  Windows®: set HORCC_MRCF=1

- Designate a group name and a remote host P-VOL a case.
  # paircreate -g Oradb -vr

- This command creates pairs for all LUs assigned to group **Oradb** in the configuration definition file (two pairs for the configuration in Figure 2.29). Designate a volume name (oradev1) and a remote host P-VOL a case.

  # paircreate -g Oradb -d oradev1 -vr

  This command creates pairs for all LUs designated as **oradev1** in the configuration definition file (CL1-A, T1, L1 and CL1-B, T2, L1 for the configuration in Figure 2.29).

- Designate a group name and display pair status.
  # pairdisplay -g Oradb

```
Group PairVol(L/R) (Port#,TID,LU-M),Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M
oradb  oradev1(L)   (CL1-B,  2,  1 - 0)    3005      20..S-VOL COPY   -----      18  -
oradb  oradev1(R)   (CL1-A,  1,  1 - 0)    3005      18..P-VOL COPY  3005      20  -
oradb  oradev2(L)   (CL1-B,  2,  2 - 0)    3005      21..S-VOL COPY   -----      19  -
oradb  oradev2(R)   (CL1-A,  1,  2 - 0)    3005      19..P-VOL COPY  3005      21  -
```

2.6.2 One Host and Two Instances

The command device is defined using the system raw device name (character-type device file name). The command device defined in the configuration definition file must be established in a way to be following either every instance. If one command device is used between different instances on the same port, then the number of instances is up to 16 per command device. If this restriction is exceeded, then use a different path for each instance. For example, the command devices for Figure 2.30 and Figure 2.31 would be:

- HP-UX®:

  HORCM_CMD for HORCMINST0 = /dev/rdsk/c0t0d0
  HORCM_CMD for HORCMINST1 = /dev/rdsk/c1t0d0

- Solaris®:

  HORCM_CMD for HORCMINST0 = /dev/rdsk/c0t0d0s2
  HORCM_CMD for HORCMINST1 = /dev/rdsk/c1t0d0s2

  You can use the command device without a label in the format command.

- AIX®:

  HORCM_CMD for HORCMINST0 = /dev/rhdisk*X*
  HORCM_CMD for HORCMINST1 = /dev/rhdisk*X*

  where *X* = device number is created automatically by AIX®

- Tru64 UNIX®:

  HORCM_CMD for HORCMINST0 = /dev/rdisk/dsk*X*c
  HORCM_CMD for HORCMINST1 = /dev/rdisk/dsk*X*c
  where *X* = device number assigned by Tru64 UNIX®

- Linux®:

  HORCM_CMD for HORCMINST0 = /dev/sd*X*
  HORCM_CMD for HORCMINST1 = /dev/sd*X*
  where *X* = device number assigned by Linux®

- IRIX®:

  HORCM_CMD for HORCMINST0 = /dev/rdsk/dks*X*d*X*l*X*vol
  OR
  HORCM_CMD for HORCMINST0 = /dev/rdsk/node_wwn/lun*X*vol/c*X*p*X*

  HORCM_CMD for HORCMINST1 = /dev/rdsk/dksXdXlXvol
  OR
  HORCM_CMD for HORCMINST1 = /dev/rdsk/node_wwn/lun*X*vol/c*X*p*X*

  where *X* = device number assigned by IRIX®

- Windows NT®/2000:

  HORCM_CMD of HORCMINST0 = \\.\PhysicalDriveX
  OR
  HORCM_CMD of HORCMINST0 = \\.\Volume{guid} (Windows 2000 only)

  HORCM_CMD of HORCMINST1 = \\.\PhysicalDriveX
  OR
  HORCM_CMD of HORCMINST1 = \\.\Volume{guid} (Windows 2000 only)

  where *X* = device number assigned by Windows® NT/2000

  The PhysicalDrive number may change at every reboot. If the number changes, use Volume{guid} for which the same name is kept.

  If "\\.\Volume{guid}" is specified, CCI changes it to "\\.\PhysicalDrive?" to be corresponded. Volume{guid} is created when you make a partition by using the Windows' Disk Management. When the **inqraid $Volume -CLI -fv** or **raidscan -x findcmddev0.?** commands are used, you can find Volume{guid}.

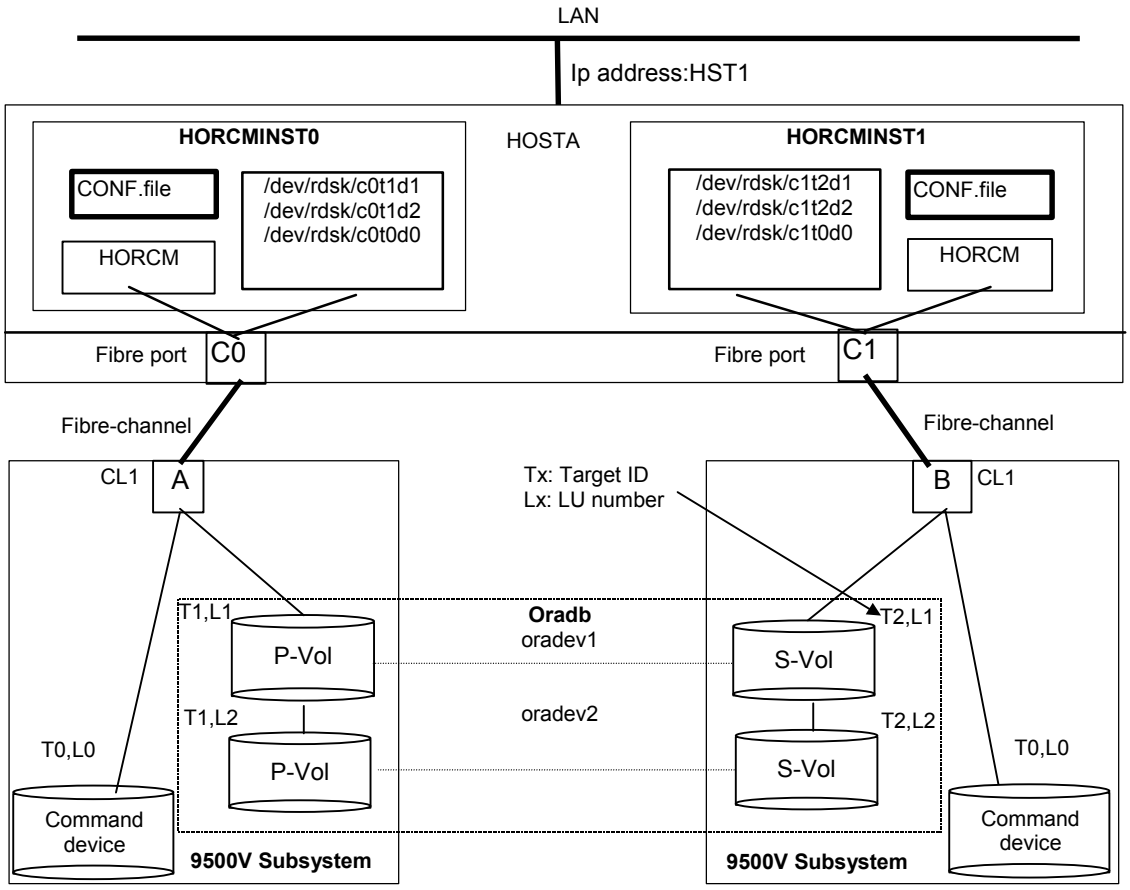

Configuration file for HORCMINST0 (horcm0.conf)

```
HORCM_MON
#ip_address  service   poll(10ms)   timeout(10ms)
HST1         horcm0    12000 [Note 1]      3000

HORCM_CMD
#dev_name
 /dev/xxx  [Note 2]

HORCM_DEV
#dev_group  dev_name  port#  TargetID  LU#  MU#
 Oradb      oradev1   CL1-A      1      1
 Oradb      oradev2   CL1-A      1      2

HORCM_INST
#dev_group      ip_address      service
 Oradb          HST1            horcm1
```

Configuration file for HORCMINST1 (horcm1.conf)

```
HORCM_MON
#ip_address  service   poll(10ms)   timeout(10ms)
HST1         horcm1    12000 [Note 1]      3000

HORCM_CMD
#dev_name
 /dev/yyy  [Note 2]

HORCM_DEV
#dev_group  dev_name  port#  TargetID  LU#  MU#
 Oradb      oradev1   CL1-B      2      1
 Oradb      oradev2   CL1-B      2      2

HORCM_INST
#dev_group      ip_address      service
 Oradb          HST1            horcm0
```

Note 1: To calculate the value for poll(10ms), see section 2.5.3.
Note 2: The command device is dedicated to CCI communications and cannot be used by any other applications (neither the user). Command devices must be set using the Resource Manager 9500V program. If setting two command devices, define two command devices in the HORCM_CMD section.
In the same line, you can also add a command device specified by different path, so that the host can use the same command device incase when one of the path cannot be used.

Figure 2.30 One Host and Two Instances Example Using Synchronous TrueCopy 9500V

Example of CCI commands with Instance-0 on HOSTA: Synchronous TrueCopy 9500V

- When the command execution environment is not set, set an instance number.
  For **# setenv HORCMINST 0**
  For Windows®: **set HORCMINST=0**

  Designate a group name (Oradb) and a local instance P- VOL a case.
  # paircreate -g Oradb -f never -vl

  This command creates pairs for all LUs assigned to group **Oradb** in the configuration definition file (two pairs for the configuration in Figure 2.30).

- Designate a volume name (oradev1) and a local instance P-VOL a case.
  # paircreate -g Oradb -d oradev1 -f never -vl

  This command creates pairs for all LUs designated as **oradev1** in the configuration definition file (CL1-A, T1, L1 and CL1-B, T2, L1 for the configuration in Figure 2.30).

- Designate a group name and display pair status.
  # pairdisplay -g Oradb

```
Group PairVol(L/R) (Port#,TID,LU),Seq#,LDEV#.P/S,Status,Fence, Seq#,P-LDEV# M
oradb  oradev1(L)    (CL1-A  ,  1,  1) 3005       18..P-VOL COPY NEVER , 3006      19  -
oradb  oradev1(R)    (CL1-B  ,  2,  1) 3006       19..S-VOL COPY NEVER , -----      18  -
oradb  oradev2(L)    (CL1-A  ,  1,  2) 3005       20..P-VOL COPY NEVER , 3006      21  -
oradb  oradev2(R)    (CL1-B  ,  2,  2) 3006       21..S-VOL COPY NEVER , -----      20  -
```

Example of CCI commands with Instance-1 on HOSTA: Synchronous TrueCopy 9500V

- When the command execution environment is not set, set an instance number.
  For **# setenv HORCMINST 1**
  For Windows®: **set HORCMINST=1**

- Designate a group name and a remote instance P-VOL a case.
  # paircreate -g Oradb -f never -vr

  This command creates pairs for all LU designated as **Oradb** in the configuration definition file (two pairs for the configuration in Figure 2.30).

- Designate a volume name (oradev1) and a remote instance P-VOL a case.
  # paircreate -g Oradb -d oradev1 -f never -vr

  This command creates pairs for all LUs designated as **oradev1** in the configuration definition file (CL1-A, T1, L1 and CL1-B, T2, L1 for the configuration in Figure 2.30).

- Designate a group name and display pair status.
  # pairdisplay -g Oradb

```
Group   PairVol(L/R) (Port#,TID,LU),Seq#,LDEV#.P/S,Status,Fence, Seq#,P-LDEV# M
oradb    oradev1(L)    (CL1-B  ,  2,  1)  3006       19..S-VOL COPY NEVER , -----      18  -
oradb    oradev1(R)    (CL1-A  ,  1,  1)  3005       18..P-VOL COPY NEVER , 3006     19  -
oradb    oradev2(L)    (CL1-B  ,  2,  2)  3006       21..S-VOL COPY NEVER , -----      20  -
oradb    oradev2(R)    (CL1-A  ,  1,  2)  3005       20..P-VOL COPY NEVER , 3006     21  -
```

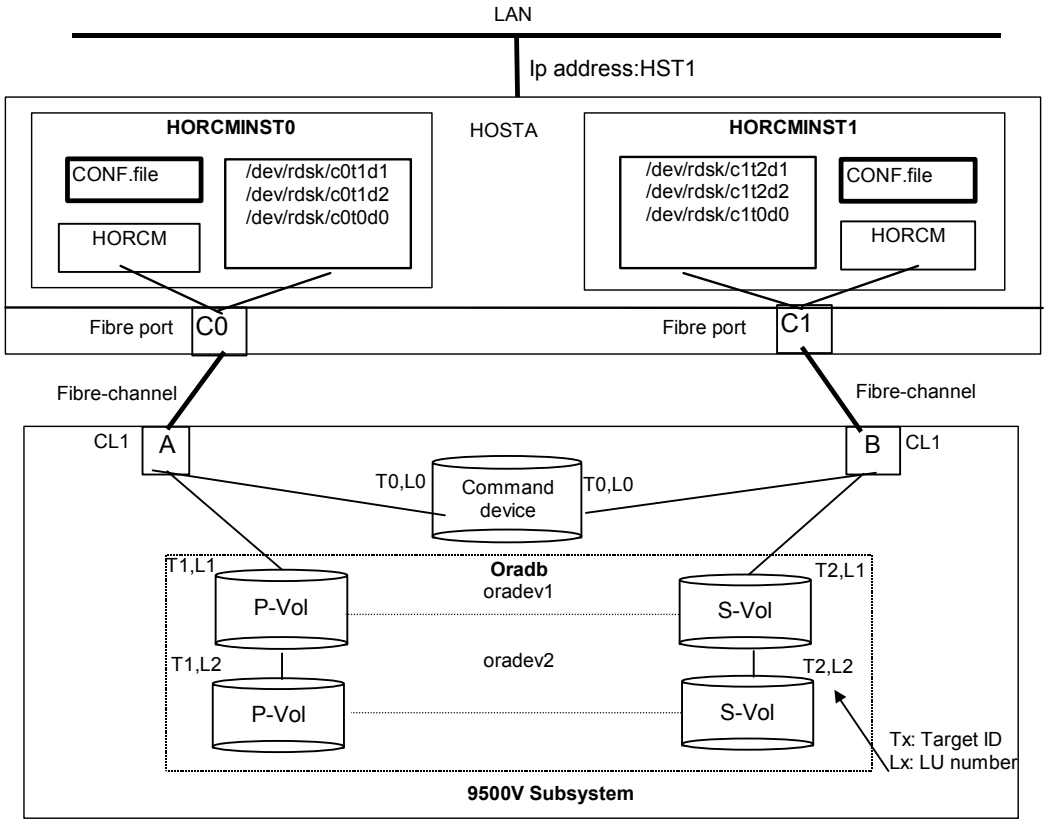

Configuration file for HORCMINST0 (horcm0.conf)

```
HORCM_MON
#ip_address  service   poll(10ms)   timeout(10ms)
HST1         horcm0    12000 [Note 1]     3000

HORCM_CMD
#dev_name
 /dev/xxx  [Note 2]

HORCM_DEV
#dev_group  dev_name  port#  TargetID  LU#  MU#
 Oradb      oradev1   CL1-A      1      1    0
 Oradb      oradev2   CL1-A      1      2    0

HORCM_INST
#dev_group     ip_address      service
 Oradb         HST1            horcm1
```

Configuration file for HORCMINST1 (horcm1.conf)

```
HORCM_MON
#ip_address  service   poll(10ms)   timeout(10ms)
HST1         horcm1    12000 [Note 1]     3000

HORCM_CMD
#dev_name
 /dev/xxx  [Note 2]

HORCM_DEV
#dev_group  dev_name  port#  TargetID  LU#  MU#
 Oradb      oradev1   CL1-B      2      1    0
 Oradb      oradev2   CL1-B      2      2    0

HORCM_INST
#dev_group     ip_address      service
 Oradb         HST1            horcm0
```

Note 1: To calculate the value for poll(10ms), see section 2.5.3.
Note 2: The command device is dedicated to CCI communications and cannot be used by any other application (neither the user). Command devices must be set using Resource Manager 9500V. When setting two command devices, define two command devices in the HORCM_CMD section.
In the same line, you can also add a command device specified by a different path so that the host can use the same command device when one of the paths cannot be used.

Figure 2.31 Two Hosts and Two Instances Using ShadowImage 9500V

Example of CCI commands with Instance-0 on HOSTA: ShadowImage 9500V

- When the command execution environment is not set, set an instance number.
  C shell: **# setenv HORCMINST 0**
  Windows®: **set HORCMINST=0**
- When the command execution environment is not set, set HORCC_MRCF to the environment variable.

  C shell: # setenv HORCC_MRCF 1
  Windows®: set HORCC_MRCF=1
- Designate a group name (Oradb) and a local instance P- VOL a case.
  # paircreate -g Oradb -vl

  This command creates pairs for all LUs assigned to group **Oradb** in the configuration definition file (two pairs for the configuration in Figure 2.31).
- Designate a volume name (oradev1) and a local instance P-VOL a case.
  # paircreate -g Oradb -d oradev1 -vl

  This command creates pairs for all LUs designated as **oradev1** in the configuration definition file (CL1-A, T1, L1 and CL1-B, T2, L1 for the configuration in Figure 2.31).
- Designate a group name and display pair status.
  # pairdisplay -g Oradb

```
Group PairVol(L/R)(Port#,TID,LU-M),Seq#,LDEV#..P/S, Status,  Seq#, P-LDEV# M
oradb oradev1(L)    (CL1-A, 1,1 - 0) 3005    18..   P-VOL COPY  3005      19    -
oradb oradev1(R)    (CL1-B, 2,1 - 0) 3005    19..   S-VOL COPY  -----     18    -
oradb oradev2(L)    (CL1-A, 1,2 - 0) 3005    20..   P-VOL COPY  3005      21    -
oradb oradev2(R)    (CL1-B, 2,2 - 0) 3005    21..   S-VOL COPY  -----     20    -
```

Example of CCI commands with Instance-1 on HOSTA: ShadowImage 9500V

- When the command execution environment is not set, set an instance number.
  C shell: **# setenv HORCMINST 1**
  Windows®: **set HORCMINST=1**
- When the command execution environment is not set, set HORCC_MRCF to the environment variable.

  C shell: # setenv HORCC_MRCF 1
  Windows®: set HORCC_MRCF=1
- Designate a group name and a remote instance P-VOL a case.
  # paircreate -g Oradb -vr

  This command creates pairs for all LU designated as **Oradb** in the configuration definition file (two pairs for the configuration in Figure 2.31).
- Designate a volume name (oradev1) and a remote instance P-VOL a case.
  # paircreate -g Oradb -d oradev1 -vr

  This command creates pairs for all LUs designated as **oradev1** in the configuration definition file (CL1-A, T1, L1 and CL1-B, T2, L1 for the configuration in Figure 2.31).

- Designate a group name and display pair status.
  # pairdisplay -g Oradb

```
Group   PairVol(L/R)(Port#,TID,LU-M),Seq#,LDEV#..P/S, Status, Seq#,P-LDEV# M
oradb   oradev1(L)   (CL1-B, 2,1 - 0)   3005   19..   S-VOL COPY  -----      18  -
oradb   oradev1(R)   (CL1-A, 1,1 - 0)   3005   18..   P-VOL COPY  3005     19  -
oradb   oradev2(L)   (CL1-B, 2,2 - 0)    3005   21..   S-VOL COPY  -----      20  -
oradb   oradev2(R)   (CL1-A, 1,2 - 0)    3005   20..   P-VOL COPY 3005      21  -
```

2.6.3 Two Command Devices

Setting two command devices enables you to use the alternate command device function.

The command device is defined by using the system raw device name (character-type device file name). The command device defined in the configuration definition file must be established to follow either every instance. If one command device is used between different instances on the same port, then the number of instances is up to 16 per command device. If this restriction is exceeded, use a different path for each instance. For example, the command devices for Figure 2.32 and Figure 2.33 would be:

- HP-UX®:

  HORCM_CMD for HORCMINST0 = /dev/rdsk/c0t0d0
  HORCM_CMD for HORCMINST0 = /dev/rdsk/c0t0d1
  HORCM_CMD for HORCMINST1 = /dev/rdsk/c1t0d0
  HORCM_CMD for HORCMINST1 = /dev/rdsk/c1t0d1

  You can use the command device without a label in the format command.

- Solaris®:

  HORCM_CMD for HORCMINST0 = /dev/rdsk/c0t0d0s2
  HORCM_CMD for HORCMINST0 = /dev/rdsk/c0t0d1s2
  HORCM_CMD for HORCMINST1 = /dev/rdsk/c1t0d0s2
  HORCM_CMD for HORCMINST1 = /dev/rdsk/c1t0d1s2

- AIX®:

  HORCM_CMD for HORCMINST0 = /dev/rhdisk*X*
  HORCM_CMD for HORCMINST0 = /dev/rhdisk*Y*
  HORCM_CMD for HORCMINST1 = /dev/rhdisk*X*
  HORCM_CMD for HORCMINST1 = /dev/rhdisk*Y*
  where *X* and *Y* = device number created automatically by AIX®

- Tru64 UNIX®:

  HORCM_CMD for HORCMINST0 = /dev/rhdisk/dsk*X*c
  HORCM_CMD for HORCMINST0 = /dev/rhdisk/dsk*Y*c
  HORCM_CMD for HORCMINST1 = /dev/rhdisk/dsk*X*c
  HORCM_CMD for HORCMINST1 = /dev/rhdisk/dsk*Y*c
  where *X* and *Y* = device number assigned by Tru64 UNIX®

- Linux®:

  HORCM_CMD for HORCMINST0 = /dev/sd*X*
  HORCM_CMD for HORCMINST0 = /dev/sd*Y*
  HORCM_CMD for HORCMINST1 = /dev/sd*X*
  HORCM_CMD for HORCMINST1 = /dev/sd*Y*
  where *X* and *Y* = device number assigned by Linux®

- IRIX®:

  HORCM_CMD for HORCMINST0 = /dev/rdsk/dks*X*d*X*l*X*vol
  OR
  HORCM_CMD for HORCMINST0 = /dev/rdsk/node_wwn/lun*X*vol/c*X*p*X*

  HORCM_CMD for HORCMINST0 = /dev/rdsk/dks*Y*d*Y*l*Y*vol
  OR
  HORCM_CMD for HORCMINST0 = /dev/rdsk/node_wwn/lun*Y*vol/c*Y*p*Y*

  HORCM_CMD for HORCMINST1 = /dev/rdsk/dks*X*d*X*l*X*vol
  OR
  HORCM_CMD for HORCMINST1 = /dev/rdsk/node_wwn/lun*X*vol/c*X*p*X*

  HORCM_CMD for HORCMINST1 = /dev/rdsk/dks*Y*d*Y*l*Y*vol
  OR
  HORCM_CMD for HORCMINST1 = /dev/rdsk/node_wwn/lun*Y*vol/c*Y*p*Y*
  where *X* and *Y* = device number assigned by IRIX®

- Windows NT®/2000:

  HORCM_CMD of HORCMINST0 = \\.\PhysicalDrive*X*

  OR
  HORCM_CMD of HORCMINST0 = \\.\Volume{guid} (Windows 2000 only)

  HORCM_CMD of HORCMINST0 = \\.\PhysicalDrive*Y*

  OR
  HORCM_CMD of HORCMINST0 = \\.\Volume{guid} (Windows 2000 only)

  HORCM_CMD of HORCMINST1 = \\.\PhysicalDrive*X*

  OR
  HORCM_CMD of HORCMINST1 = \\.\Volume{guid} (Windows 2000 only)

  HORCM_CMD of HORCMINST1 = \\.\PhysicalDrive*Y*

  OR
  HORCM_CMD of HORCMINST1 = \\.\Volume{guid} (Windows 2000 only)

  where *X* and *Y* = device number assigned by Windows NT®/2000

The PhysicalDrive number may change at every reboot. If the number changes, use Volume{guid} for which the same name is kept.

If "\\.\Volume{guid}" is specified, CCI changes it to "\\.\PhysicalDrive?" to be corresponded. Volume{guid} is created when you make a partition by using the Windows' Disk Management. When the **inqraid $Volume -CLI -fv** or **raidscan -x findcmddev0.?** commands are used, you can find Volume{guid}.

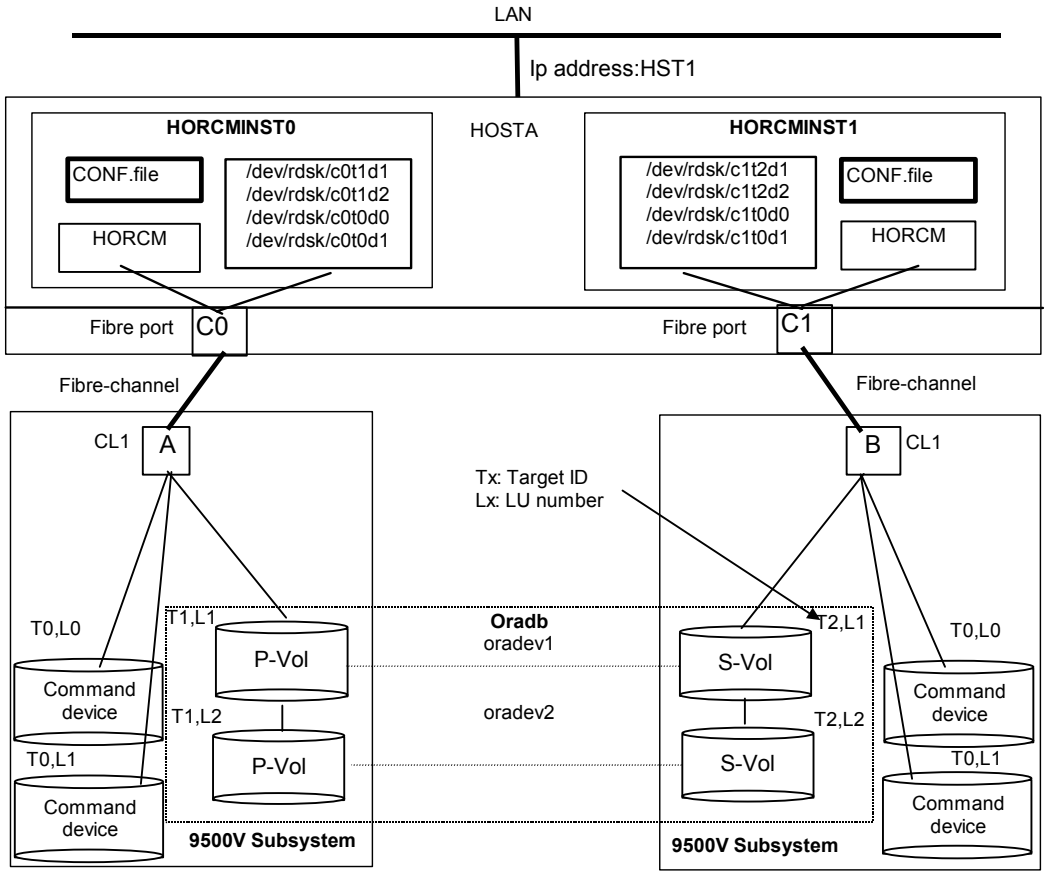

Configuration file for HORCMINST0 (horcm0.conf)

```
HORCM_MON
#ip_address  service   poll(10ms)   timeout(10ms)
HST1         horcm0    12000 [Note 1]    3000

HORCM_CMD
#dev_name
 /dev/xxx  /dev/yyy  [Note 2]

HORCM_DEV
#dev_group  dev_name  port#   TargetID  LU#  MU#
 Oradb      oradev1   CL1-A       1     1
 Oradb      oradev2   CL1-A       1     2

HORCM_INST
#dev_group      ip_address      service
 Oradb          HST1            horcm1
```

Configuration file for HORCMINST1 (horcm1.conf)

```
HORCM_MON
#ip_address  service   poll(10ms)   timeout(10ms)
HST1         horcm1    12000 [Note 1]    3000

HORCM_CMD
#dev_name
 /dev/zzz  /dev/uuu  [Note 2]

HORCM_DEV
#dev_group  dev_name  port#   TargetID  LU#  MU#
 Oradb      oradev1   CL1-B       2     1
 Oradb      oradev2   CL1-B       2     2

HORCM_INST
#dev_group      ip_address      service
 Oradb          HST1            horcm0
```

Note 1: To calculate the value for poll(10ms), see section 2.5.3.
Note 2: The command device is dedicated to CCI communications and cannot be used by any other applications (neither the user). Command devices must be set using the Resource Manager 9500V program. If setting two command devices, define two command devices in the HORCM_CMD section.
In the same line, you can also add a command device specified by different path, so that the host can use the same command device when one of the paths cannot be used.

Figure 2.32 Two Command Devices Example Using Synchronous TrueCopy 9500V

Example of CCI commands with Instance-0 on HOSTA: Synchronous TrueCopy 9500V

- When the command execution environment is not set, set an instance number.
  C shell: **# setenv HORCMINST 0**
  Windows®: **set HORCMINST=0**

- Designate a group name (Oradb) and a local instance P- VOL a case.
  # paircreate -g Oradb -f never -vl

  This command creates pairs for all LUs assigned to group **Oradb** in the configuration definition file (two pairs for the configuration in Figure 2.32).

- Designate a volume name (oradev1) and a local instance P-VOL a case.
  # paircreate -g Oradb -d oradev1 -f never -vl

  This command creates pairs for all LUs designated as **oradev1** in the configuration definition file (CL1-A, T1, L1 and CL1-B, T2, L1 for the configuration in Figure 2.32).

- Designate a group name and display pair status.
  # pairdisplay -g Oradb

```
Group  PairVol(L/R) (Port#,TID,LU),Seq#,LDEV#.P/S,Status,Fence, Seq#,P-LDEV# M
oradb   oradev1(L)   (CL1-A , 1,  1) 3005      18..P-VOL COPY NEVER , 3006   19  -
oradb   oradev1(R)   (CL1-B , 2,  1) 3006      19..S-VOL COPY NEVER , -----   18  -
oradb   oradev2(L)   (CL1-A , 1,  2) 3005      20..P-VOL COPY NEVER , 3006   21  -
oradb   oradev2(R)   (CL1-B , 2,  2) 3006      21..S-VOL COPY NEVER , -----    20  -
```

Example of CCI commands with Instance-1 on HOSTA: Synchronous TrueCopy 9500V

- When the command execution environment is not set, set an instance number.
  C shell: **# setenv HORCMINST 1**
  Windows®: **set HORCMINST=1**

- Designate a group name and a remote instance P-VOL a case.
  # paircreate -g Oradb -f never -vr

  This command creates pairs for all LU designated as **Oradb** in the configuration definition file (two pairs for the configuration in Figure 2.32).

- Designate a volume name (oradev1) and a remote instance P-VOL a case.
  # paircreate -g Oradb -d oradev1 -f never -vr

  This command creates pairs for all LUs designated as **oradev1** in the configuration definition file (CL1-A, T1, L1 and CL1-B, T2, L1 for the configuration in Figure 2.32).

- Designate a group name and display pair status.
  # pairdisplay -g Oradb

```
Group  PairVol(L/R) (Port#,TID,LU),Seq#,LDEV#.P/S,Status,Fence, Seq#,P-LDEV# M
oradb   oradev1(L)   (CL1-B , 2,  1) 3006      19..S-VOL COPY NEVER , -----    18  -
oradb   oradev1(R)   (CL1-A , 1,  1) 3005      18..P-VOL COPY NEVER , 3006   19  -
oradb   oradev2(L)   (CL1-B , 2,  2) 3006      21..S-VOL COPY NEVER , -----    20  -
oradb   oradev2(R)   (CL1-A , 1,  2) 3005      20..P-VOL COPY NEVER , 3006   21  -
```

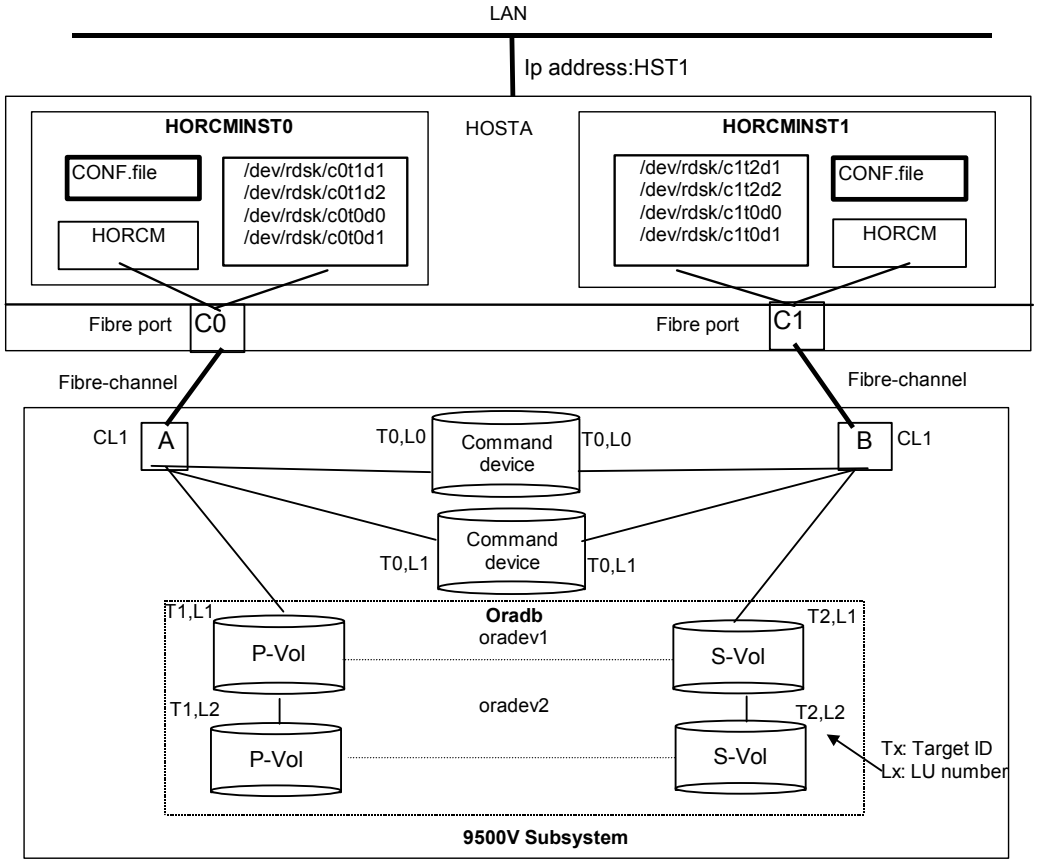

Configuration file for HORCMINST0 (horcm0.conf)

```
HORCM_MON
#ip_address  service   poll(10ms)   timeout(10ms)
HST1         horcm0    12000 [Note 1]    3000

HORCM_CMD
#dev_name
 /dev/xxx  /dev/yyy  [Note 2]

HORCM_DEV
#dev_group  dev_name  port#   TargetID  LU#  MU#
 Oradb      oradev1   CL1-A       1     1    0
 Oradb      oradev2   CL1-A       1     2    0

HORCM_INST
#dev_group      ip_address      service
 Oradb          HST1            horcm1
```

Configuration file for HORCMINST1 (horcm1.conf)

```
HORCM_MON
#ip_address  service   poll(10ms)   timeout(10ms)
HST1         horcm1    12000 [Note 1]    3000

HORCM_CMD
#dev_name
 /dev/xxx  /dev/yyy  [Note 2]

HORCM_DEV
#dev_group  dev_name  port#   TargetID  LU#  MU#
 Oradb      oradev1   CL1-B       2     1    0
 Oradb      oradev2   CL1-B       2     2    0

HORCM_INST
#dev_group      ip_address      service
 Oradb          HST1            horcm0
```

Note 1: To calculate the value for poll(10ms), see section 2.5.3.
Note 2: The command device is dedicated to CCI communications and cannot be used by any other application (neither the user). Command devices must be set using Resource Manager 9500V. When setting two command devices, define two command devices in the HORCM_CMD section.
In the same line, you can also add a command device specified by a different path so that the host can use the same command device when one of the paths cannot be used.

Figure 2.33 Two Command Devices Example Using ShadowImage 9500V

Example of CCI commands with Instance-0 on HOSTA: ShadowImage 9500V

- When the command execution environment is not set, set an instance number.
  C shell: **# setenv HORCMINST 0**
  Windows®: **set HORCMINST=0**
- When the command execution environment is not set, set HORCC_MRCF to the environment variable.

  C shell: **# setenv HORCC_MRCF 1**
  Windows®: **set HORCC_MRCF=1**
- Designate a group name (Oradb) and a local instance P- VOL a case.
  # paircreate -g Oradb -vl

  This command creates pairs for all LUs assigned to group **Oradb** in the configuration definition file (two pairs for the configuration in Figure 2.33).
- Designate a volume name (oradev1) and a local instance P-VOL a case.
  # paircreate -g Oradb -d oradev1 -vl

  This command creates pairs for all LUs designated as **oradev1** in the configuration definition file (CL1-A, T1, L1 and CL1-B, T2, L1 for the configuration in Figure 2.33).
- Designate a group name and display pair status.
  # pairdisplay -g Oradb

```
Group PairVol(L/R) (Port#,TID,LU-M), Seq#, LDEV#..P/S, Status,  Seq#, P-LDEV# M
oradb oradev1(L)    (CL1-A, 1,1 - 0)   3005   18..   P-VO  COPY  3005      19      -
oradb oradev1(R)    (CL1-B, 2,1 - 0)   3005   19..   S-VOL COPY  -----      18      -
oradb oradev2(L)    (CL1-A, 1,2 - 0)   3005   20..   P-VOL COPY  3005      21      -
oradb oradev2(R)    (CL1-B, 2,2 - 0)   3005   21..   S-VOL COPY  -----      20      -
```

Example of CCI commands with Instance-1 on HOSTA: ShadowImage 9500V

- When the command execution environment is not set, set an instance number.
  C shell: **# setenv HORCMINST 1**
  Windows®: **set HORCMINST=1**
- When the command execution environment is not set, set HORCC_MRCF to the environment variable.

  C shell: **# setenv HORCC_MRCF 1**
  Windows®: **set HORCC_MRCF=1**
- Designate a group name and a remote instance P-VOL a case.
  # paircreate -g Oradb -vr

  This command creates pairs for all LU designated as **Oradb** in the configuration definition file (two pairs for the configuration in Figure 2.33).
- Designate a volume name (oradev1) and a remote instance P-VOL a case.
  # paircreate -g Oradb -d oradev1 -vr

  This command creates pairs for all LUs designated as **oradev1** in the configuration definition file (CL1-A, T1, L1 and CL1-B, T2, L1 for the configuration in Figure 2.33).

- Designate a group name and display pair status.
  # pairdisplay -g Oradb

```
Group   PairVol(L/R) (Port#,TID,LU-M),Seq#,LDEV#..P/S,  Status, Seq#,P-LDEV# M
oradb   oradev1(L)   (CL1-B, 2,1 - 0)    3005   19..   S-VOL  COPY  -----       18  -
oradb   oradev1(R)   (CL1-A, 1,1 - 0)    3005   18..    P-VOL  COPY  3005       19  -
oradb   oradev2(L)   (CL1-B, 2,2 - 0)    3005    21..    S-VOL  COPY  -----       20  -
oradb   oradev2(R)   (CL1-A, 1,2 - 0)    3005   20..    P-VOL   COPY  3005       21  -
```

2.6.4 Synchronous TrueCopy 9500V/ShadowImage 9500V Configuration with Cascade Pairs

The command device is defined using the system raw device name (character-type device file name). The command device defined in the configuration definition file must be established in a way to be following either every instance. If one command device is used between different instances on the same port, then the number of instances is up to 16 per command device. If this restriction is exceeded, use a different path for each instance. For example, the command devices for would be as follows:

- HP-UX®:

  HORCM_CMD of HOSTA (/etc/horcm.conf) = /dev/rdsk/c0t0d1
  HORCM_CMD of HOSTB (/etc/horcm.conf) = /dev/rdsk/c1t0d1
  HORCM_CMD of HOSTB (/etc/horcm0.conf) = /dev/rdsk/c1t0d1

- Solaris®:

  HORCM_CMD of HOSTA (/etc/horcm.conf) = /dev/rdsk/c0t0d1s2
  HORCM_CMD of HOSTB (/etc/horcm.conf) = /dev/rdsk/c1t0d1s2
  HORCM_CMD of HOSTB (/etc/horcm0.conf) = /dev/rdsk/c1t0d1s2

  You can use the command device without a label in the format command.

- AIX®:

  HORCM_CMD of HOSTA (/etc/horcm.conf) = /dev/rhdisk*X*
  HORCM_CMD of HOSTB (/etc/horcm.conf) = /dev/rhdisk*X*
  HORCM_CMD of HOSTB (/etc/horcm0.conf) = /dev/rhdisk*X*

  where *X* = device number assigned by AIX®

- Tru64 UNIX®:

  HORCM_CMD of HOSTA (/etc/horcm.conf) = /dev/rhdisk/dsk*X*c
  HORCM_CMD of HOSTB (/etc/horcm.conf) = /dev/rhdisk/dsk*X*c
  HORCM_CMD of HOSTB (/etc/horcm0.conf) = /dev/rhdisk/dsk*X*c

  where *X* = device number assigned by Tru64 UNIX®

- Linux®:

  HORCM_CMD of HOSTA (/etc/horcm.conf) = /dev/sd*X*
  HORCM_CMD of HOSTB (/etc/horcm.conf) = /dev/sd*X*
  HORCM_CMD of HOSTB (/etc/horcm0.conf) = /dev/sd*X*

  where *X* = device number assigned by Linux®

- IRIX®:

  HORCM_CMD of HOSTA (/etc/horcm.conf) = /dev/rdsk/dks*X*d*X*l*X*vol
  OR
  HORCM_CMD of HOSTA (/etc/horcm.conf) = /dev/rdsk/node_wwn/lun*X*vol/c*X*p*X*

  HORCM_CMD of HOSTB (/etc/horcm.conf) = /dev/rdsk/dks*X*d*X*l*X*vol
  OR
  HORCM_CMD of HOSTB (/etc/horcm.conf) = /dev/rdsk/node_wwn/lun*X*vol/c*X*p*X*

  HORCM_CMD of HOSTB (/etc/horcm0.conf) = /dev/rdsk/dks*X*d*Y*l*X*vol
  OR
  HORCM_CMD of HOSTB (/etc/horcm0.conf) = /dev/rdsk/node_wwn/lun*X*vol/c*Y*p*X*

  where *X* = device number assigned by IRIX®

- Windows® NT/2000:

  HORCM_CMD of HOSTA (/etc/horcm.conf) = \\.\PhysicalDrive*X*
  OR
  HORCM_CMD of HOSTA (/etc/horcm.conf) = \\.\Volume{guid} (Windows 2000 only)

  HORCM_CMD of HOSTB (/etc/horcm.conf) = \\.\PhysicalDrive*X*
  OR
  HORCM_CMD of HOSTB (/etc/horcm.conf) = \\.\Volume{guid} (Windows 2000 only)

  HORCM_CMD of HOSTB (/etc/horcm0.conf) = \\.\PhysicalDrive*X*
  OR
  HORCM_CMD of HOSTB (/etc/horcm0.conf) = \\.\Volume{guid} (Windows 2000 only)

  where *X* = device number assigned by Windows® NT/2000

  The PhysicalDrive number may change at every reboot. If the number changes, use Volume{guid} for which the same name is kept.

  If "\\.\Volume{guid}" is specified, CCI changes it to "\\.\PhysicalDrive?" to be corresponded. Volume{guid} is created when you make a partition by using the Windows' Disk Management. When the **inqraid $Volume -CLI -fv** or **raidscan -x findcmddev0.?** commands are used, you can find Volume{guid}.

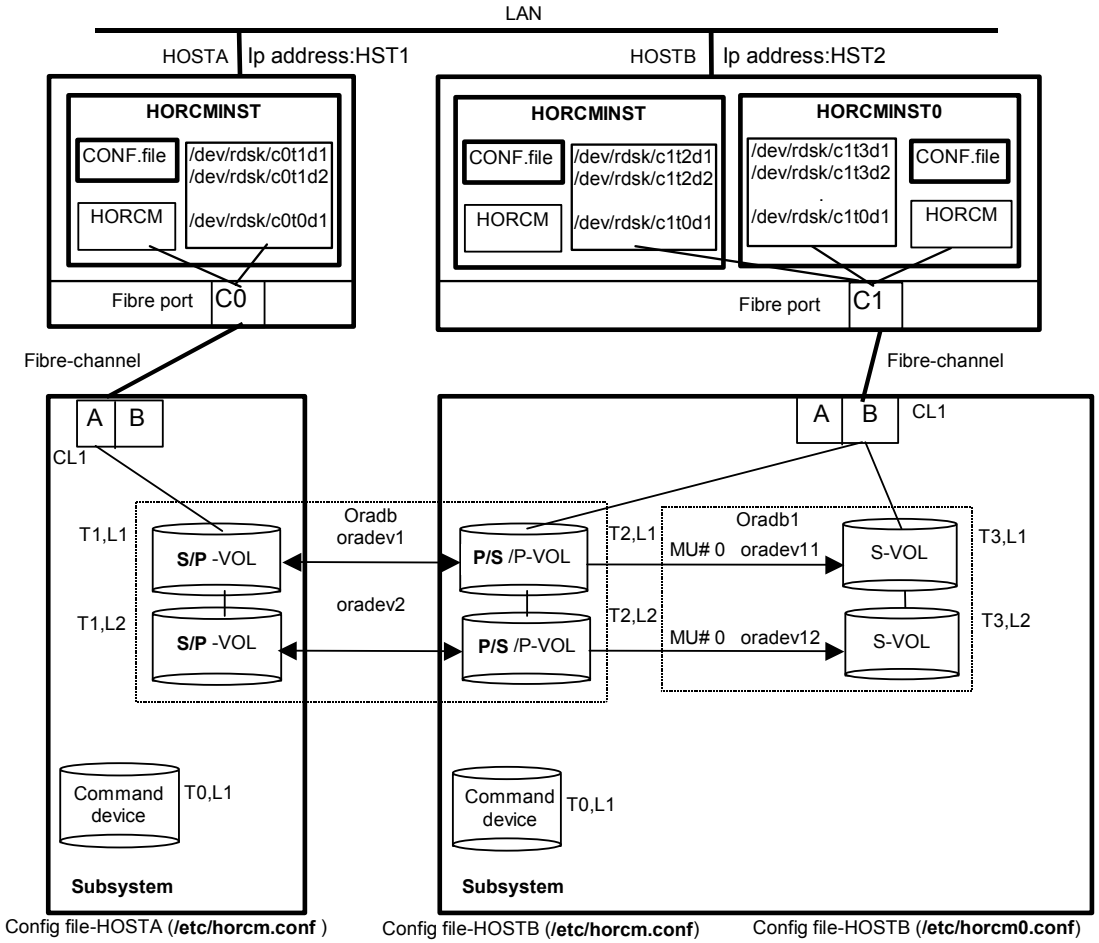

Config file-HOSTA (**/etc/horcm.conf**)

```
HORCM_MON
#ip_address service poll(10ms) timeout(10ms)
HST1        horcm    12000  [Note 1]   3000
HORCM_CMD
#dev_name
 /dev/xxx  [Note 2]
HORCM_DEV
#dev_groupdev_name port#  TID LU# MU#
Oradb       oradev1   CL1-A  1     1
Oradb       oradev2   CL1-A  1     2

HORCM_INST
#dev_group      ip_address      service
 Oradb          HST2            horcm
 Oradb          HST2            horcm0
```

Config file-HOSTB (**/etc/horcm.conf**)

```
HORCM_MON
#ip_address service poll(10ms) timeout(10ms)
HST2        horcm    12000  [Note 1]   3000
HORCM_CMD
#dev_name
 /dev/xxx  [Note 2]
HORCM_DEV
#dev_groupdev_name port#  TID LU# MU#
Oradb       oradev1   CL1-B  2     1
Oradb       oradev2   CL1-B  2     2
Oradb1      oradev11  CL1-B  2     1     0
Oradb1      oradev12  CL1-B  2     2     0

HORCM_INST
#dev_group      ip_address      service
 Oradb          HST1            horcm
 Oradb1         HST2            horcm0
```

Config file-HOSTB (**/etc/horcm0.conf**)

```
HORCM_MON
#ip_address service poll(10ms) timeout(10ms)
HST2        horcm0   12000  [Note 1]   3000
HORCM_CMD
#dev_name
 /dev/xxx  [Note 2]
HORCM_DEV
#dev_groupdev_name port#  TID LU# MU#
Oradb       oradev1   CL1-B  2     1
Oradb       oradev2   CL1-B  2     2
Oradb1      oradev11  CL1-B  3     1     0
Oradb1      oradev12  CL1-B  3     2     0

HORCM_INST
#dev_group      ip_address      service
 Oradb          HST1            horcm
 Oradb1         HST2            horcm
```

Shaded portions: If HORCMINST0 needs to operate TrueCopy's paired volume, then describe **oradb**.
Note 1: To calculate the value for poll(10ms), see section 2.5.3.
Note 2: The command device is dedicated to CCI communications and cannot be used by any other applications (neither the user). Command devices must be set using the Resource Manager 9500V program. If you are setting two command devices, define two command devices in the HORCM_CMD section. In the same line, you can also add a command device specified by different path, so that the host can use the same command device when one of the paths cannot be used.

Figure 2.34 Synchronous TrueCopy 9500V/ShadowImage 9500V Configuration Example with Cascade Pairs

Example of CCI commands with HOSTA and HOSTB:

- Designate a group name (Oradb) on Synchronous TrueCopy 9500V environment of HOSTA.
  # paircreate -g Oradb -f never -vl
- Designate a group name (Oradb1) on MRCF-Lite environment of HOSTB. When the command execution environment is not set, set HORCC_MRCF.

  C shell: **# setenv HORCC_MRCF 1**
  Windows®: **set HORCC_MRCF=1**
  # paircreate -g Oradb1 -vl

  These commands create pairs for all LUs assigned to groups **Oradb** and **Oradb1** in the configuration definition file (four pairs for the configuration in Figure 2.34).
- Designate a group name and display pair status on HOSTA.
  # pairdisplay -g oradb -m all

```
Group   PairVol(L/R) (Port#,TID,LU-M),Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M
oradb    oradev1(L)   (CL1-A , 1,  1-0) 3052   26..SMPL  ----,    -----    ----    -
oradb    oradev1(L)   (CL1-A , 1,  1)    3052   26..P-VOL COPY,  3053    28     -
oradb1  oradev11(R) (CL1-B , 2,  1-0)  3053   28..P-VOL COPY, 3053    30     -
oradb    oradev1(R)   (CL1-B , 2,  1)    3053   28..S-VOL COPY,  -----    26     -
oradb    oradev2(L)   (CL1-A , 1,  2-0) 3052   27..SMPL  ----,    -----   ----    -
oradb    oradev2(L)   (CL1-A , 1,  2)    3052   27..P-VOL COPY,  3053    29     -
oradb1  oradev12(R) (CL1-B , 2,  2-0)  3053   29..P-VOL COPY, 3053    31     -
oradb    oradev2(R)   (CL1-B , 2,  2)    3053   29..S-VOL COPY,  -----    27     -
```

Example of CCI commands with HOSTB:

- Designate a group name (oradb) on Synchronous TrueCopy 9500V environment of HOSTB.
  # paircreate -g Oradb -f never -vr
- Designate a group name (Oradb1) on MRCF-Lite environment of HOSTB. When the command execution environment is not set, set HORCC_MRCF.
  C shell: **# setenv HORCC_MRCF 1**
  Windows®: **set HORCC_MRCF=1**
  # paircreate -g Oradb1 -vl

  This command creates pairs for all LUs assigned to group **Oradb1** in the configuration definition file (four pairs for the configuration in Figure 2.34).
- Designate a group name and display pair status on Synchronous TrueCopy 9500V environment of HOSTB.
  # pairdisplay -g oradb -m all

```
Group   PairVol(L/R) (Port#,TID,LU-M),Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M
oradb1  oradev11(L) (CL1-B , 2, 1-0)   3053   28..P-VOL  PAIR, 3053     30     -
oradb    oradev1(L)  (CL1-B , 2,  1)     3053   28..S-VOL   PAIR, -----     26     -
oradb    oradev1(R)  (CL1-A , 1,  1-0)  3052   26..SMPL    ----, -----     ----   -
oradb    oradev1(R)  (CL1-A , 1,  1)     3052   26..P-VOL  PAIR, 3053     28     -
oradb1  oradev12(L) (CL1-B , 2,  2-0)  3053   29..P-VOL  PAIR, 3053     31     -
oradb    oradev2(L)  (CL1-B , 2,  2)      3053  29..S-VOL  PAIR,  -----    27     -
oradb    oradev2(R)  (CL1-A , 1,  2-0)   3052  27..SMPL   ----,  -----    ----   -
oradb    oradev2(R)  (CL1-A , 1,  2)     3052   27..P-VOL  PAIR, 3053     29     -
```

- Designate a group name and display pair status on the ShadowImage 9500V environment of HOSTB.
  # pairdisplay -g oradb1 -m all

```
Group   PairVol(L/R) (Port#,TID,LU-M),Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M
oradb1  oradev11(L) (CL1-B , 2,  1-0)   3053   28..P-VOL PAIR,   3053     30    -
oradb   oradev1(L)  (CL1-B , 2,  1)     3053   28..S-VOL PAIR,   -----     26    -
oradb1  oradev11(R) (CL1-B , 3,  1-0)   3053   30..S-VOL PAIR,   -----     28    -
oradb1  oradev12(L) (CL1-B , 2,  2-0)   3053   29..P-VOL PAIR,   3053     31    -
oradb   oradev2(L)  (CL1-B , 2,  2)     3053   29..S-VOL PAIR,   -----     27    -
oradb1  oradev12(R) (CL1-B , 3,  2-0)   3053   31..S-VOL PAIR,   -----     29    -
```

- Designate a group name and display pair status on the ShadowImage 9500V environment of HOSTB (HORCMINST0).
  # pairdisplay -g oradb1 -m all

```
Group   PairVol(L/R) (Port#,TID,LU-M),Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M
oradb1  oradev11(L) (CL1-B , 3,  1-0)   3053   30..S-VOL PAIR,   -----     28    -
oradb1  oradev11(R) (CL1-B , 2,  1-0)   3053   28..P-VOL PAIR,   3053     30    -
oradb   oradev1(R)  (CL1-B , 2,  1)     3053   28..S-VOL PAIR,   -----     26    -
oradb1  oradev12(L) (CL1-B , 3,  2-0)   3053   31..S-VOL PAIR,   -----     29    -
oradb1  oradev12(R) (CL1-B , 2,  2-0)   3053   29..P-VOL PAIR,   3053     31    -
oradb   oradev2(R)  (CL1-B , 2,  2)     3053   29..S-VOL PAIR,   -----     27    -
```

2.7 Error Monitoring and Configuration Confirmation

HORCM supports error monitoring and configuration confirmation commands for linkage with the system operation management of the UNIX®/PC server.

2.7.1 Paired Volume Error Monitoring

The HORC Manager (HORCM) monitors all volumes defined in the configuration definition file at a certain interval regardless of the Synchronous TrueCopy 9500V/ShadowImage 9500V commands.

- **Objects and scope of monitoring:** The HORCM operates as a daemon process on the host server and monitors all the paired volumes defined in the configuration definition file, not the volume groups. The HORC Manager's monitoring applies to the primary volumes only (since the primary volumes control the status). The HORC Manager monitors the changes in the pair status of these volumes. Only when the PAIR status changes to the PSUS status and that change is caused by an error (such as P-VOL error or S-VOL's SUS), does the HORC Manager regard the change as an error.
- **Monitoring time and interval:** This command always issues I/O instructions to the 9500V in order to obtain information for monitoring. It is possible to specify the monitoring interval in the configuration definition file to adjust the daemon load.
- **Error notification by HORCM:** Since the operation management of the UNIX® server checks Syslog to find system errors in many cases, Synchronous TrueCopy 9500V/ShadowImage 9500V error messages are output to Syslog for linkage with the system operation management.
- **Error notification command:** The CCI supports the error notification function using commands in order to allow the UNIX® server client to monitor errors. This command is connected to the HORCM (daemon) to obtain the transition of the pairing status and report it. When an error is detected, this command outputs an error message. This command waits until an error occurs or reports that "No" error occurs if it finds no errors in pairing status transition queue of the HORCM's pairing monitor. Operations can be specified using certain options. If the command finds the status transition data in the status transition queue, it displays the data of all volumes. Specifying the option of this command can erase data in the HORCM's status transition queue.

2.7.2 Pair Status Display and Configuration Confirmation

The configuration definition file combines physical volumes in the 9500V used independently by the servers. Therefore, be certain that the server volumes are combined as intended by the server system administrator.

The pairdisplay command displays the pairing status to enable you to verify the completion of pair creation or pair resynchronization (see Figure 2.35). This command is also used to confirm the configuration of the paired volume connection path (physical link of paired volumes among the servers). For further information on the pairdisplay command, see section 4.9.

```
                          --Link information of ①--  ----Link information of ②---
Group PairVol(L/R) (Port#,TID,LU),Seq#,LDEV#.P/S,Status,Fence, Seq#,P-LDEV# M
G1    Oradb1(L)    (P1 ,  T1,  L1),Seq#,  1..P-VOL, Pair NEVER , Seq#,   2  -
G1    Oradb1(R)    (P2 ,  T2,  L2),Seq#,  2..S-VOL, Pair NEVER , Seq#,   1  -
```

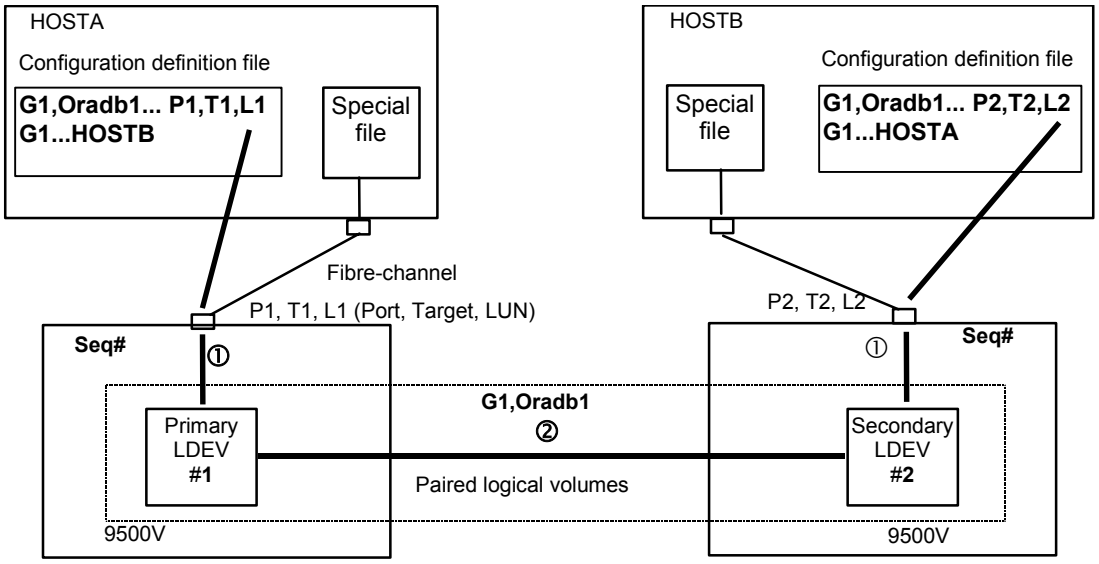

Figure 2.35 Pair Configuration Confirmation (Pairdisplay)

The raidscan command displays the Fibre port, target ID, LDEVs mapped to LUNs, and status of those LDEVs, regardless of the configuration definition file. When a port number is specified, this command displays information about all target IDs and LUNs of that port. For further information on the raidscan command, see section 4.12.1.

2.8 Recovery Procedures for HA Configurations

After configuring and starting Synchronous TrueCopy 9500V operations, the system administrator should conduct operational tests for possible failures in the system. In normal operation, service personnel obtain information for identifying the failure cause on the HITRACK and SNMP Agent or Web. However, a motive for the action above should be given by the Synchronous TrueCopy 9500V operation command.

Figure 2.36 shows the system failover and recovery procedure. Figure 2.37 shows the regression and Synchronous TrueCopy 9500V recovery procedure.

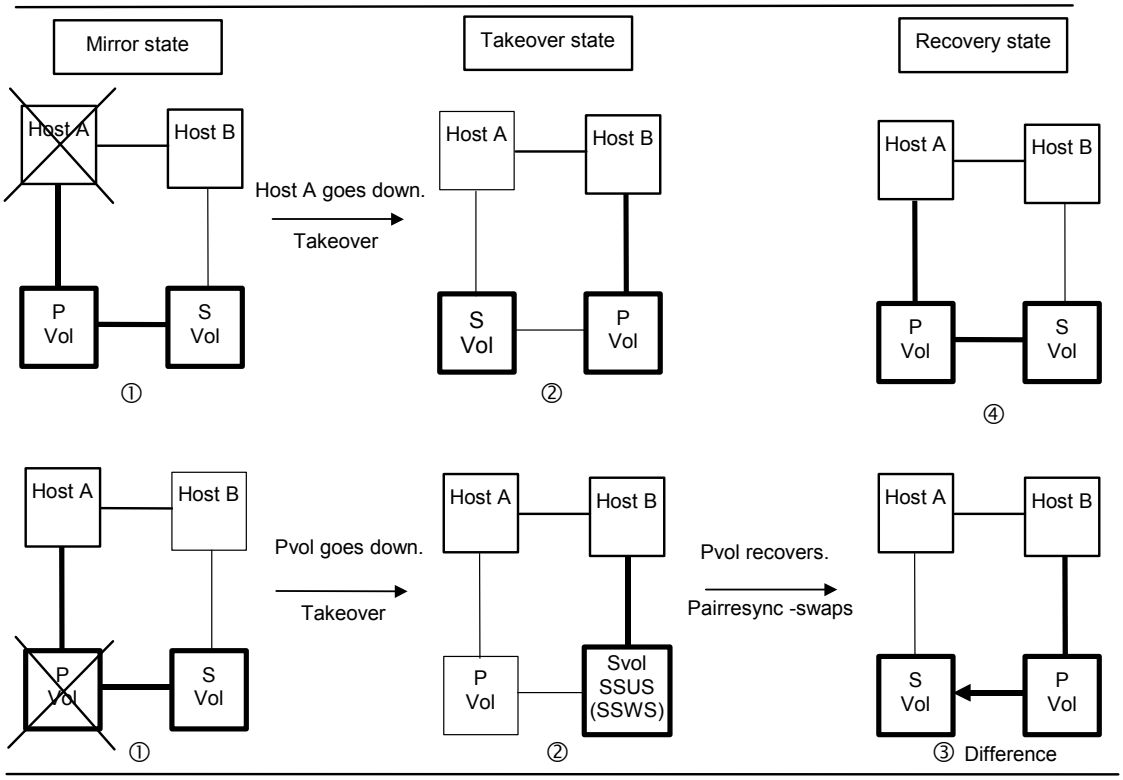

① A failure occurs in the host A server (1-top) or in the Pvol (1-bottom).

② Host B detects the failure of host A or the Pvol and issues a takeover command to make the SVOL usable. Host B takes over processing from host A. In the case of host A failure (1-top), the Swap-takeover command will be executed. In the case of Pvol failure (1-bottom), the SVOL-SSUS-takeover command will be executed.

③ While host B continues processing, PVOL and SVOL are swapped (pairresync -swaps), and the delta data (BITMAP) updated by host B is fed back to host A.

④ After host A or the Pvol has recovered, host A can take over processing from host B by executing the swap-takeover (horctakeover) command.

Figure 2.36 System Failover and Recovery

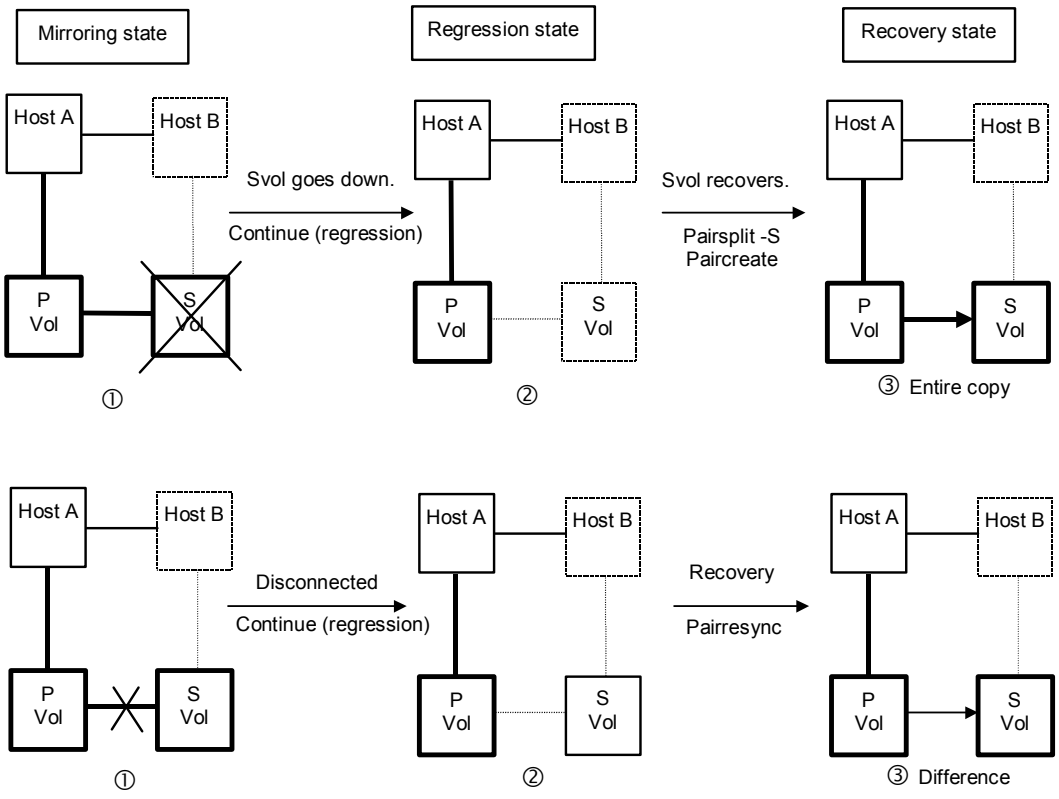

① The PVOL detects a failure in the SVOL and causes suspension of the duplicated writing. (The fence level determines whether host A continues processing or host B takes over the processing from host A.)

② The PVOL changes the paired volume status to PSUE and keeps track of the difference data. The HORCM detects the status change and outputs a message to syslog. If the client of host A has initiated the monitoring command, the message concerned is displayed on the screen of the client.

③ The SVOL recovers from the failure. The host A issues the pairsplit -S, paircreate -vl, or pairresync command to update the PVOL data by copying entire data or copying differential data only. The updated data is fed back to the SVOL.

Figure 2.37 Degeneracy and Recovery in Case of System Error

Chapter 3 Preparing for CCI Operations

3.1 System Requirements

CCI operations involve the CCI software on the UNIX®/PC server host and the Thunder 9500™ V subsystem(s) containing the command device(s) set using the Resource Manager 9500, and the Synchronous TrueCopy 9500V and/or ShadowImage 9500V primary and secondary volumes. From the Resource Manager 9500V program, specify a local and a remote port (path) of the subsystem.

The system requirements for CCI are:

- **CCI software product.** The CCI software is supplied on CD-ROM. The CCI software files take up 2.5 MB of space. The log files can take up to 3 MB of space.
- **Host platform.** CCI is supported on several UNIX®-based and PC server platforms, including Solaris®, HP-UX®, AIX®, Linux®, Tru64 UNIX®, IRIX®, Windows® 2000, and Windows NT® systems.
  - Root/administrator access to the host is required to perform CCI operations.
  - Static memory capacity: minimum = 300 kB, maximum = 500 kB
    Dynamic memory capacity (depends on the setting in HORCM_CONF):
    maximum = 500 kB per unit ID.
  - Host servers which are combined when paired logical volumes are defined should run on the operating system of the same architecture. If not, one host may not be able to recognize a paired volume of another host, even though HORCM can run properly.
- **Thunder 9500™ V subsystem.** The Thunder 9500™ V subsystems support CCI operations.
  - The 9500V command device must be defined and accessed as a raw device (no file system, no mount operation).
- **Resource Manager 9500V.** The Resource Manager 9500V is used to define the LUs, path using Synchronous TrueCopy 9500V, and the CCI command device.

Note: Windows® 2000 dynamic disks are not supported.

3.2 Hardware Installation

Installation of the hardware required for CCI is performed by the user and the Hitachi Data Systems representative. To install the hardware required for CCI operations:

User:

a) Identify the Synchronous TrueCopy 9500V and/or ShadowImage 9500V primary and secondary volumes, so that the CCI hardware and software components can be installed and configured properly.

b) Verify that the UNIX®/PC server hardware and software are properly installed and configured (see section 3.1 for system requirements).

Hitachi Data Systems representative:

a) Connect the Thunder 9500™ V subsystem(s) to the UNIX®/PC server host(s).

b) Install and enable the Synchronous TrueCopy 9500V and/or ShadowImage 9500V features on the Thunder 9500™ V subsystem(s).

3.3 Software Installation

Installation of the CCI software on the host server(s) is performed by the user, with assistance as needed from the Hitachi Data Systems representative.

3.3.1 Software Installation for UNIX® Systems

If you are installing CCI from CD-ROM, please use the **RMinstsh** and **RMuninst** scripts on the CD-ROM to automatically install and uninstall the CCI software. For other media, please use the following instructions.

Note: The following instructions refer to UNIX® commands, which may be different on your platform. Please consult your operating system documentation (e.g., UNIX® man pages) for platform-specific command information.

3.3.1.1 New Installation into Root Directory

1. Insert the installation CD into the CD-ROM drive.
2. Mount the CD manually if it is not automounted by the OS. Create mount point if one is not available

   Example (Unix):

   #mkdir /cdrom

   #mount /dev/dsk/c0t1d0s2 /cdrom
3. Move to the root directory.

   Example: cd /
4. Copy all files from the installation medium using the **cpio** command:

   # cpio -idmu < /cdrom/XX_XX_XX/Solaris/rmhorc
5. Execute the HORCM installation command: **# /HORCM/horcminstall.sh**
6. Verify installation of the proper version using the raidqry command:

   # raidqry -h
   Model: RAID-Manager/HP-UX
   Ver&Rev: XX-XX-XX/XX XX = product version number
   Usage: raidqry [options]

3.3.1.2 New Installation into Non-Root Directory

1. Insert the installation medium (e.g., CD-ROM) into the proper I/O device.
2. Move to the desired directory for CCI. The specified directory must be mounted by a partition of except root disk or an external disk.

   # cd /Specified Directory
3. Copy all files from the installation medium using the **cpio** command:

   # cpio -idmu < /cdrom/XX_XX_XX/Solaris/rmho
4. Make a symbolic link for /HORCM:

   # ln -s /*Specified Directory*/HORCM /HORCM
5. Execute the HORCM installation command:

   # /HORCM/horcminstall.sh
6. Verify installation of the proper version using the **raidqry** command:

   # raidqry -h
   Model: RAID-Manager/HP-UX
   Ver&Rev: XX-XX-XX/XX XX = product version number
   Usage: raidqry [options]

3.3.2 Changing the CCI User for UNIX® Systems

The CCI software is initially configured to allow only the **root** user (system administrator) to execute CCI commands. If desired (e.g., CCI administrator does not have **root** access), the system administrator can change the CCI user from **root** to another user name.

To change the CCI user:

1. Change the owner of the following CCI files from the **root** user to the desired user name:
   /HORCM/etc/horcmgr
   All CCI commands in the /HORCM/usr/bin directory
   All CCI log directories in the /HORCM/log directories
2. Change the owner of the raw device file of the HORCM_CMD command device in the configuration definition file from the **root** user to the desired user name.
3. ***Optional***: Establishing the HORCM (/etc/horcmgr) start environment. If users have designation of the full environment variables (HORCM_LOG HORCM_LOGS), they start horcmstart.sh command without an argument. In this case, the HORCM_LOG and HORCM_LOGS directories must be owned by the CCI administrator. The environment variable (HORCMINST, HORCM_CONF) is established as necessary.
4. ***Optional***: Establishing the command execution environment. If users have designation of the environment variables (HORCC_LOG), the HORCC_LOG directory must be owned by the CCI administrator. The environment variable (HORCMINST) is established as necessary.

3.3.3 Software Installation for Windows® NT/2000 Systems

Install CCI on all servers involved in CCI operations. If the network (TCP/IP) is not established, install a network of the Windows NT®/2000 attachment, and add TCP/IP protocol.

To install the CCI software on a Windows NT®/2000 system:

1. If a previous version of CCI is already installed, uninstall (remove) it as follows:
   a) Confirm that HORCM is not running. If it is running, shut it down:

      One CCI instance: **D:\HORCM\etc> horcmshutdown**

      Two CCI instances: **D:\HORCM\etc> horcmshutdown 0 1**

   b) If Synchronous TrueCopy 9500V/ShadowImage 9500V commands are running in the interactive mode, terminate the interactive mode and exit these commands using the **-q** option.

   c) Remove the previous version of CCI using the **Add/Remove Programs** control panel.
2. Insert the installation medium (e.g., CD-ROM) into the proper I/O device.
3. Run **Setup.exe** and follow the instructions on screen to complete the installation.
4. Verify installation of the proper version using the **raidqry** command:

   D:\HORCM\etc> raidqry -h
   Model: RAID-Manager/HP-UX
   Ver&Rev: XX-XX-XX/XX XX = product version number
   Usage: raidqry [options]

3.4 Creating/Editing the Configuration File

The configuration definition file is a text file that is created and/or edited using any standard text editor (e.g., UNIX® vi editor, Windows® Notepad). A sample configuration definition file, HORCM_CONF (/HORCM/etc/horcm.conf), is included with the CCI software. This file should be used as the basis for creating your configuration definition file(s). The system administrator should copy the sample file, set the necessary parameters in the copied file, and place the copied file in the proper directory.

See section 2.5.3 for a detailed description of editing and creating the configuration definition file(s), and see section 2.6 for sample CCI configurations.

Caution: Do not edit the configuration definition file while HORCM is running. Shut down HORCM, edit the configuration file as needed, and then restart HORCM.

Table 3.1 lists the parameters defined in the configuration file and specifies the default value, type, and limit for each parameter.

Table 3.1 Configuration (HORCM_CONF) Parameters

Parameter	Default value	Type	Limit
ip_address	None	Character string	63 characters
service	None	Character string or numeric value*1	15 characters
Poll (10 ms)	1000*2	Numeric value*1	None*3
timeout (10 ms)	3000	Numeric value*1	None
dev_name	None	Character string	31 characters Recommended value = 8 or less.
dev_group	None	Character string	31 characters Recommended value = 8 or less.
port #	None	Character string	31 characters
Target ID	None	Numeric value*1	7 characters
LU#	None	Numeric value*1	7 characters
MU#	0	Numeric value*1	7 characters
HORCM CMDdev name	None	Character string	63 characters

*1 ***Note:*** Use decimal notation for numeric values (not hexadecimal).
*2 ***Note:*** For Synchronous TrueCopy/ShadowImage 9500V operations, you must change the default value using the equation described in section 2.5.3. Setting the value incorrectly may cause a conflict in the internal process, which suspends the process temporary and stops the internal process of the subsystem.
*3 ***Note:*** For details on calculating the value and the equation for poll(10ms), see section 2.5.3.

3.5 CCI Startup

After you have installed the CCI software (see section 3.3) and set the configuration definition file(s) (see section 3.4), you can begin using the CCI software (HORCM) to perform Synchronous TrueCopy 9500V and/or ShadowImage 9500V operations on the attached Thunder 9500™ V subsystems.

3.5.1 Startup for UNIX® Systems

3.5.1.1 One Instance

To start up one instance of CCI on a UNIX® system:

1. Modify **/etc/services** to register the port name/number (service) of the configuration definition file. Make the port name/number the same on all servers.
   horcm xxxxx/udp xxxxx = the port name/number of horcm.conf
2. If you want HORCM to start automatically each time the system starts up, add **/etc/horcmstart.sh** to the system automatic start-up file (e.g., **/sbin/rc**).
3. Execute the **horcmstart.sh** script manually to start the CCI instance:
   # horcmstart.sh
4. Set the log directory (**HORCC_LOG**) in the command execution environment as needed.
5. If you want to perform Synchronous TrueCopy 9500V operations, do not set the **HORCC_MRCF** environment variable. If you want to perform ShadowImage 9500V operations, set the **HORCC_MRCF** environment variable for the HORCM execution environment.

   For B shell:

   # HORCC_MRCF=1
   # export HORCC_MRCF

   C shell:

   # setenv HORCC_MRCF 1
6. Execute the pairdisplay command to verify the configuration.
   # pairdisplay -g xxxx xxxx = group-name

3.5.1.2 Two Instances

To start up two instances of CCI on a UNIX® system:

1. Make two copies of the sample configuration definition file.

   # cp /etc/horcm.conf /etc/horcm0.conf

   # cp /etc/horcm.conf /etc/horcm1.conf

2. Modify **/etc/services** to register the port name/number (service) of each configuration definition file. The port name/number must be different for each CCI instance.

   horcm0 xxxxx/udp xxxxx = the port name/number for horcm0.conf

   horcm1 yyyyy/udp yyyyy = the port name/number for horcm1.conf

3. If you want HORCM to start automatically each time the system starts up, add **/etc/horcmstart.sh 0 1** to the system automatic start-up file (e.g., **/sbin/rc**).

4. Execute the **horcmstart.sh** script manually to start the CCI instances:

   # horcmstart.sh 0 1

5. Set an instance number to the environment which executes a command:

   For B shell:

   # HORCMINST=X **X = instance number = 0 or 1**

   # export HORCMINST

   C shell:

   # setenv HORCMINST X

6. Set the log directory (**HORCC_LOG**) in the command execution environment as needed.

7. If you want to perform Synchronous TrueCopy 9500V operations, do not set the **HORCC_MRCF** environment variable. If you want to perform ShadowImage 9500V operations, set the **HORCC_MRCF** environment variable for the HORCM execution environment.

   For B shell:

   # HORCC_MRCF=1

   # export HORCC_MRCF

   C shell:

   # setenv HORCC_MRCF 1

8. Execute the pairdisplay command to verify the configuration.

   # pairdisplay -g xxxx xxxx = group-name

3.5.2 Startup for Windows® Systems

3.5.2.1 One Instance

To start up one instance of CCI on a Windows NT®/2000 system:

1. Copy the sample file (\HORCM\etc\horcm.conf) to the specified directory.

   D:\HORCM\etc> \HORCM\etc\horcm.conf \WINNT\horcm.conf

2. Modify **\WINNT\system32\drivers\etc\services** to register the port name/number (service) of the configuration definition file. Make the port name/number the same on all servers:

   horcm xxxxx/udp xxxxx = the port name/number of horcm.conf

3. If you want HORCM to start automatically each time the system starts up, add **\HORCM\etc\horcmstart** to the system automatic start-up file (e.g., **\autoexec.bat**).

4. Execute the **horcmstart** script manually to start CCI:

   D:\HORCM\etc> horcmstart

5. Set the log directory (**HORCC_LOG**) in the command execution environment as needed.

6. If you want to perform Synchronous TrueCopy 9500V operations, do not set the **HORCC_MRCF** environment variable. If you want to perform ShadowImage 9500V operations, set the **HORCC_MRCF** environment variable for the HORCM execution environment:

   D:\HORCM\etc> set HORCC_MRCF=1

7. Execute the pairdisplay command to verify the configuration.

   D:\HORCM\etc> pairdisplay -g xxxx xxxx = group name

3.5.2.2 Two Instances

To start up two instances of CCI on a Windows NT®/2000 system:

1. Make two copies of the sample configuration definition file.

   D:\HORC\etc> copy \HORCM\etc\horcm.conf \WINNT\horcm0.conf

   D:\HORC\etc> copy \HORCM\etc\horcm.conf \WINNT\horcm1.conf

2. Modify **\WINNT\system32\drivers\etc\services** to register the port name/number (service) of the configuration definition files. Make sure that the port name/number is different for each instance:

   horcm0 xxxxx/udp xxxxx = the port name/number of horcm0.conf

   horcm1 yyyyy/udp yyyyy = the port name/number of horcm1.conf

3. If you want HORCM to start automatically each time the system starts up, add **\HORCM\etc\horcmstart 0 1** to the system automatic start-up file (e.g., **\autoexec.bat**).

4. **4D:\HORCM\etc> horcmstart 0 1**

5. Set an instance number to the environment which executes a command:

   D:\HORCM\etc> set HORCMINST=X X = instance number = 0 or 1

6. Set the log directory (**HORCC_LOG**) in the command execution environment as needed.

7. If you want to perform Synchronous TrueCopy 9500V operations, do not set the **HORCC_MRCF** environment variable.

8. If you want to perform ShadowImage 9500V operations, set the **HORCC_MRCF** environment variable for the HORCM execution environment:

   D:\HORCM\etc> set HORCC_MRCF=1

9. Execute the pairdisplay command to verify the configuration.

   D:\HORCM\etc> pairdisplay -g xxxx xxxx = group name

Chapter 4 Performing CCI Operations

4.1 Important Notice

Note the following when performing CCI operations:

- Do not execute Synchronous TrueCopy 9500V/ShadowImage 9500V operations while formatting the volume. Formatting takes priority and Synchronous TrueCopy 9500V/ShadowImage 9500V operations will be suspended.
- Remember to change the default value for the poll(10ms) parameter in the configuration definition file. For details on calculating the value for poll(10ms), see section 2.5.3. For details on configuration parameters, see section 3.4.
- When an internal process conflict occurs between the CCI and the subsystem, the processing of the subsystem is temporarily suspended. If the conflict continues, internal processing may not proceed. Therefore, when monitoring (polling) the status of the subsystem (by creating a script using the CCI commands) set the display-information-based commands (e.g. pairdisplay, raidscan, raidar, and raidqry) to be issued more than or equal to a minute.
- Commands that change the status of pairs (paircreate, pairsplit, pairresync) cannot be executed while changing the microprogram online.
- Do not change the microprogram online while executing commands that change the status of pairs (paircreate, pairsplit, pairresync). The execution time for the copying process varies; changing the microprogram online suspends the copying operation temporarily.

  The processing time for changing the microprogram online is 2 min./CTL (the copying process will be suspended for 2 minutes per CTL). The waiting process in the batch file may end abnormally when executing a copy (using a batch file) by designating a specific time.
- All command options displayed in help do not perform.

  The options that can perform are limited to the options described in this manual.

4.2 Environment Variables

When activating HORCM or initiating a command, users can specify any of the environment variables shown in Table 4.1.

Table 4.1 Environment Variables (Continued on next page)

Variable	Functions
HORCM (**/etc/horcmgr**) environment variables	**$HORCM_CONF**: Names the HORCM configuration file, default = /etc/horcm.conf. **$HORCM_LOG**: Names the HORCM log directory, default = /HORCM/log/curlog. **$HORCM_TRCSZ**: Specifies the size of the HORCM trace file in kB, default = 1 MB. The trace file size cannot be changed using the **horcctl** command. **$HORCM_TRCLVL**: Specifies the HORCM trace level (0 - 15), default = 4. If a negative value is specified, trace mode is canceled. The trace level can be changed using the **horcctl -c -l** command. **$HORCM_TRCBUF**: Specifies the HORCM trace mode. If this variable is specified, data is written in the trace file in the non-buffer mode. If not, data is written in the buffer mode. The trace mode can be changed using the **horcctl -c -b** command. **$HORCM_TRCUENV**: Specifies whether or not to succeed the trace control parameters (TRCLVL and TRCBUF) as they are when a command is issued. When this variable is specified, the HORCM default trace control parameters are used to the trace control parameters of HORCM as global parameters. If not, the default trace control parameters for HORCM commands are used and tracing level = 4, trace mode = buffer mode. **$HORCMFCTBL**: Changes the fibre address conversion table number, used when the target ID indicated by the raidscan command is different than the TID on the system (see Appendix D). **$HORCMPROMOD**: Sets HORCM forcibly to protection mode. Command devices in non-protection mode can be used as protection mode also. **$HORCMPERM**: Specifies the file name for the protected volumes. When this variable is not specified, the default name is as follows (* as an instance number): For UNIX systems: /etc/horcmperm.conf or /etc/horcmperm*.conf For Windows NT/2000 systems: \WINNT\horcmperm.conf or \WINNT\horcmperm*.conf When the variable is set HORCMPERM = MGFNOINST, the built-in command will not execute. Set this variable when you want to execute any command from the user shell script. For details on protected volumes, see section 2.5.6.
Synchronous TrueCopy 9500V command environment variables	**$HORCC_LOG**: Specifies the command log directory name, default = /HORCM/log* (* = instance number). When a magic character STDERROUT is set for this variable, the log output changes to standard error output. Use this character when you want to execute commands from the script file knowing that the error would occur, in order to inhibit log output. **$HORCC_TRCSZ**: Specifies the size of the command trace file in kB, default = HORCM trace file size. The default HORCM trace file size can be changed using **horcctl -d -s**. **$HORCC_TRCLVL**: Specifies the command trace level (0 = 15), default = 4 or the specified HORCM trace level. If a negative value is specified, trace mode is canceled. The default trace level for HORCM commands can be changed using the **horcctl -d -l**. **$HORCC_TRCBUF**: Specifies the command trace mode. If specified, data is written in the trace file in the non-buffer mode. If not, the HORCM trace mode is used. The default trace mode for HORCM commands can be changed using the **horcctl -d -b**.

Table 4.1 Environment Variables (continued)

Variable	Functions
Synchronous TrueCopy 9500V instance environment variable	**$HORCMINST**: Specifies the instance number when using two or more CCI instances on the same server. The command execution environment and the HORCM activation environment require an instance number to be specified. Set the configuration definition file (HORCM_CONF) and log directories (HORCM_LOG and HORCC_LOG) for each instance.
ShadowImage 9500V command environment variables	**$HORCC_MRCF**: Sets the command execution environment of the ShadowImage 9500V commands. The selection whether the command functions as that of the Synchronous TrueCopy 9500V or the ShadowImage 9500V is made according to this variable. The HORCM is not affected by this variable. When issuing a Synchronous TrueCopy 9500V command, do not set the HORCC_MRCF variable for the execution environment of the command. When issuing a ShadowImage 9500V command, set the environmental variable HORCC_MRCF=1 for the execution environment of the command.

4.3 Creating Pairs (Paircreate)

WARNING: Use the paircreate command with caution. The paircreate command starts the Synchronous TrueCopy 9500V/ShadowImage 9500V initial copy operation, which overwrites all data on the secondary volume. If the primary and secondary volumes are not identified correctly, or if the wrong options are specified (e.g., **vl** instead of **vr**), data will be transferred in the wrong direction.

The **paircreate** command generates a new volume pair from two unpaired volumes. The paircreate command can create either a paired logical volume or a group of paired volumes. The paircreate command allows you to specify the direction (local or remote) of the pair generation (see Figure 4.1). If local (**vl** option) is specified, the server issuing the paircreate command has the primary volume. If remote (**vr** option) is specified, the remote server has the primary volume. The **-split** option of the paircreate command (ShadowImage 9500V only) allows you to simultaneously create and split pairs using a single CCI command. When **-split** is used, the pair status changes from COPY to PSUS (instead of PAIR) when the initial copy operation is complete. Table 4.2 lists and describes the paircreate command parameters and returned values.

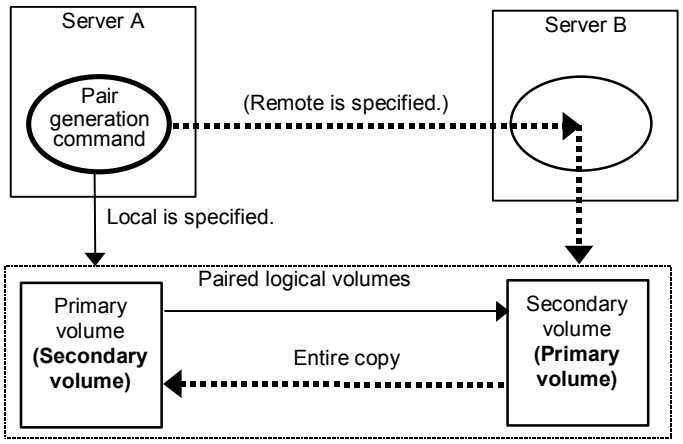

Figure 4.1 Pair Creation

Before issuing the paircreate command, verify that the secondary volume is not mounted on any system. If the secondary volume is mounted after paircreate, delete the pair (pairsplit -S), unmount the secondary volume, then reissue the paircreate command.

Note: The paircreate command terminates before the initial copy operation is complete (except when the **nocopy** option is specified). Use the pair event waiting or pair display command to verify that the initial copy operation completed successfully (status changes from COPY to PAIR, or from COPY to PSUS if the **-split** option was specified). The execution log file also shows completion of the initial copy operation.

A maximum number of 15 must always be specified for the copying pace. For details, see the -c<size> option in Table 4.2.

The operation of creating a pair is shown below.

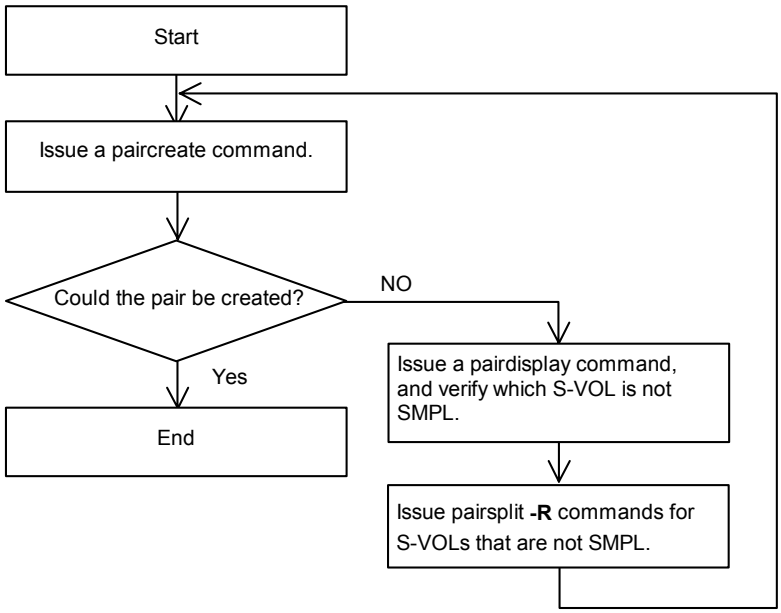

Figure 4.2 Pair Creation Operation

Table 4.2 Paircreate Command Parameters (Continued on next page)

Parameter	Value
Command Name	paircreate
Format	paircreate { -h | -q | -z | -g <group> | -d <pair Vol> | -d[g] <raw_device> [MU#] | -d[g] <seq#> <LDEV#> [MU#] -f <fence> [CTGID] | -v | -c <size> | nocopy | -nomsg | -split | [-m <mode>] }
Options	**-h**: Displays Help/Usage and version information. **-q**: Terminates the interactive mode and exits the command. **-z** or **-zx**: Makes the paircreate command enter the interactive mode. The **-zx** option guards performing of the HORCM in the interactive mode. When this option detects a HORCM shut down, interactive mode terminates. **-g <group>**: Specifies a group name defined in the configuration definition file. The command is executed for the specified group unless the -d <pair Vol> option is specified. **-d <pair Vol>**: Specifies paired logical volume name defined in the configuration definition file. When this option is specified, the command is executed for the specified paired logical volume. **-d[g] <raw_device> [MU#]**: Searches a group on the configuration definition file (local instance) for the specified raw_device, and if the specified raw_device is contained in the group, the target volume is executed as the paired logical volume (-d) or group (-dg). This option is effective without specification of "-g <group> " option. If the specified the raw_device is contained in two or more groups, the command is executed on the first group. **-d[g] <seq#> <LDEV#> [MU#]**: Searches a group on the configuration definition file (local instance) for the specified LDEV, and if the specified LDEV is contained in the group, the target volume is executed as the paired logical volume (-d) or group (-dg). This option is effective without specification of "-g <group> " option. If the specified LDEV is contained in two or more groups, the command is executed on the first group. The **<seq #> <LDEV #>** values can be specified in hexadecimal (by addition of "0x ") or decimal notation. **-f <fence> [CTGID]** (Synchronous TrueCopy 9500V only): Specifies the level for assuring the consistency of paired volume data. A fence level of "never" or "data" must be specified. This option is required for Synchronous TrueCopy 9500V. If the CTGID (CT group ID) option is not specified, the pair is automatically assigned to a new group. If CTGID is not specified and the maximum number of CT groups (16) exists. The CTGID option is used to make paired volumes of specified group forcibly by given CTGID of another group. **-vl** or **-vr**: Specifies the data flow direction and must always be specified. The -vl option specifies "local" and the host which issues the command possesses the primary volume. The -vr option specifies "remote" and the remote host possesses the primary volume while the local host possesses the secondary volume. **-c <size>**: This option must always be specified. You can use this option to specify the copying pace (1 - 15) to be used for the initial data copy. You can shorten the copy time by specifying a large number. Do not use the default value 3. For Synchronous TrueCopy 9500V/ShadowImage 9500V operations, the maximum value 15 must always be specified. **-nocopy** (Synchronous TrueCopy 9500V): Creates paired volumes without copying data in the case in which the data consistency of simplex volumes is assured by the user. **-nomsg**: Suppresses messages to be displayed when this command is executed. It is used to execute this command from a user program. This option must be specified at the beginning of a command argument. The command execution log is not affected by this option. **-split** (ShadowImage 9500V only): Splits the paired volume after the initial copy operation is complete. **-m <mode> (ShadowImage 9500V only)**: For mode, you can specify the following option: **noread:** Specifies the noread mode for hiding the secondary volume. The secondary volume becomes read-disabled when this mode option is specified. The secondary volume is read-enabled when this mode option is omitted. ***Note:*** The primary volume becomes read-disabled only during a reverse resync operation (**restore** option of pairresync command).

Table 4.2 Paircreate Command Parameters (continued)

Parameter	Value
Returned values	Normal termination: 0. When creating groups, 0 = normal termination for all pairs. Abnormal termination: Returns a common error code and a unique error code for this command. For details on error codes, see Chapter 5.

```
C:\HORCM\etc>paircreate –g VG01 –vl –c 15
C:\HORCM\etc>pairevtwait –g VG01 –s pair –t 300 10
pairevtwait : Wait status done.

C:\HORCM\etc>pairdisplay –g VG01
Goup PairVol(L/R) (Port#,TID,LU-M),Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M
VG01 oradb1(L)    (CL1-A , 1, 1-0)  174   1..P-VOL  PAIR, 174     2   -
VG01 oradb1(R)    (CL1-A , 1, 2-0)  174   2..S-VOL  PAIR,-----    1   -
```

Figure 4.3 Paircreate Command

4.4 Splitting Pairs (Pairsplit)

The **pairsplit** command stops updates to the secondary volume of a pair (see Figure 4.4). The pairsplit command allows read/write access to the secondary volume. The pairsplit command can be applied to a paired logical volume or a group of paired volumes. Table 4.3 lists and describes the pairsplit command parameters and returned values.

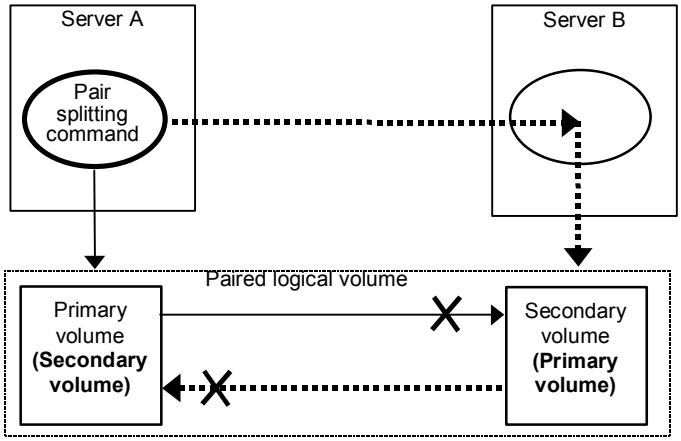

Figure 4.4 Pair Splitting

The pairsplit command allows read/write access to the secondary volume, depending on the selected options (-r, -rw, -S, -R, -P, -E). The primary volume's server is automatically detected by the pairsplit command, so the server does not need to be specified in the pairsplit command parameters. If the -S option (simplex) is used, the volume pair is deleted, the volumes are returned to the simplex state, and the primary and secondary volume status is lost. Paired volumes are split as soon as the pairsplit command is issued. If you want to synchronize the volumes, the pairsplit command must be issued after write I/Os to the paired volume have been completed.

When splitting a pair, whether or not you can change the pair status of S-VOL, changing the pair status of P-VOL to SMPL takes priority. Therefore, if the pair status of S-VOL cannot be changed to SMPL, the pair status of P-VOL might not correspond with that of S-VOL.

When a path failure has occurred or when S-VOL is not used temporarily (for example, while formatting), the pair status of S-VOl cannot be changed to SMPL.

When you want to change the status of P-VOL and S-VOL to SMPL, a pairdisplay command is issued and a pairsplit **-R** command is issued for S-VOL that is not SMPL. (This can be specified for Synchronous TrueCopy 9500V).

When the pairsplit **-R** is issued for S-VOL in paired state, a pairsplit **-S** is issued for P-VOL to change the status of P-VOL to SMPL.

Table 4.3 Pairsplit Command Parameters (Continued on next page)

Parameter	Value
Command Name	pairsplit
Format	pairsplit {-h | -q | -z | -g <group> | -d <pair Vol> | -d[g] <raw_device> [MU#] | -FHORC | -FMRCF [MU#] | -d[g] <seq#> <LDEV#> [MU#] –r | -rw | -S | -R | -P | -I | -nomsg | -C <size> | -E }
Options ***Note:*** Only one pairsplit option (-r, -rw,-S, -P, -C, -E) can be specified. If more than one option is specified, only the last option will be executed.	**-h**: Displays Help/Usage and version information. **-q**: Terminates the interactive mode and exits this command. **-z** or **-zx**: Makes the pairsplit command enter the interactive mode. The **-zx** option guards performing of the HORCM in the interactive mode. When this option detects a HORCM shut down, interactive mode terminates. **-g <group>**: Specifies a group name defined in the configuration definition file. This option must always be specified. The command is executed for the specified group unless the -d <pair Vol> option is specified. **-d <pair Vol>**: Specifies the paired logical volume name defined in the configuration definition file. When this option is specified, the command is executed for the specified paired logical volumes. **-d[g] <raw_device> [MU#]**: Searches a group on the configuration definition file (local instance) for the specified raw_device, and if the specified raw_device is contained in the group, the target volume is executed as the paired logical volume (-d) or group (-dg). This option is effective without specification of "-g <group>" option. If the specified the raw_device is contained in two or more groups, the command is executed on the first group. **FHORC** or **-FCA**: Forcibly specifies a cascading Synchronous TrueCopy 9500V volume for specified pair logical volumes on a ShadowImage 9500V environment (see example in Figure 4.4). If the **-I** option is specified, this option splits a cascading Synchronous TrueCopy 9500V volume on a local host (near site). If no **-I** option is specified, this option splits a cascading Synchronous TrueCopy 9500V volume on a remote host (far site). The target Synchronous TrueCopy 9500V volume must be a P-VOL and the **-R** option cannot be specified. **-FMRCF [MU#]** or **–FBC [MU#]**: Forcibly specifies a cascading ShadowImage 9500V volume for specified pair logical volumes on a Synchronous TrueCopy 9500V environment (see example in Figure 4.7). If the **-I** option is specified, this option splits a cascading ShadowImage 9500V volume on a local host (near site). If no **-I** option is specified, this option splits a cascading ShadowImage 9500V volume on a remote host (far site). The target ShadowImage 9500V volume must be a P-VOL and the **-E** option cannot be specified. **-d[g] <seq#> <LDEV#> [MU#]**: Searches a group on the configuration definition file (local instance) for the specified LDEV, and if the specified LDEV is in the group, the target volume is executed as the paired logical volume (-d) or group (-dg). This option is effective without specification of "-g <group>" option. If the specified LDEV is contained in two or more groups, the command is executed on the first group. The **<seq #> <LDEV #>** values can be specified in hexadecimal (by addition of "0x") or decimal notation. **-r or –rw** (for Synchronous TrueCopy 9500V): Specifies a mode of access to the S-VOL after paired volumes are split. The -r option (default) allows read-only from the S-VOL. The -rw option enables read and write access for the S-VOL. **-S**: Selects simplex mode (releases the pair). When the pairing direction is reversed among the hosts (e.g., disaster recovery), this mode is established once, and then the paircreate command is issued. If you want to re-establish a pair which has been released, you must use the paircreate command (not pairresync). **-R** (for Synchronous TrueCopy 9500V): Brings the secondary volume into the simplex mode forcibly. It is issued by the secondary host, if the host possessing the primary volume is down or has failed. **-P**: Brings the primary volume into the blocking mode (PSUS) forcibly. It is issued by the secondary host to suppress data updating by the host possessing the primary volume. **-I**: When this command cannot utilize the remote host for host down, this option enables a pairsplit operation by local host only. The target volume of local host can be P-VOL or S-VOL. **-nomsg**: Suppresses messages to be displayed when this command is executed. It is used to execute a command from a user program. This option must be specified at the beginning of a command argument. The command execution log is not affected by this option.

Table 4.3 Pairsplit Command Parameters (continued)

Parameter	Value
Options (continued)	**-C <size>**: This option does not affect on operation even if it specified or omitted. **-E (ShadowImage 9500V only)**: Suspends a paired volume forcibly when a failure occurs. Not normally used.
Returned values	Normal termination: 0. When creating groups, 0 = normal termination for all pairs. Abnormal termination: Returns a common error code and a unique error code for this command. For details on error codes, see Chapter 5.

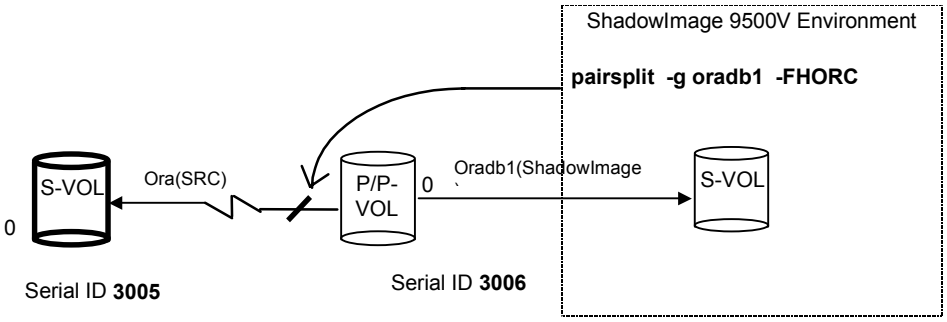

Figure 4.5 Example of -FHORC Option for Pairsplit

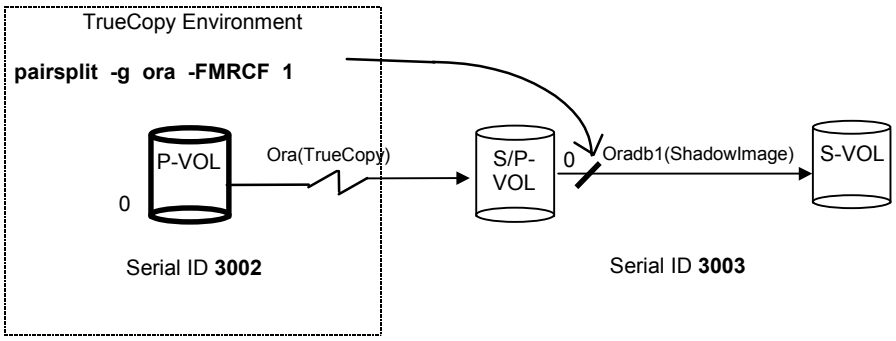

Figure 4.6 Example of -FMRCF Option for Pairsplit

```
C:\HORCM\etc>pairsplit -g VG01

C:\HORCM\etc>pairdisplay -g VG01
Group PairVol(L/R)  (Port#,TID,LU-M),Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M
VG01 oradb1(L)      (CL1-A , 1,  1-0) 174   1..P-VOL PSUS,  174    2    -
VG01 oradb1(R)      (CL1-A , 1,  2-0) 174   2..S-VOL SSUS,-----    1    -
```

Figure 4.7 Pair Split Command

```
C:\HORCM\etc>pairsplit -g VG01 –E

C:\HORCM\etc>pairdisplay -g VG01
Group PairVol(L/R)  (Port#,TID,LU-M),Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M
VG01 oradb1(L)      (CL1-A , 1,  1-0) 174   1..P-VOL PSUE,  174    2    -
VG01 oradb1(R)      (CL1-A , 1,  2-0) 174   2..S-VOL PSUE,-----    1    -
```

Figure 4.8 Pair Split Command –E Option

```
C:\HORCM\etc>pairsplit -g VG01 –S

C:\HORCM\etc>pairdisplay -g VG01
Group PairVol(L/R)  (Port#,TID,LU-M),Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M
VG01 oradb1(L)      (CL1-A , 1,  1-0) 174   1..SMPL  ----,-----    2    -
VG01 oradb1(R)      (CL1-A , 1,  2-0) 174   2..SMPL  ----,-----    1    -
```

Figure 4.9 Pair Split Command -S Option

4.4.1 Timing of Pairsplit Operations

Since a pair is split at the time when the pairsplit command is issued, verify that writing onto pair volumes is fixed before issuing the command.

- Instantaneous offline backup of the UNIX® file system:
  - Unmount the primary volume, and then split the volume pair.
  - Mount the primary volume (**mount**).
  - Verify that the pairsplit is complete, and mount the secondary volume (**mount -r**).
  - Execute the backup.
  - Restore the volumes to their previous state, and resynchronize the volume pair.
- Online backup of the UNIX® file system:
  - Issue the **sync** command to a mounted primary volume to flush the file system buffer, then split the volume pair in Read/Write mode.
  - Verify that the pairsplit is complete, then use the **fsck** command to check the consistency of the secondary volume file system.
  - Mount (**mount**) the secondary volume.
  - Execute the backup.
  - Restore the volumes to their previous state and resynchronize the volume pair.
- Instantaneous offline backup of the Windows NT®/2000 file system:
  - Execute **-x umount** of the primary volume, then split the volume pair.
  - Execute **x-mount** of the primary volume.
  - Verify that the pairsplit is complete, then execute **-x mount** of the secondary volume.
  - Execute the backup.
  - Restore the volumes to their previous state, and resynchronize the volume pair.
- Online backup of the Windows NT®/2000 file system:
  - Issue the **x-sync** command to a mounted primary volume to flush the file system buffer, then split the volume pair in Read/Write mode.
  - Verify that the pairsplit is complete, then use **x-mount** of the secondary volume.
  - Execute the backup.
  - Restore the volumes to their previous state and resynchronize the volume pair.

Note: If the primary volume is divided by LVM or partition, the control information of LVM or partition on the primary volume is also copied to the secondary volume. When executing the backup from the secondary volume, import this control information and execute pairsplit when activating the secondary volume.

4.5 Resynchronizing Pairs (Pairresync)

The **pairresync** command re-establishes a split pair, then restarts the update copy operations to the secondary volume (see Figure 4.10). The pairresync command can resynchronize either a paired logical volume or a group of paired volumes. The normal direction of resynchronization is from the primary volume to the secondary volume. If the **-restore** option is specified (ShadowImage 9500V only), the pair is resynchronized in the reverse direction (i.e., secondary volume to primary volume). Figure 4.11 shows the normal and restore resync operations. The primary volume is read and write accessible during pairresync. The secondary volume becomes write-disabled when the pairresync command is issued.

Table 4.4 lists and describes the pairresync command parameters and returned values. The primary volume's server is automatically detected by the pairresync command, so the server does not need to be specified in the pairresync command parameters.

The pairresync command terminates before resynchronization of the secondary (or primary) volume is complete. Use the pair event waiting or pair display command to verify that the resync operation completed successfully (status changes from COPY to PAIR). The execution log file also shows completion of the resync operation. The status transition of the paired volume is judged by the status of the primary volume.

If no data was written to the secondary volume while the pair was split, the differential data on the primary volume is copied. If data was written to the secondary volume, the differential data on the primary volume and secondary volume is copied to the secondary volume. This process is reversed when the ShadowImage 9500V **-restore** option is specified.

Before issuing the pairresync command (normal or reverse direction), make sure that the secondary volume is not mounted on any UNIX® system. Before issuing a reverse pairresync command, make sure that the primary volume is not mounted on any UNIX® system.

Note: A maximum number of 15 must always be specified for the copying pace. For details, see the –c<size> option in Table 4.4.

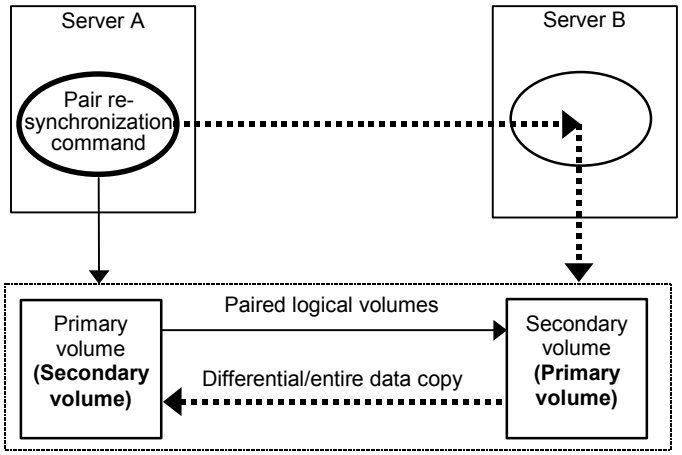

Figure 4.10 Pair Resynchronization

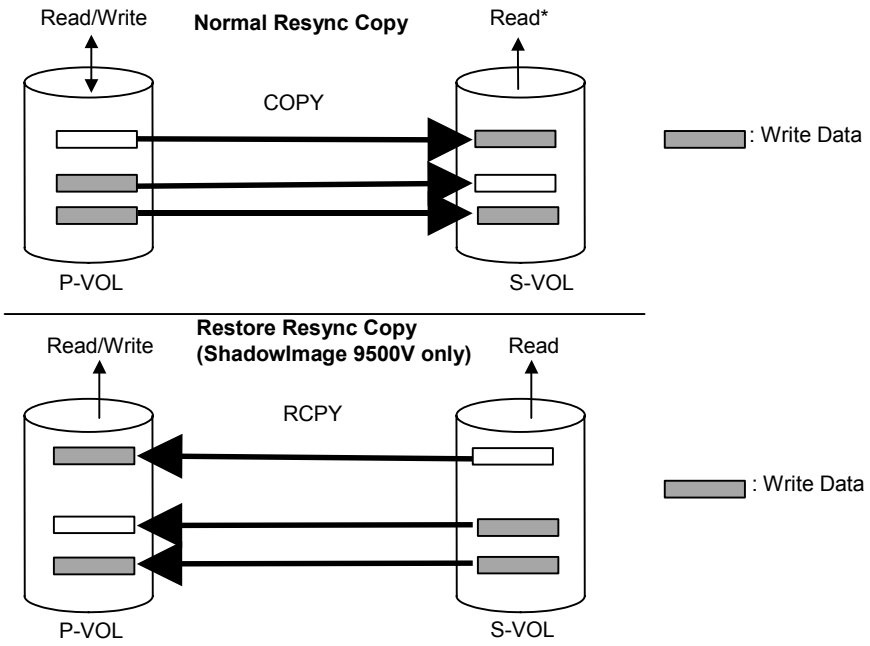

Read*: Read disabled when the paircreate –m noread option is specified.

Figure 4.11 Normal Resync and ShadowImage 9500V Restore Resync

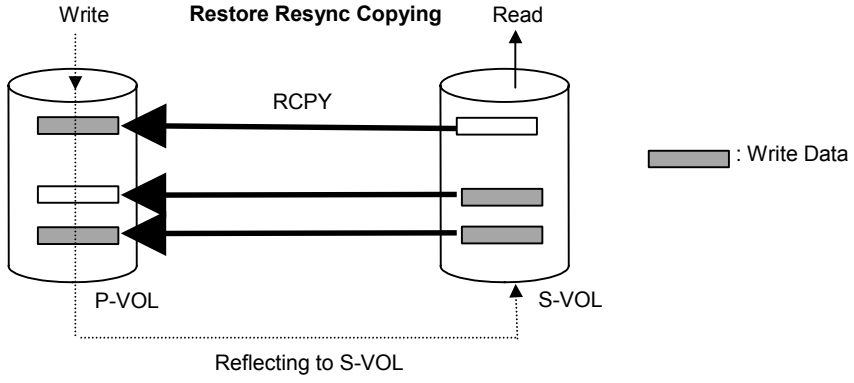

Figure 4.12 Reflecting Write Data to P-VOL during ShadowImage 9500V Restore Resync Copy

Table 4.4 Pairresync Command Parameters (Continued on next page)

Parameter	Value
Command Name	pairresync
Format	pairresync { -h | -q | -z | -g <group> | -d <pair Vol> | -d[g] <raw_device> [MU#] | -FHORC | -FMRCF [MU#] | -d[g] <seq#> <LDEV#> [MU#] | -c <size> | -nomsg | -l | -restore | -swaps | -swapp }
Options	**-h**: Displays Help/Usage and version information. **-q**: Terminates the interactive mode and exits this command. **-z** or **-zx**: Makes the pairresync command enter the interactive mode. The **-zx** option guards performing of the HORCM in the interactive mode. When this option detects a HORCM shut down, interactive mode terminates. **-g <group>**: This option is used to specify a group name defined in the configuration definition file. This option must always be specified. The command is executed for the specified group unless the -d <pair Vol> option is specified. **-d <pair Vol>**: Specifies a paired logical volume name defined in the configuration definition file. When this option is specified, the command is executed for the specified paired logical volumes. **-d[g] <raw_device> [MU#]**: Searches a group on the configuration definition file (local instance) for the specified raw_device, and if the specified raw_device is contained in the group, the target volume is executed as the paired logical volume (-d) or group (-dg). This option is effective without specification of "-g <group>" option. If the specified the raw_device is contained in two or more groups, the command is executed on the first group. **-FHORC** or **-FCA**: Forcibly specifies a cascading Synchronous TrueCopy 9500V volume for specified pair logical volumes on a ShadowImage 9500V environment (see example in Figure 4.13). If the **-l** option is specified, this option resyncs a cascading Synchronous TrueCopy 9500V volume on a local host (near site). If no **-l** option is specified, this option resyncs a cascading Synchronous TrueCopy 9500V volume on a remote host (far site). The target Synchronous TrueCopy 9500V volume must be a P-VOL. **-FMRCF [MU#]** or **–FBC [MU#]**: Forcibly specifies a cascading ShadowImage 9500V volume for specified pair logical volumes on a Synchronous TrueCopy 9500V environment (see example in Figure 4.14). If the **-l** option is specified, this option resyncs a cascading ShadowImage 9500V volume on a local host (near site). If no **-l** option is specified, this option resyncs a cascading ShadowImage 9500V volume on a remote host (far site). The target ShadowImage 9500V volume must be a P-VOL. **-d[g] <seq#> <LDEV#> [MU#]**: Searches a group on the configuration definition file (local instance) for the specified LDEV, and if the specified LDEV is contained in the group, the target volume is executed as the paired logical volume (-d) or group (-dg). This option is effective without specification of the "-g <group>" option. If the specified LDEV is contained in two or more groups, the command is executed on the first group. The **<seq #> <LDEV #>** values can be specified in hexadecimal (by addition of "0x") or decimal notation. **<seq #> <LDEV #>** values can be specified in hexadecimal (by addition of "0x ") or decimal notation. **-c <size>**: This option must always be specified. You can use this option to specify the copying pace for the resync operation (range = 1 to 15 track extents). For ShadowImage 9500V operations, the maximum value 15 must always be specified. If omitted, the value used for paircreate will be used. **-nomsg**: Suppresses messages to be displayed when this command is executed. It is used to execute this command from a user program. This option must be specified at the beginning of a command argument. The command execution log is not affected by this option. **-l**: When this command cannot utilize the remote host for host down, this option enables a pairresync operation by the local host only. The target volume of the local host can be P-VOL or S-VOL. **-restore**: Performs reverse resync (from secondary volume to primary volume). **-swaps** (Synchronous TrueCopy 9500V only): Executed from the S-VOL side when there is no host on the P-VOL side to help. Typically executed in **PSUS** state to facilitate "fast failback" without requiring a full copy. In Figure 4.16, the left side shows T0 for both the P-VOL and S-VOL (before command execution), and the right side shows T1, after the command has executed. For both -swaps and -swapp, the delta data from the original S-VOL becomes dominant and is copied to the original P-VOL, then the S/P-VOL designations are swapped.

Table 4.4 Pairresync Command Parameters (continued)

Parameter	Value
Options (continued)	**-swapp** (Synchronous TrueCopy 9500V only): Executes the equivalent of a -swaps from the original P-VOL side. Unlike -swaps, -swapp does require the cooperation of hosts at both sides.
Returned Values	Normal termination: 0. When resynchronizing groups, 0 = normal termination for all pairs. Abnormal termination: Returns a common error code and a unique error code for this command. For details on error codes, see Chapter 5.

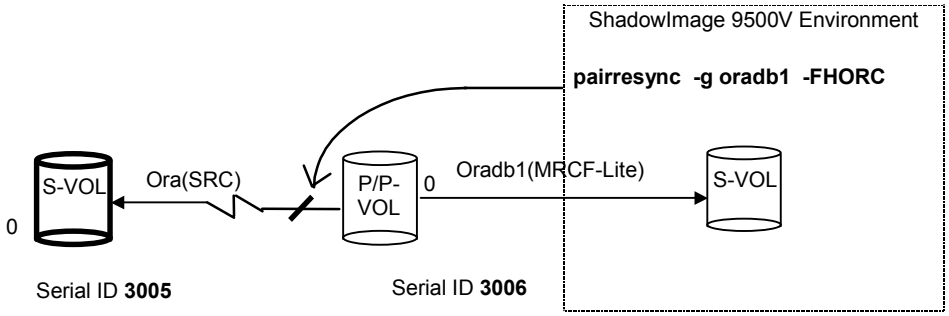

Figure 4.13 Example of -FHORC Option for Pairresync

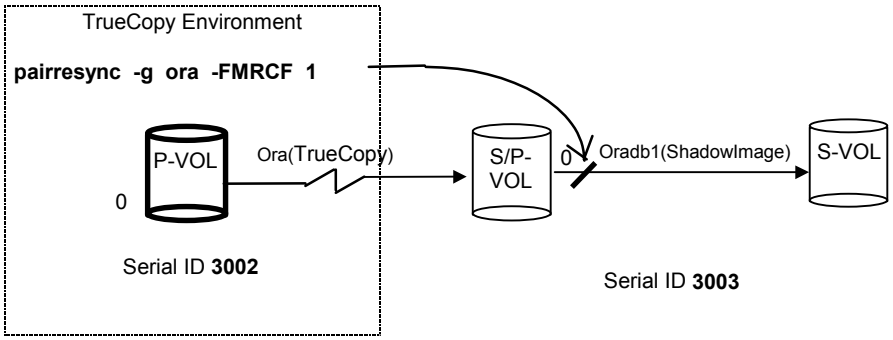

Figure 4.14 Example of -FMRCF Option for Pairresync

```
C:\HORCM\etc>pairresync -g VG01 -c 15

C:\HORCM\etc>pairevtwait -g VG01 -s pair -t 300 10
pairevtwait : Wait status done.

C:\HORCM\etc>pairdisplay -g VG01
Group PairVol(L/R)  (Port#,TID,LU-M),Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M
VG01 oradb1(L)      (CL1-A ,1,  1-0) 174     1..P-VOL PAIR,  174    2    -
VG01 oradb1(R)      (CL1-A ,1,  2-0) 174     2..S-VOL PAIR,-----    1    -
```

Figure 4.15 Pairresync Command

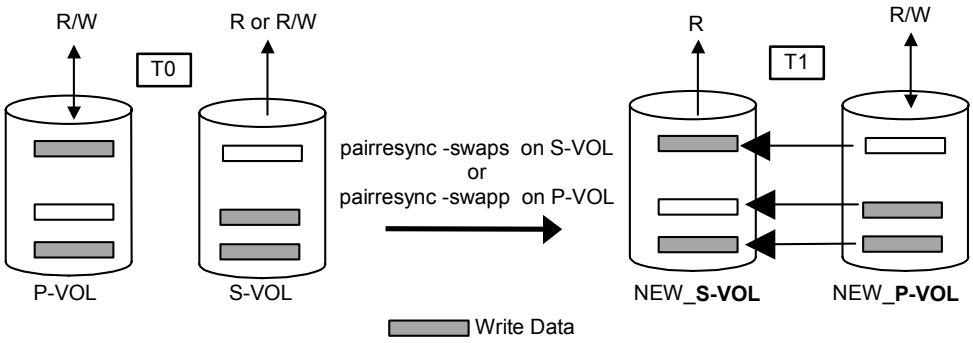

Figure 4.16 Swap Operation

4.6 Confirming Pair Operations (Pairevtwait)

The pair event waiting (**pairevtwait**) command is used to wait for completion of pair creation and pair resynchronization and to check status (see Figure 4.17). It waits ("sleeps") until the paired volume status becomes identical to a specified status, then completes. The pairevtwait command can be used for a paired logical volume or a group of paired volumes. The primary volume's server is automatically detected by the pair event waiting command, so the server does not need to be specified in the pair event waiting command parameters.

Table 4.5 lists and describes the pair event waiting command parameters and returned values. The pair event waiting command waits until the specified status is established, and terminates abnormally if an abnormal status is detected. The transition of the paired volume status is judged by the status of the primary volume. If the event waiting command is issued for a group, the command waits until the status of each volume in the group becomes identical to the specified status. When the event waiting command with the **-nowait** option is issued for a group, the status is returned if the status of each volume in the group is identical. This command must be used to confirm a pair status transition.

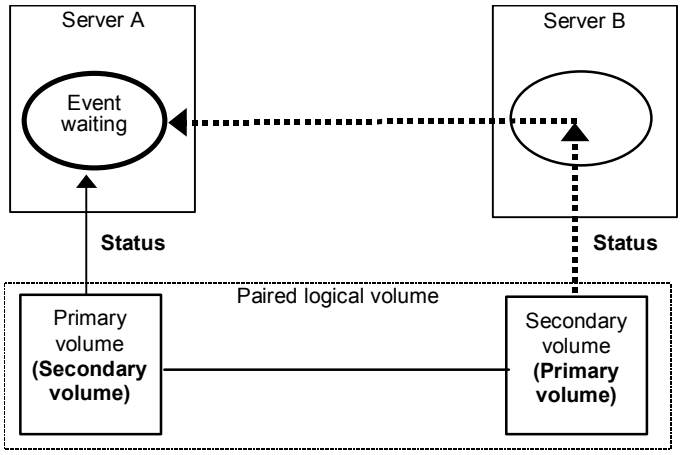

Figure 4.17 Pair Event Waiting

Table 4.5 Pairevtwait Command Parameters (Continued on next page)

Parameter	Value
Command Name	pairevtwait
Format	pairevtwait { -h | -q | -z | -g <group> | -d <pair Vol> | -d[g] <raw_device> [MU#] | -FHORC | -FMRCF [MU#] | -d[g] <seq#> <LDEV#> [MU#] | -s <status> ... | -t <timeout>[interval] | -nowait | -I | -nomsg }
Options	**-h**: Displays Help/Usage and version information. **-q**: Terminates the interactive mode and exits this command. **-z** or **-zx**: Makes the pairevtwait command enter the interactive mode. The **-zx** option guards performing of the HORCM in the interactive mode. When this option detects a HORCM shut down, interactive mode terminates. **-g <group>**: Specifies a group name defined in the configuration definition file. This option must always be specified. The command is executed for the specified group unless the -d <pair Vol> option is specified. **-d <pair Vol>**: Specifies a paired logical volume name defined in the configuration definition file. When this option is specified, the command is executed for the specified paired logical volumes. **-d[g] <raw_device> [MU#]**: Searches a group on the configuration definition file (local instance) for the specified raw_device, and if the specified raw_device is contained in the group, the target volume is executed as the paired logical volume (-d) or group (-dg). This option is effective without specifying the "-g <group>" option. If the specified the raw_device is contained in two or more groups, the command is executed on the first group. **-FHORC** or **-FCA**: Forcibly specifies a cascading Synchronous TrueCopy 9500V volume for specified pair logical volumes on a ShadowImage 9500V environment (see example in Figure 4.18). If the **-I** option is specified, this option tests status of a cascading Synchronous TrueCopy 9500V volume on a local host (near site). If no **-I** option is specified, this option tests status of a cascading Synchronous TrueCopy 9500V volume on a remote host (far site). The target Synchronous TrueCopy 9500V volume must be P-VOL or SMPL. **-FMRCF [MU#]** or **–FBC [MU#]**: Forcibly specifies a cascading ShadowImage 9500V volume for specified pair logical volumes on a Synchronous TrueCopy 9500V environment (see example in Figure 4.19). If the **-I** option is specified, this option tests status of a cascading ShadowImage 9500V volume on a local host (near site). If no **-I** option is specified, this option tests status of a cascading MRCF-Lite volume on a remote host (far site). The target ShadowImage 9500V volume must be P-VOL or SMPL. **-d[g] <seq#> <LDEV#> [MU#]**: Searches a group on the configuration definition file (local instance) for the specified LDEV, and if the specified LDEV is contained in the group, the target volume is executed as the paired logical volume (-d) or group (-dg). This option is effective without specifying the "-g <group>" option. If the specified LDEV is contained in two or more groups, the command is executed on the first group. The **<seq #> <LDEV #>** values can be specified in hexadecimal (by addition of "0x") or decimal notation. **-s <status>**: **..**Specifies the waiting status, which is "smpl", "copy/rcpy", "pair", "psus", or "psue/pdub". If two or more statuses are specified following -s, waiting is done according to the logical OR of the specified statuses. This option is valid when the -nowait option is not specified. **-t <timeout> [interval]**: Specifies the interval of monitoring a status specified using the -s option and the time-out period in units of 1 sec. Unless [interval] is specified, the default value is used. This option is valid when the -nowait option is not specified. **-nowait**: When this option is specified, the pair status at that time is reported without waiting. The pair status is set as a returned value for this command. When this option is specified, the -t and -s options are not needed. **-I**: When this command cannot utilize a remote host for host down, this option executes this command by a local host only. The target volume of the local host can be P-VOL or S-VOL. **-nomsg**: Suppresses messages to be displayed when this command is executed. It is used to execute a command from a user program. This option must be specified at the beginning of a command argument. The command execution log is not affected by this option.

Table 4.5 Pairevtwait Command Parameters (Continued)

Parameter	Value
Returned values	When the **-nowait** option is specified: Normal termination: 1: The status is **SMPL**. 2: The status is **COPY** or **RCPY**. 3: The status is **PAIR**. 4: The status is **PSUS**. 5: The status is **PSUE**. When monitoring groups, 1/2/3/4/5 = normal termination for all pairs. Abnormal termination: Returns a common error code and a unique error code for this command. For details on error codes, see Chapter 5. When the **-nowait** option is not specified: Normal termination: 0. When monitoring groups, 0 = normal termination for all pairs. Abnormal termination: Returns a common error code and a unique error code for this command. For details on error codes, see Chapter 5.

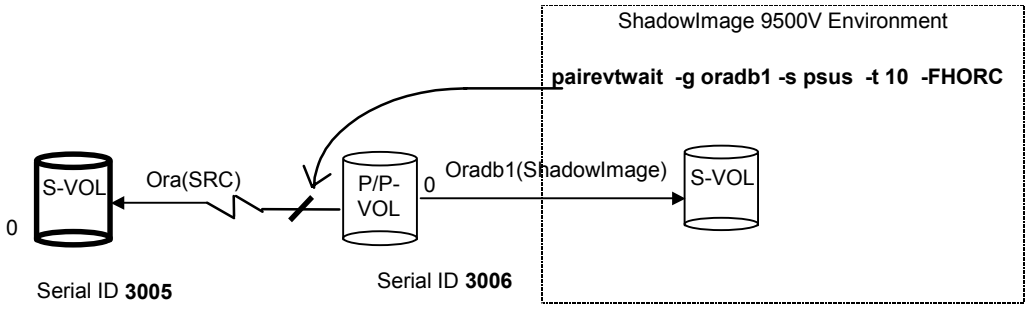

Figure 4.18 Example of -FHORC Option for Pairevtwait

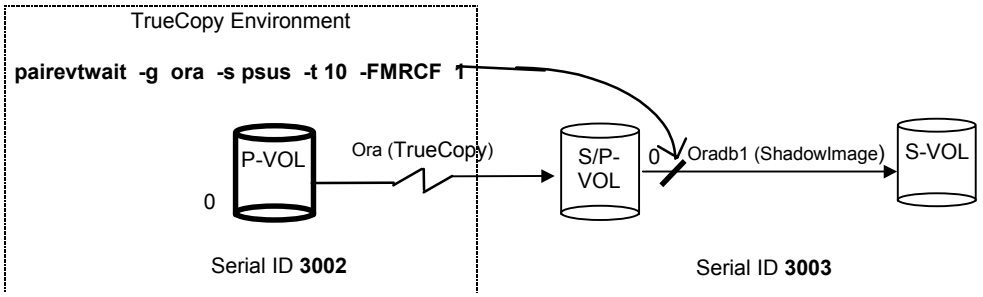

Figure 4.19 Example of -FMRCF Option for Pairevtwait

```
C:\HORCM\etc>paircreate –g VG01 –c 15
C:\HORCM\etc>pairevtwait –g VG01 –s pair –t 300 10
pairevtwait : Wait status done.
```

Figure 4.20 Pairevtwait Command

4.7 Monitoring Pair Activity (Pairmon)

The **pairmon** command, connected to the HORCM daemon, obtains the pair status transition of each volume pair and reports it. If the pair status changes (due to an error or a user-specified command), the pairmon command issues a message. Table 4.6 lists and describes the pairmon command parameters. Figure 4.21 shows an example of the pairmon command and its output. Table 4.7 lists the results of the command options.

The pair status transition events exist in the HORCM pair status transition queue. The **-resevt** option (reset event) deletes one/all events from the HORCM pair status transition queue. If reset event is not specified, the pair status transition queue is maintained. If the **-s** option is not specified, pairmon displays all events for which it receives information from HORCM. If the **-s** option is specified, only the specified status transitions are displayed.

The CCI software supports the error monitoring and configuration confirmation commands for linkage with the system operation management of the UNIX® server.

Table 4.6 Pairmon Command Parameters

Parameter	Value
Command Name	pairmon
Format	pairmon { -h | -q | -z | -D | -allsnd | -resevt | -nowait | -s <status> ... }
Options	**-h**: Displays Help/Usage and version information. **-q**: Terminates the interactive mode and exits this command. **-z** or **-zx**: Makes the pairmon command enter the interactive mode. The **-zx** option guards performing of the HORCM in the interactive mode. When this option detects a HORCM shut down, interactive mode terminates. **-D**: Selects the default report mode. In the default mode, if there is pair status transition information to be reported, one event is reported and the event is reset. If there is no pair status transition information to be reported, the command waits. The report mode consists of the three flags: -allsnd, - resevt, and -nowait options. **-allsnd**: Reports all events if there is pair status transition information. **-resevt**: Reports events if there is pair status transition information, then resets all events. **-nowait**: When this option is specified, the command does not wait when there is no pair status transition information. **-s <status> ...**: Specifies the pair status transition to be reported: smpl, copy (includes rcpy), pair, psus, psue. If two or more statuses are specified following -s, masking is done according to the logical OR of the specified statuses. If this option is not specified, pairmon displays all events which received information from HORCM.

```
# pairmon  -allsnd   -nowait
Group Pair vol      Port    targ#   lun#    LDEV#…  Oldstat code –> Newstat code
oradb oradb1 CL1-A  1       5       5…      SMPL    0x00 –> COPY    0x01
oradb oradb2 CL1-A  1       6       6…      PAIR    0x02 –> PSUS    0x04
```

Figure 4.21 Pairmon Command

The output of the pairmon command includes:

- **Group**: Shows the group name (dev_group) described in the configuration definition file.
- **Pair vol**: Shows the paired volume name (dev_name) in the specified group which is described in the configuration definition file.
- **Port targ# lun#**: Shows the port ID, TID, and LUN described in the configuration definition file. For further information on Fibre-to-SCSI address conversion, see Appendix D.
- **LDEV#**: Shows the Thunder 9500V LDEV ID for the specified device. LDEV indicates LU for Synchronous TrueCopy 9500V/ShadowImage 9500V operations.
- **Oldstat**: Shows the old pair status when the status of the volume is changed.
- **Newstat**: Shows the new pair status when the status of the volume is changed.
- **Code**: Shows the 9500V-internal code for the specified status.

Table 4.7 Results of Pairmon Command Options

-D	-nowait	-resevt	-allsnd	Actions
-D				When HORCM does not have an event, this option waits until an event occurs. If one or more events exist, it reports one event and resets the event which it reported.
Invalid			-allsnd	When HORCM does not have an event, this option waits until an event occurs. If one or more events exist, it reports all events.
Invalid		-resevt		When HORCM does not have an event, this option waits until an event occurs. If one or more events exist, it reports one event and resets all events.
Invalid		-resevt	-allsnd	When HORCM does not have an event, this option waits until an event occurs. If one or more events exist, it reports all events and resets all events.
Invalid	-nowait			When HORCM does not have an event, this option reports event nothing. If one or more events exist, it reports one event and resets the event which it reported.
Invalid	-nowait		-allsnd	When HORCM does not have an event, this option reports event nothing. If one or more events exist, it reports all events.
Invalid	-nowait	-resevt		When HORCM does not have an event, this option reports event nothing. If one or more events exist, it reports one event and resets all events.
Invalid	-nowait	-resevt	-allsnd	When HORCM does not have an event, this option reports event nothing. If one or more events exist, it reports all events and resets all events.

4.8 Checking Attribute and Status (Pairvolchk)

The **pairvolchk** command acquires and reports the attribute of a volume or group connected to the local host (issuing the command) or remote host. The volume attribute is SMPL (simplex), P-VOL (primary volume), or S-VOL (secondary volume). The **-s[s]** option reports the pair status in addition to the attribute. Table 4.8 lists and describes the pairvolchk command parameters and returned values. Figure 4.24 shows an example of the pairvolchk command and its output. Table 4.9 shows the truth table for pairvolchk group status display.

Table 4.8 Pairvolchk Command Parameters (Continued on next page)

Parameter	Value
Command Name	pairvolchk
Format	pairvolchk { -h | -q | -z | -g <group> | -d <pair Vol> | -d[g] <raw_device> [MU#] | -FHORC | -FMRCF [MU#] | -d[g] <seq#> <LDEV#> [MU#] | -c | -ss | -nomsg }
Options	**-h**: Displays Help/Usage and version information. **-q**: Terminates the interactive mode and exits the pair volume check command. **-z** or **-zx**: Makes the pairvolchk command enter the interactive mode. The **-zx** option guards performing of the HORCM in the interactive mode. When this option detects a HORCM shut down, interactive mode terminates. **-g <group>**: Specifies the group name defined in the configuration definition file. This option must always be specified. The command is executed for the specified group unless the -d <pair Vol> option is specified. **-d <pair Vol>**: This option is used to specify the paired logical volume name defined in the configuration definition file. When this option is specified, the command is executed for the specified paired logical volumes. **-d[g] <raw_device> [MU#]**: Searches a group on the configuration definition file (local instance) for the specified raw_device, and if the specified raw_device is contained in the group, the target volume is executed as the paired logical volume (-d) or group (-dg). This option is effective without specifying the "-g <group>" option. If the specified the raw_device is contained in two or more groups, the command is executed on the first group. **-FHORC** or **-FCA**: Forcibly specifies a cascading Synchronous TrueCopy 9500V volume for specified pair logical volumes on a ShadowImage 9500V environment (see example in Figure 4.22). If no **-c** option is specified, this option acquires the attributes of a cascading Synchronous TrueCopy 9500V volume on a local host (near site). If the **-c** option is specified, this option acquires the attributes of a cascading Synchronous TrueCopy 9500V volume on a remote host (far site). **-FMRCF [MU#]** or **-FBC [MU#]**: Forcibly specifies a cascading ShadowImage 9500V volume for specified pair logical volumes on a Synchronous TrueCopy 9500V environment (see example in Figure 4.23). If no **-c** option is specified, this option acquires the attributes of a cascading ShadowImage 9500V volume on a local host (near site). If the **-c** option is specified, this option acquires the attributes of a cascading ShadowImage 9500V volume on a remote host (far site). **-d[g] <seq#> <LDEV#> [MU#]**: Searches a group on the configuration definition file (local instance) for the specified LDEV, and if the specified LDEV is contained in the group, the target volume is executed as the paired logical volume (-d) or group (-dg). This option is effective without specifying the "-g <group> " option. If the specified LDEV is contained in two or more groups, the command is executed on the first group. The **<seq #> <LDEV #>** values can be specified in hexadecimal (by addition of "0x") or decimal notation. **-c**: Checks the conformability of the paired volumes of the local and remote hosts and reports the volume attribute of the remote host. If this option is not specified, the volume attribute of the local host is reported. **-ss**: Used to acquire the pair status of a volume in addition to the volume attribute. If this option is not specified, only the volume attribute is reported. **-nomsg**: Suppresses messages to be displayed when this command is executed. It is used to execute a command from a user program. This option must be specified at the beginning of a command argument. The command execution log is not affected by this option.

Table 4.8 Pairvolchk Command Parameters (continued)

Parameter	Value
Returned values	When the -ss option is not specified: Normal termination: 1: The volume attribute is SMPL. 2: The volume attribute is P-VOL. 3: The volume attribute is S-VOL. Abnormal termination: Returns a common error code and a unique error code for this command. For details on error codes, see Chapter 5. When the -ss option is specified: Normal termination: 11: The status is SMPL. 22: The status is PVOL_COPY or PVOL_RCPY. 23: The status is PVOL_PAIR. 24: The status is PVOL_PSUS. 25: The status is PVOL_PSUE. 32: The status is SVOL_COPY or SVOL_RCPY. 33: The status is SVOL_PAIR. 34: The status is SVOL_PSUS. 35: The status is SVOL_PSUE. Abnormal termination: Returns a common error code and a unique error code for this command. For details on error codes, see Chapter 5.

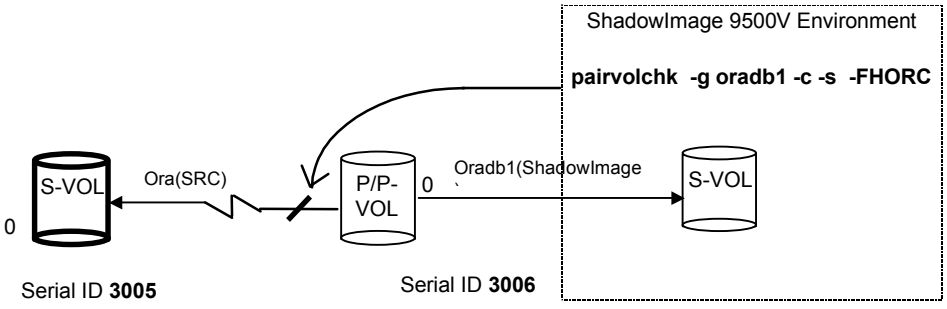

Figure 4.22 Example of -FHORC Option for Pairvolchk

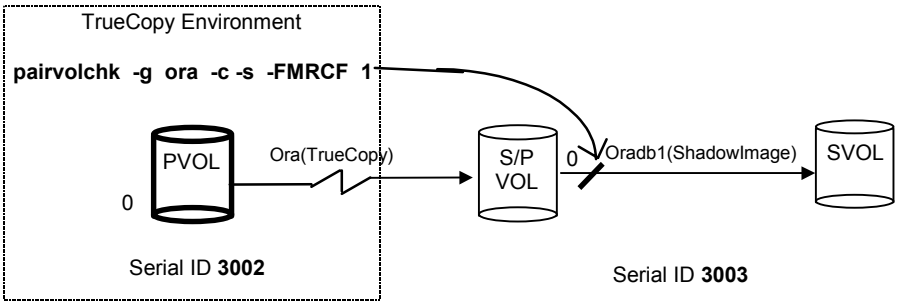

Figure 4.23 Example of -FMRCF Option for Pairvolchk

```
# pairvolchk  -g   oradb
Pairvolchk : Volstat is P-VOL.
```

Figure 4.24 Pairvolchk Command

Table 4.9 Truth Table for Pairvolchk Group Status Display

Option	Status of Each Volume in the Group				Group Status
	COPY*	PSUE	PSUS	PAIR	
Note*	TRUE	x	x	x	COPY*
	false	TRUE	x	x	PSUE
	false	false	TRUE	x	PSUS
	false	false	false	TRUE	PAIR
-ss	TRUE	x	x	x	COPY*
	false	TRUE	x	x	PSUE
	false	false	x	TRUE	PAIR
	false	false	TRUE	false	PSUS

*Note = This option can be executed when the pairvolchk –s command has a USEOLDVCHK environment variable.
*COPY = COPY or RCPY
x = true or false (does not matter).

4.9 Displaying Pair Status (Pairdisplay)

The **pairdisplay** command displays the pair status; this allows you to verify that pair operations are complete (e.g., paircreate, pairresync). The pairdisplay command is also used to confirm the configuration of the pair connection path (the physical link of paired volumes and servers). The pairdisplay command can be used for a paired volume or a group of paired volumes. Table 4.10 lists and describes the pairdisplay command parameters and returned values. Figure 4.25 shows examples of the pairdisplay command and its output.

Table 4.10 Pairdisplay Command Parameters (Continued on next page)

Parameter	Value
Command Name	pairdisplay
Format	pairdisplay { -h | -q | -z | -g <group> | -d <pair Vol> | -d[g] <raw_device> [MU#] | -FHORC | -FMRCF [MU#] | -d[g] <seq#> <LDEV#> [MU#] | -c | -I | -f[xcd] | -CLI | -m <mode> }
Options	**-h**: Displays Help/Usage and version information. **-q**: Terminates the interactive mode and exits the pair volume check command. **-z** or **-zx**: Makes the pairdisplay command enter the interactive mode. The **-zx** option guards performing of the HORCM in the interactive mode. When this option detects a HORCM shut down, interactive mode terminates. **-g <group>**: Specifies the group name defined in the configuration definition file. This option must always be specified. The command is executed for the specified group unless the -d <pair Vol> option is specified. **-d <pair Vol>**: This option is used to specify the paired logical volume name defined in the configuration definition file. When this option is specified, the command is executed for the specified paired logical volumes. **-d[g] <raw_device> [MU#]**: Searches a group on the configuration definition file (local instance) for the specified raw_device, and if the specified raw_device is contained in the group, the target volume is executed as the paired logical volume (-d) or group (-dg). This option is effective without specification of "-g <group>" option. If the specified the raw_device is contained in two or more groups, the command is executed on the first group. **-FHORC** or **-FCA**: Forcibly specifies a cascading Synchronous TrueCopy 9500V volume for specified pair logical volumes on a ShadowImage 9500V environment. If the **-I** option is specified, this option displays status of a cascading Synchronous TrueCopy 9500V volume on a local host (near site). If no **-I** option is specified, this option displays status of a cascading Synchronous TrueCopy 9500V volume on a remote host (far site). This option cannot be specified with **-m <mode>** option on the same command line. **-FMRCF [MU#]** or **–FBC [MU#]**: Forcibly specifies a cascading ShadowImage 9500V volume for specified pair logical volumes on a Synchronous TrueCopy 9500V environment. If the **-I** option is specified, this option displays status of a cascading ShadowImage 9500V volume on a local host (near site). If no **-I** option is specified, this option displays status of a cascading ShadowImage 9500V volume on a remote host (far site). This option cannot be specified with **-m <mode>** option on the same command line. **-d[g] <seq#> <LDEV#> [MU#]**: Searches a group on the configuration definition file (local instance) for the specified LDEV, and if the specified LDEV is contained in the group, the target volume is executed as the paired logical volume (-d) or group (-dg). This option is effective without specification of "-g <group>" option. If the specified LDEV is contained in two or more groups, the command is executed on the first group. The **<seq #> <LDEV #>** values can be specified in hexadecimal (by addition of "0x") or decimal notation. **-c**: Checks the configuration of the paired volume connection path (physical link of paired volume among the servers) and displays illegal pair configurations. If this option is not specified, the status of the specified paired volume is displayed without checking the path configuration. **-I**: Displays the paired volume status of the local host (which issues this command).

Table 4.10 Pairdisplay Command Parameters (continued)

Parameter	Value
Options (continued)	**-fc**: Displays the copy operation progress or the bitmap percentage. Used to confirm **SSWS** state as indication of SVOL_SSUS-takeover after. **-fx**: Displays the LDEV ID as a hexadecimal number. **-fd**: Displays the group-based device file of the configuration file registered in HORCM and the relation between the groups. If Unknown is displayed in the DEVICE_FILE shown below, the pair operation (except for the local option) is rejected as protection mode because the volume is not registered. ***Example:*** # pairdisplay –g oradb –fd Group PairVol(L/R) Device_File M,Seq#, LDEV#.P/S, Status, Seq#, P-LDEV# M oradb oradb1(L) c0t3d0 0 5013 17..P-VOL COPY, 5013 18 - oradb oradb1(R) Unknown 0 5013 ****..------ -------, ------ ----- **--CLI**: Used to specify display for command line interface (CLI). This option displays to the same position that defined number of columns, and displays one header. The delimiters between columns are displayed as spaces or hyphens (-). Display example: Group PairVol L/R Port# TID LU-M Seq# LDEV# P/S Status Seq# P-LDEV# M homrcf1 deva0 L CL1-A 3 5 0 3005 5 P-VOL PAIR 3005 3 - homrcf1 deva1 L CL1-A 3 5 0 3005 5 SMPL - - - - homrcf1 deva2 L CL1-A 3 5 0 3005 5 SMPL - - - - **-m <mode>**: Used to display a paired status of each mirror descriptors for the specified pair logical volume, and used for paired status indication of the cascading volume. **<mode>**option can be designated “all “ (see Figure 4.26): “**all**” option is used to display a paired status of all mirror descriptors (MU#).
Returned values	1: The volume attribute is **SMPL**. 2: The volume attribute is **P-VOL**. 3: The volume attribute is **S-VOL**. When displaying groups, 1/2/3 = normal termination for all pairs. Abnormal termination (other than 0 to 127): refer to the execution log files for error details.

```
# pairdisplay  -g oradb
Group
p PairVol(L/R) (Port#,TID,LU),Seq#,LDEV#….P/S,Status,  Fence, Seq#,P-LDEV# M
oradb oradb1(L)   (CL1-A  , 1, 1) 3005  1…….P-VOL COPY   Never, 3006  19     -
oradb oradb1(R)   (CL1-B  , 2, 1) 3006  1…….S-VOL COPY   Never, ----  18     -
```

Figure 4.25 Synchronous TrueCopy 9500V Pairdisplay Command

The output of the pairdisplay command includes:

- **Group**: Shows the group name (dev_group) described in the configuration definition file
- **Pair Vol(L/R)**: Shows the paired volume name (dev_name) described in the configuration definition file. (L) indicates the local host; (R) indicates the remote host
- **(Port#,TID,LU) (Synchronous TrueCopy 9500V)**: Shows the port number, target ID and LU number as described in the configuration definition file
- **(Port#,TID,LU-M) (ShadowImage 9500V)**: Shows the port number, target ID, LU number, and MU number as described in the configuration definition file
- **Seq#**: Shows the serial number (Serial ID) of the Thunder 9500™ V subsystem

- **LDEV#**: Shows the logical device number. LDEV indicates LU for Synchronous TrueCopy 9500V/ShadowImage 9500V operations.
- **P/S**: Shows the volume attribute
- **Status**: Shows the status of the paired volume
- **Fence (for Synchronous TrueCopy 9500V)**: Shows fence level
- **%:** Shows copy operation completion or percent pair synchronization

State	Synchronous TrueCopy 9500V			ShadowImage 9500V		
VOL.	COPY	PAIR	OTHER	COPY	PAIR	OTHER
P-VOL	CR	BMP	BMP	CR	CR	CR
S-VOL	-	BMP	BMP	CR	CR	CR

CR: Shows the copy operation rate (identical rate of a pair).
BMP: Shows the identical percentage of BITMAP both P-VOL and S-VOL.

- **P-LDEV#**: Shows the LDEV number of the partner volume of the pair
- **M**
  M = W (PSUS only): For P-VOLs, shows suspension with the S-VOL read/write enabled. For S-VOLs, shows that write I/Os were accepted at the S-VOL.
  M = N (COPY, PAIR, PSUE only): Shows that read access is disabled

Figure 4.26 shows examples of the **-m** option of the pairdisplay command.

Display examples for -m all:

```
# pairdisplay -g oradb -m all
Group   PairVol(L/R) (Port#,TID,LU-M), Seq#,  LDEV#..P/S,  Status, Seq#, P-LDEV# M
oradb   oradev1(L)   (CL1-A , 3,  0-0) 3005     0....SMPL  ----,   ----  ----    -
oradb   oradev1(L)   (CL1-A , 3,  0-0)   3005   0....P-VOL PAIR,   3005  26      -
oradb1  oradev11(R)  (CL1-A , 3,  2-0) 3005     2....P-VOL COPY,   3005  27      -
oradb   oradev1(R)   (CL1-A , 3,  2-0)   3005   2....S-VOL COPY,   ----- 25      -
# pairdisplay -d /dev/rdsk/c0t3d0 -l -m all
Group  PairVol(L/R)  (Port#,TID,LU-M), Seq#,  LDEV#..P/S,  Status, Seq#, P-LDEV# M
oradb  oradev1(L)    (CL1-A , 3,  0-0) 3005     0....SMPL  ----,   ----   ----   -
oradb  oradev1(L)    (CL1-A , 3,  0-0)   3005   0....P-VOL PAIR,   3005   26     -
```

Figure 4.26 Pairdisplay -m

4.10 Checking Synchronous TrueCopy 9500V Pair Currency (Paircurchk)

The Synchronous TrueCopy 9500V paircurchk command checks the currency of the Synchronous TrueCopy 9500V secondary volume(s) by evaluating the data consistency based on pair status and fence level. Table 4.11 specifies the data consistency for each possible state of a Synchronous TrueCopy 9500V volume. A paired volume or group can be specified as the target of the paircurchk command. The paircurchk command assumes that the target is an S-VOL. If the paircurchk command is specified for a group, the data consistency of each volume in the group is checked, and all inconsistent volumes are found in the execution log file and displayed. Paircurchk is also executed as part of the Synchronous TrueCopy 9500V takeover command (see next section).

The Synchronous TrueCopy 9500V supports the paircurchk command.

Table 4.11 Data Consistency Displayed by the Paircurchk Command

Object Volume			Currency	
Attribute	**Status**	**Fence**	**Paircurchk**	**SVOL_takeover**
SMPL	–	–	To be confirmed	–
P-VOL	–	–	To be confirmed	–
S-VOL	COPY	Data		
		Never		
	PAIR	Data	OK	OK
		Never	To be analyzed	To be analyzed
	PSUS	Data	Suspected	Suspected
		Never	Suspected	Suspected
	PSUE	Data	OK	OK
		Never	Suspected	Suspected
	SSWS	Data	Suspected	–
		Never	Suspected	

Notes:
1. To be confirmed = It is necessary to check the object volume, since it is not the secondary volume.
2. To be analyzed = Whether S-VOL has a mirror consistency or not cannot be judged by the status of S-VOL.
If the status of P-VOL is PAIR, the mirror consistency is OK. If the status of P-VOL is PSUS or PSUE, the mirror consistency is suspected.
3. Suspected = S-VOL has no mirror consistency.
4. OK = Mirroring consistency is assured.

Table 4.12 Paircurchk Command Parameters

Parameter	Value
Command Name	paircurchk
Format	paircurchk { -h | -q | -z | -g <group> | -d <pair Vol> | -d[g] <raw_device> [MU#] | -d[g] <seq#> <LDEV#> [MU#] | -nomsg }
Options	**-h**: Displays Help/Usage and version information. **-q**: Terminates the interactive mode and exits the pair volume check command. **-z** or **-zx**: Makes the pairvolchk command enter the interactive mode. The **-zx** option guards performing of the HORCM in the interactive mode. When this option detects a HORCM shut down, interactive mode terminates. **-g <group>**: Specifies the group name defined in the configuration definition file. This option must always be specified. The command is executed for the specified group unless the -d <pair Vol> option is specified. **-d <pair Vol>**: This option is used to specify the paired logical volume name defined in the configuration definition file. When this option is specified, the command is executed for the specified paired logical volumes. **-d[g] <raw_device> [MU#]**: Searches a group on the configuration definition file (local instance) for the specified raw_device, and if the specified raw_device is contained in the group, the target volume is executed as the paired logical volume (-d) or group (-dg). This option is effective without specification of "-g <group> " option. If the specified the raw_device is contained in two or more groups, the command is executed on the first group. **-d[g] <seq#> <LDEV#> [MU#]**: Searches a group on the configuration definition file (local instance) for the specified LDEV, and if the specified LDEV is contained in the group, the target volume is executed as the paired logical volume (-d) or group (-dg). This option is effective without specification of "-g <group> " option. If the specified LDEV is contained in two or more groups, the command is executed on the first group. The **<seq #> <LDEV #>** values can be specified in hexadecimal (by addition of "0x ") or decimal notation. **-nomsg**: Suppresses messages to be displayed when this command is executed. It is used to execute a command from a user program. This option must be specified at the beginning of a command argument. The command execution log is not affected by this option.
Returned values	Normal termination: 0. Abnormal termination: Returns a common error code and a unique error code for this command. For details on error codes, see Chapter 5.

```
# paircurchk  -g  oradb
paircurchk : Volume currency error.
Group  Pair vol  Port  targ# lun# LDEV#...P/S  Status  Fence   To be...
oradb  oradb1    CL1-A  1    5       30...SMPL ...     ...     Confirmed
paircurchk: [EX_VOLCUR] S-Vol currency error
```

The output of the paircurchk command includes:

- **Group**: Shows the group name (dev_group) described in the configuration definition file
- **Pair Vol**: Shows the paired volume name (dev_name) described in the configuration definition file
- **Port#, targ#, lun#**: Shows the port number, target ID, and LU number as described in the configuration definition file
- **LDEV#**: Shows the logical device number. LDEV indicates LU for Synchronous TrueCopy 9500V/ShadowImage 9500V operations.
- **P/S**: Shows the volume attribute
- **Status**: Shows the status of the paired volume
- **Fence**: Shows the fence level
- **To be...**: Evaluating the mirror consistency of S-VOL shows the currency of the volume.

4.11 Performing Synchronous TrueCopy 9500V Takeover Operations (Horctakeover)

The horctakeover command is a scripted command for executing several HORC operations. The takeover command checks the specified volume's or group's attributes (paircurchk), decides the takeover function based on the attributes, executes the chosen takeover function, and returns the result. The four Synchronous TrueCopy 9500V takeover functions designed for HA software operation are: takeover-switch, swap-takeover, PVOL-takeover, and SVOL-takeover. A paired volume or a group can be specified as the target of the Synchronous TrueCopy 9500V takeover command. If Synchronous TrueCopy 9500V SVOL-takeover is specified for a group, the data consistency check is executed for all volumes in the group, and all inconsistent volumes are found in the execution log file and displayed (same as paircurchk command).

The takeover command allows swapping of the primary and secondary volumes, so that if the primary or secondary volume is switched due to a server error or package transfer, duplex operations can be continued using the reversed volumes. When control is handed over to the current node, swapping the volumes again eliminates the need to copy them. The takeover command also allows the secondary volume to be separated for disaster recovery operations.

Table 4.13 lists and describes the Synchronous TrueCopy 9500V takeover command parameters and returned values.

Table 4.13 Horctakeover Command Parameters (Continued on next page)

Parameter	Value
Command Name	horctakeover
Format	horctakeover { -h | -q | -z | -g <group> | -d <pair Vol> | -d[g] <raw_device> [MU#] | -d[g] <seq#> <LDEV#> [MU#] |-S | -I | -t<timeout> | -nomsg }
Options	**-h**: Displays Help/Usage and version information. **-q**: Terminates the interactive mode and exits the pair volume check command. **-z** or **-zx**: Makes the pairvolchk command enter the interactive mode. The **-zx** option guards performing of the HORCM in the interactive mode. When this option detects a HORCM shut down, interactive mode terminates. **-g <group>**: Specifies the group name defined in the configuration definition file. This option must always be specified. The command is executed for the specified group unless the -d <pair Vol> option is specified. **-d <pair Vol>**: This option is used to specify the paired logical volume name defined in the configuration definition file. When this option is specified, the command is executed for the specified paired logical volumes. **-d[g] <raw_device> [MU#]**: Searches a group on the configuration definition file (local instance) for the specified raw_device, and if the specified raw_device is contained in the group, the target volume is executed as the paired logical volume (-d) or group (-dg). This option is effective without specification of "-g <group> " option. If the specified the raw_device is contained in two or more groups, the command is executed on the first group. **-d[g] <seq#> <LDEV#> [MU#]**: Searches a group on the configuration definition file (local instance) for the specified LDEV, and if the specified LDEV is contained in the group, the target volume is executed as the paired logical volume (-d) or group (-dg). This option is effective without specification of "-g <group> " option. If the specified LDEV is contained in two or more groups, the command is executed on the first group. The **<seq #> <LDEV #>** values can be specified in hexadecimal (by addition of "0x ") or decimal notation. **-S**: Selects and executes SVOL-takeover. The target volume of the local host must be an S-VOL. If this option is specified, then the following "-I" option is invalid.

Table 4.13 Horctakeover Command Parameters (continued)

Parameter	Value
Options (continued)	**-l**: Enables read and write to the primary volume(s) by a local host only without a remote host, and executes PVOL-takeover when the primary volume cannot be used because it is fenced (fence = DATA, state = PSUE, or PSUE volume is contained in the group). If the primary volume can be accessed, nop-takeover is executed. The target volume of the local host must be a P-VOL. **-nomsg**: Suppresses messages to be displayed when this command is executed. It is used to execute a command from a user program. This option must be specified at the beginning of a command argument. The command execution log is not affected by this option.
Returned values	Normal termination: 0: Nop-takeover (no operation). 1: Swap-takeover was successfully executed. 2: SVOL-takeover was successfully executed. 3: PVOL-SMPL-takeover was successfully executed. 4: PVOL-PSUE-takeover was successfully executed. 5: SVOL-SSUS-takeover was successfully executed. Abnormal termination: other than 0-5, Returns a common error code and a unique error code for this command. For details on error codes, see Chapter 5.

4.11.1 Horctakeover Command Functions

4.11.1.1 Takeover-Switch Function

The control scripts activated by the HA software are used the same way by all nodes of a cluster; they do not discriminate between primary and secondary volumes. The takeover command, when activated by a control script, checks the combination of attributes of the local and remote volumes and determines the proper takeover action.

Table 4.14 lists the volume attributes and specifies the Synchronous TrueCopy 9500V takeover action for each combination of attributes.

Table 4.14 Volume Attributes and Takeover Actions

Local Node (Takeover Node)		Remote Node		
Volume Attribute	**Fence Level and Status**	**Volume Attribute**	**P-VOL Status**	**Takeover Action**
SMPL	-	SMPL	-	NG [1]
		P-VOL	-	**Nop-Takeover** [2]
		S-VOL	-	Volumes not conform [3]
		Unknown [4]	-	NG
P-VOL (primary)	Fence = Data Status = PSUE	SMPL	-	NG
		P-VOL	-	Volumes not conform
		S-VOL	-	**PVOL-Takeover**
		Unknown Status (e.g. LAN down)	-	**PVOL-Takeover**
	Fence = Never Status = others	SMPL	-	NG
		P-VOL	-	Volumes not conform
		S-VOL	-	**Nop-Takeover**
		Unknown Status (e.g. LAN down)	-	**Nop-Takeover**
S-VOL (secondary)	Status = SSWS [5] After SVOL_SSUS-takeover	Any	-	**Nop-Takeover**
	Others	SMPL	-	Volumes not conform
		P-VOL	PAIR	**Swap-Takeover**
			Others	**SVOL-Takeover**
		S-VOL	-	Volumes not conform
		Unknown	-	**SVOL-Takeover**

Notes:

1. NG = The takeover command is rejected, and the operation terminates abnormally.
2. Nop-Takeover = The takeover command is accepted, but no operation is performed.
3. Volumes not conform = The volumes are not in sync, and the takeover command terminates abnormally.
4. Unknown = The remote node attribute is unknown and cannot be identified. The remote node system is down or cannot communicate.
5. SSWS = Suspend for Swapping with S-VOL side only. The SSWS state is displayed as SSUS (SVOL_PSUS) by ALL commands except the **-fc** option of the pairdisplay command.

4.11.1.2 Swap-Takeover Function

When the P-VOL status of the remote node is PAIR and the S-VOL data is consistent, it is possible to swap the primary and secondary volumes. The swap-takeover function is used by the HA control script when a package is manually moved to an alternate data center while all hardware is operational. Swap-takeover can be specified for a paired volume or a group.

The swap-takeover function internally executes the following commands to swap the primary and secondary volumes:

1. Execute **Suspend for Swapping** for the local volume (S-VOL). If this step fails, swap-takeover is disabled and an error is returned.
2. Execute **Resync for Swapping** to switch to the primary volume for which the local volume (S-VOL) is swapped as the NEW_P-VOL. Re-synchronizes the NEW_S-VOL based on the NEW_P-VOL. As for copy pace, if the remote host is known, the command will use the value of P-VOL specified at paircreate time. If the remote host is unknown, the command will use the default number of pace (three). If this step fails, swap-takeover returns at **SVOL-SSUS-takeover**, and the local volume (S-VOL) is maintained in SSUS(PSUS) state which allows and keeps track of write I/Os using a bitmap for the S-VOL This special state is displayed as **SSWS** using the **-fc** option of the pairdisplay command.

Note: The swap-takeover function does not use **SMPL** or **No Copy** mode for swapping to guarantee mirror consistence. This is included as a function of SVOL-takeover.

4.11.1.3 SVOL-Takeover Function

The SVOL-takeover function allows the takeover node to use the secondary volume (except in COPY state) in SSUS(PSUS) state (i.e., reading and writing are enabled), on the assumption that the remote node (possessing the primary volume) cannot be used. The data consistency of the Synchronous TrueCopy 9500V S-VOL is evaluated by its pair status and fence level (same as paircurchk). If the primary and secondary volumes are not consistent, the SVOL-takeover function fails. If primary and secondary volumes are consistent, the SVOL-takeover function attempts to switch to the primary volume using **Resync for Swapping**. If successful, the SVOL-takeover function returns **Swap-takeover** as the return value of the horctakeover command. If not successful, the SVOL-takeover function returns **SVOL-SSUS-takeover** as the return value of the horctakeover command. In case of a host failure, **Swap-takeover** is returned. In case of P-VOL site failure, **SVOL-SSUS-takeover** is returned.

SVOL-takeover can be specified for a paired volume or a group. If the SVOL-takeover is specified for a group, a data consistency check is executed for all volumes in the group, and all inconsistent volumes are displayed.

4.11.1.4 PVOL-Takeover Function

The PVOL-takeover function releases the pair state as a group, since that maintains the consistency of the secondary volume at having accepted horctakeover command when the primary volume is fenced ("data" and "PSUE" state, "PSUE" volume are contained in the group). This function allows the takeover node to use the primary volume (i.e., reading and writing are enabled), on the assumption that the remote node (possessing the secondary volume) cannot be used. PVOL-takeover can be specified for a paired volume or a group.

The PVOL-takeover function executes the following two commands:

- **PVOL-PSUE-takeover**: Changes the primary volume to the suspend (PSUE, PSUS) state which enables write I/Os to all primary volumes of the group. The action of the PVOL-PSUE-Takeover causes PSUE and/or PSUS to be intermingled in the group. This intermingled pair status is PSUE as the group status; therefore, pairvolchk command returned gives priority to PSUE rather than PSUS as the group status. This special state turns back to the original state when the pairresync command is issued.
- **PVOL-SMPL-takeover**: Changes the primary volume to the simplex (SMPL) state. First, PVOL-takeover executes PVOL-PSUE-takeover further than PVOL-SMPL-takeover. If the PVOL-PSUE-takeover function fails, the PVOL-SMPL-takeover function is executed.

4.11.2 Applications of the Horctakeover Command

The basic Synchronous TrueCopy 9500V commands (takeover, pair creation, pair splitting, pair resynchronization, event waiting) can be combined to enable recovery from a disaster, backup of paired volumes, and many other operations (e.g., restoration of paired volumes based on the secondary volume, swapping of the paired volumes). Figure 4.27 illustrates the flow of starting operations on a UNIX® server at the secondary site using the Synchronous TrueCopy 9500V takeover command. Figure 4.28 illustrates the flow of starting operations on a Windows® NT/2000 server at the secondary site using the Synchronous TrueCopy 9500V takeover command.

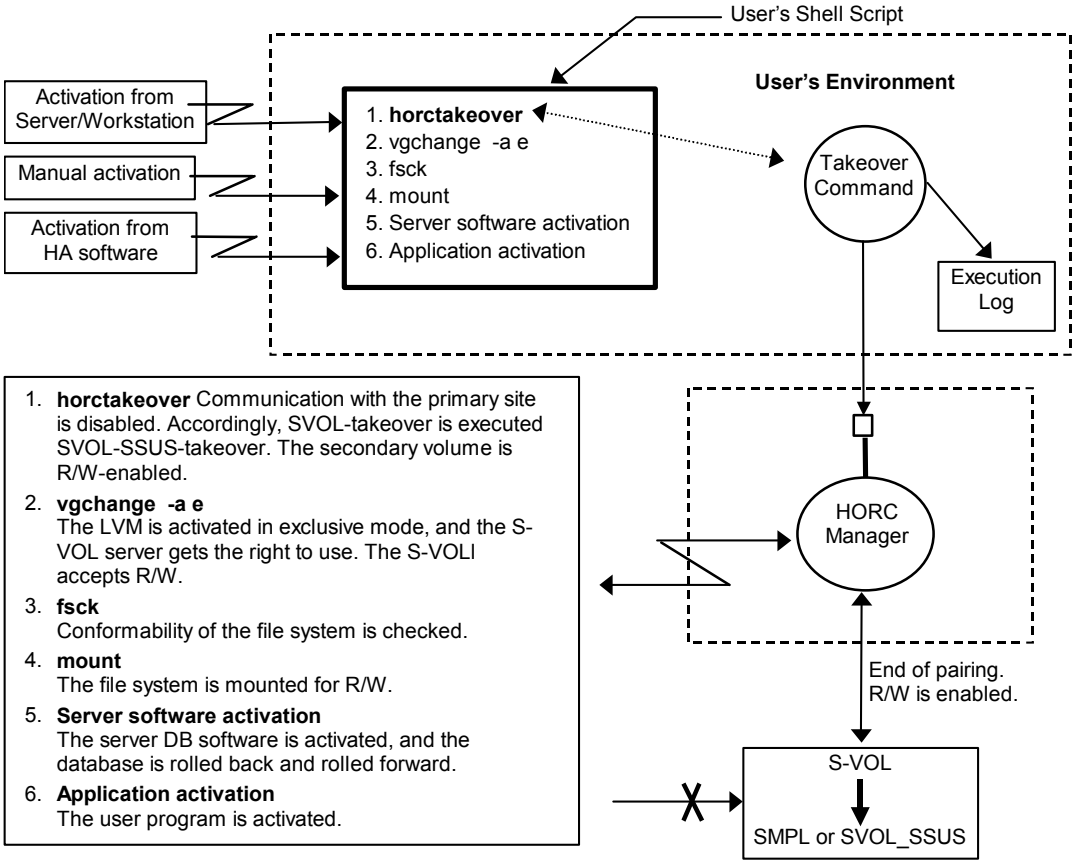

Figure 4.27 Application/Example of Synchronous TrueCopy 9500V Takeover (UNIX® -based System)

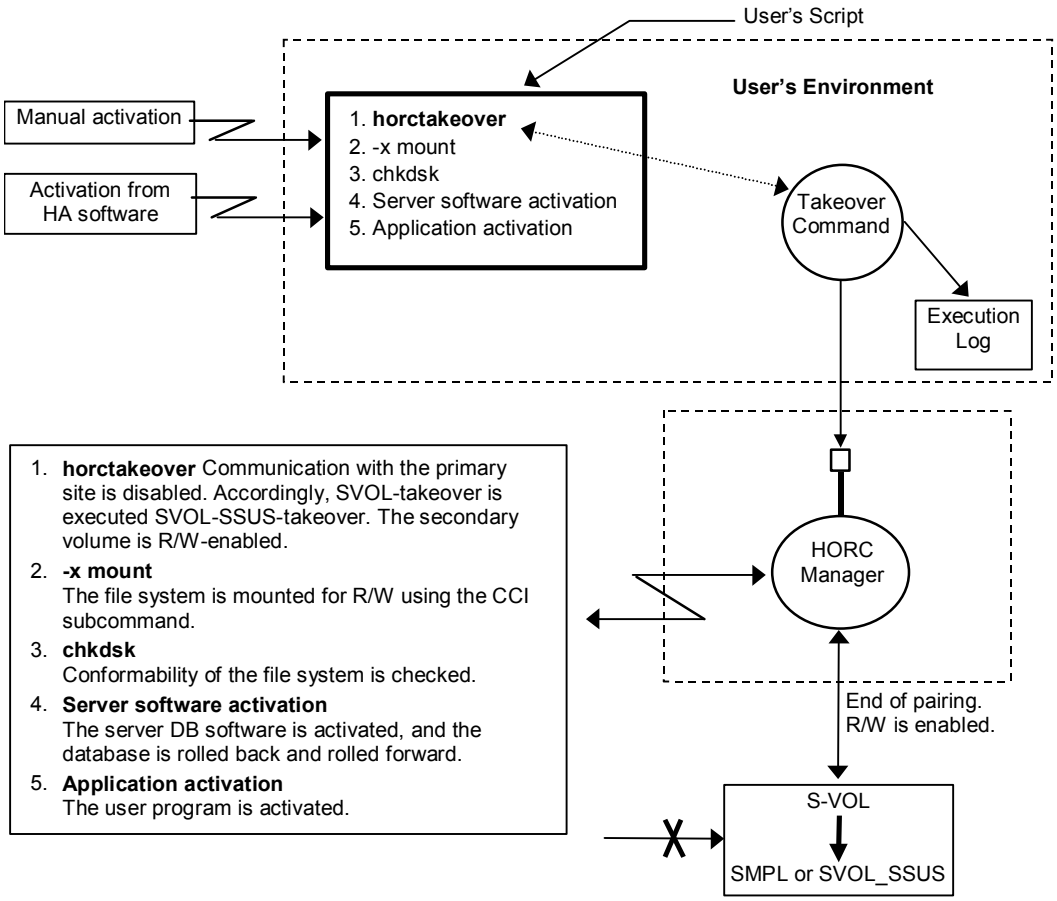

Figure 4.28 Application/Example of Synchronous TrueCopy 9500V Takeover（Windows® -based System）

4.12 Displaying Configuration Information

4.12.1 Raidscan Command

The **raidscan** command displays configuration and status information for the specified 9500V port/TID(s)/device(s). The information is acquired directly from the 9500V subsystem (not the configuration definition file). Table 4.15 lists and describes the raidscan command parameters. Figure 4.29 to Figure 4.35 list examples of the raidscan command and its output. Note that LDEV indicates LU for Synchronous TrueCopy 9500V/ShadowImage 9500V operations.

Table 4.15 Raidscan Command Parameters (Continued on next page)

Parameter	Value
Command Name	raidscan
Format	raidscan { -h | -q | -z | -p <port> [hgrp] | -pd <raw_device> | -s <Seq#> | -t <targ> | -l <lun> | [-f[xfgd]] | -CLI | -find [op] [MU#] [-g group] | -pi <strings> }

Options	**-h**: Displays Help/Usage and version information.
	-q: Terminates the interactive mode and exits the command.
	-z or **-zx**: Makes the raidscan command enter the interactive mode. The **-zx** option guards performing of the HORCM in the interactive mode. When this option detects a HORCM shut down, interactive mode terminates.
	-p <port> [hgrp]: Specifies the port ID of the 9500V port to be scanned. Valid ports are CL1-A to CL1-B and CL2-A to CL2-B. This option must always be specified. Do not specify the **[hgrp].**
	-pd <raw_device>: Specifies the raw device name. This option finds Seq# (Equipment ID) and port_name of the 9500V that the specified device can be connected, and scans the port of the 9500V which corresponds with the unit ID that searches the unit ID from Seq# (Equipment ID). This option must always be specified.
	-s <Seq#>: Used to specify the Seq# (Equipment ID) of the 9500V when this option can't specify the unit ID which is contained for "-p <port>" option. This option scans the port specified by "-p <port>" option of the 9500V which corresponds with the unit ID that searches the unit ID from Seq# (Equipment ID). If this option is specified, then the unit ID which is contained in "-p <port>" option is invalid.
	-t <targ>: Specifies a target ID (0 to 15) of the specified port. If this option is not specified, the command applies to all target IDs.
	-l <lun>: Specifies a LUN (0 to 7) of the specified target ID. If this option is not specified, the command applies to all LUNs. If this option is specified, the TID must also be specified.
	-fx: Displays the LDEV number in hexadecimal notation.
	-f or **-ff**: Specifies display of volume-type for a display column. If this option is specified, then the following **-fg** and **–fd** option is invalid.
	-fg: Specifies display of group_name for a display column. This option searches a group on the configuration definition file (local CCI instance) from the scanned LDEV, and displays a group_name when the scanned LDEV is contained in the group. If this option is specified, the **-f[f]** option is excluded. If this option is specified, the **–fd** option is invalid.
	-fd: Displays the device file registered in HORCM. This option is invalid when the **–ff** or the **–fg** option is specified.
	-CLI: Specifies display for command line interface (CLI). This option displays to the same position that defined number of columns, and displays one header. The delimiters between columns are displayed as spaces or hyphens (-). Display example:
	`Port# TargetID# Lun# Seq# Num LDEV# P/S Status P-Seq# P-LDEV#`
	`CL1-A 1 0 3005 1 0 SMPL - - -`
	`CL1-A 2 2 3005 1 2 P-VOL PAIR 3005 26`
	`CL1-A 2 3 3005 1 3 P-VOL PAIR 3005 27`

Table 4.15 Raidscan Command Parameters (continued)

Parameter	Value

Options (Continued)	**-pi<strings>**: Used to change the STDIN of the –find option to argument input. If this option is specified, the STDIN becomes invalid. Specify **<strings>** in less than or equal to 255 characters.
	
	-find [op] [MU#]: Executes the operation specified in [op] using a special file (raw device file) provided via STDIN. If the –pi <strings> option is specified, the STDIN becomes invalid and changes to <strings>.
	Restriction: Special files via STDIN are specified in the following ways:
	HP-UX® systems: /dev/rdsk/*
	Solaris® systems: /dev/rdsk/*s2, c*s2
	AIX® systems: /dev/rhdisk, /dev/hdisk.hdisk...
	Linux® systems: /dev/sd... or /dev/rd...
	Tru64 UNIX® systems: /dev/rdisk/dsk*c
	IRIX® systems: /dev/rdsk/*vol, /dev/rdsk/node_wwn/*vol/*, /dev/dsk/*vol, /dev/dsk/node_wwn/*vol/*
	Windows® NT systems: hd0-10, harddisk0...(numbers indicate the drive number) $LETALL, $Phys, D:\DskX\pY, \DskX\pY
	Windows® 2000 systems: hd0-10, harddisk0...(numbers indicate the drive number) $LETALL, $Volume, $Phys, D:\Vol(Dms,Dmt, Dmr)X\DskY, \Vol(Dms,Dmt, Dmr)X\DskY
	For Windows® 2000 LDM volumes, see Appendix E.
	
	-find: Displays the port, target ID, and LUN (9500V notation) which was mapped for LDEV using a special file (raw device file) provided via STDIN (see Figure 4.30). If **target ID and LUN are Unknown** for the target device file, the user must start HORCM without a description for HORCM_DEV and HORCM_INST, and should describe the shown Port, target ID, and LUN for HORCM_DEV. This option also uses the **-fx** option to display the LDEV numbers in hexadecimal.
	
	-find inst: Obtains the Ser# and LDEV# of the volume using a special file (raw device file) provided via STDIN, and then checks the consistency with the volume described in the configuration definition file. Then registers (permits) the special file name to HORCM, and displays the relation between the configuration definition file. This option is not usually used because it is automatically used by /etc/hormgr. This option also uses the **-fx** option to display the LDEV numbers in hexadecimal (see Figure 4.31).
	Note: When activating HORCM, the CCI automatically registers the device file by executing raidscan –find inst command. If this command is executed manually even if the device file has already been registered, a message "The registration has been canceled due to enough for HORCM." is displayed and the command will be rejected.
	
	-find verify [MU#]: Displays the relation between the volumes described in the configuration definition file after obtaining the Ser# and LDEV# of the volume using a special file (raw device file). This option also uses the **-fx** option to display the LDEV numbers in hexadecimal. This option is effected by the command execution environment (HORCC_MRCF).
	
	-find conf [MU#] [-g <group>]: Displays the image of the port, target ID, and LUN (9500V notation) which was mapped for LDEV using a special file (raw device file) provided via STDIN. Before displaying the information, this option edits the information into an image that looks like the configuration definition file. If **target ID and LUN are Unknown** for the target device file, the user must start HORCM without a description for HORCM_DEV and HORCM_INST, and should describe the shown Port, target ID, and LUN for HORCM_DEV. This option also uses the **-fx** option to display the LDEV numbers in hexadecimal. The **-g group** option specifies the group name where the name should be described in the configuration definition file. If omitted, "VG" is used for the group name.
	
	-find sync [MU#] [–g <group>]: This option reads the search conditions ($Physical, $Volume, $LETALL) from the standard input, searches for a logical drive that corresponds to the group (<group>) defined in the configuration definition file, and then sends the unwritten data on the file system buffer to the logical drive (physical disk) to synchronize the pair (see Figure 4.34). The **–g <group>** option specifies all the group names on the configuration definition file. If omitted, the –find sync [MU#] searches for a logical drive that corresponds to all groups in the CCI local instance, and sends the unwritten data on the file system buffer to the logical drive (physical disk) to synchronize the pair (see Figure 4.35).
	Notes:
	- Windows NT system does not support LDM volumes. Use $LETALL instead of $Volume.
	`raidscan –pi $LETALL –find sync –g ORA`
	`[SYNC] : ORA ORA_000[-] -> F:\Dsk1\p1 : F:`

Table 4.15 Raidscan Command Parameters (continued)

Parameter	Value
Options (Continued)	- The sync option executes the following procedures depending on the execution condition. If the logical device that corresponds to the <group> described in the configuration definition file is closed from the application, the sync option flushes the system buffer and changes the logical device to "Dismount" status. If the logical device is opened from the application, the sync option flushes the system buffer only. In this case, [FLUSH] will be displayed. [FLUSH] : ORA ORA_000[-1] -> \Vol44\Dsk0 : Volume{56e4954a-28d5-4824-a408-3ff9a6521e5d} **Restriction**: - All logical drives that corresponds to the group defined in the configuration definition file must be closed from the application. - This option cannot specify the device object name (shown below) for the search conditions: D:\Vol(Dms,Dmt, Dmr)X\DskY, \Vol(Dms,Dmt, Dmr)X\DskY

```
# raidscan  -p  cl1-a
PORT# /ALPA/C,TID#,LU#.Num(LDEV#....)...P/S, Status,LDEV#,P-Seq#,P-LDEV#
CL1-A / ef/ 5, 1,  0-0.1(0).............P-VOL ----  ----,  -----  ----
CL1-A / ef/ 5, 1,  1-0.1(1).............SMPL  ----  ----,  -----  ----
```

Figure 4.29 Raidscan Command for Fibre-Channel Ports

```
# ioscan –fun | grep rdsk | raidscan –find
DEVICE_FILE     UID  S/F PORT   TARG  LUN    SERIAL  LDEV  PRODUCT_ID
/dev/rdsk/c0t0d4  0   F  CL1-A     0    4     3116      4  DF600F-CM
/dev/rdsk/c0t0d2  0   F  CL1-A     0    2     3116      2  DF600F
/dev/rdsk/c1t0d3  -   -  CL1-A     -    -     3117      -  DF600F
```

Figure 4.30 Raidscan Command -find Option

The output of the raidscan command with the –find option includes:

- **UID**: Shows the unit ID for multiple subsystem configurations. If UID is displayed as '-', the command device for HORCM_CMD is not found.
- **S/F**: Shows whether the PORT is SCSI or fibre
- **PORT**: Shows the 9500V port number
- **TARG**: Shows the target ID (which was converted by the fibre conversion table)
- **LUN**: Shows the LUN (which was converted by the fibre conversion table)
- **SERIAL**: Shows the production (serial#) number (Equipment ID) of the 9500V subsystem
- **LDEV**: Shows the LDEV# within the 9500V subsystem. LDEV indicates LU for ShadowImage 9500V operations
- **PRODUCT_ID**: Shows the product-id field in the SCSI inquiry page

```
# iocsan –fun | grep rdsk | raidscan –find
DEVICE_FILE      Group  PairVol   PORT     TARG  LUN M  SERIAL  LDEV
/dev/rdsk/c0t3d0  oradb  oradev1  CL1-A     3      0 -  3116    17
/dev/rdsk/c0t3d0  oradb  oradev1  CL1-A     3      0 0  3116    17
```

Figure 4.31 Raidscan Command –find inst Option

Note: If multiple device files were shared (linked) within the same LDEV, the first one founded would be registered as the device file name.

The output of the raidscan command with –find inst option includes:

- **Group**: Shows the group name (dev_group) described in the configuration definition file. If "-" is displayed, the device that corresponds to this group does not exist in the configuration definition file
- **PairVol**: Shows the dev_name described in the configuration definition file
- **Port**: Shows the port name described in the configuration definition file
- **TARG**: Shows the target ID described in the configuration definition file
- **LUN**: Shows the LU number described in the configuration definition file
- **M**: Shows the MU# described in the configuration definition file
- **SERIAL**: Shows the production (serial#) number (Equipment ID) of the 9500V subsystem
- **LDEV**: Shows the LDEV# within the 9500V subsystem. LDEV indicates LU for Synchronous TrueCopy 9500V/ShadowImage 9500V operations.

```
# iocsan –fun | grep rdsk | raidscan –find verify
DEVICE_FILE      Group  PairVol   PORT     TARG  LUN M  SERIAL  LDEV
/dev/rdsk/c0t3d0  oradb  oradev1  CL1-A     3      0 0  3501    17
/dev/rdsk/c0t3d1  oradb  oradev2  CL1-A     3      1 0  3501    18
/dev/rdsk/c0t3d2  -      -        -         -      - 0  3501    19

# iocsan –fun | grep rdsk | raidscan –find verify 1 –fd
DEVICE_FILE       Group  PairVol   Device_File   M  SERIAL  LDEV
/dev/rdsk/c0t3d0  oradb  oradev1   C0t3d0        1  3501    17
/dev/rdsk/c0t3d1  oradb  oradev2   Unknown       1  3501    18
/dev/rdsk/c0t3d2  -      -         -             1  3501    19
```

Figure 4.32 Raidscan Command –find verify Option

Note: If the contents displayed in DEVICE_FILE and Device_File is different, it indicates that the volumes are shared (linked) within the same LDEV. If "Unknown" is displayed in Device_File, this volume is not registered, so the pair operation (except for local options) is rejected in the protection mode.

The output of the raidscan command with **-find** verify option includes:

- **Group**: Shows the group name (dev_group) described in the configuration definition file. If "-" is displayed, the device that corresponds to this group does not exist in the configuration definition file
- **PairVol**: Shows the dev_name described in the configuration definition file
- **Device_File**: Shows the device file registered in HORCM
- **Port**: Shows the port name described in the configuration definition file
- **TARG**: Shows the target ID described in the configuration definition file
- **LUN**: Shows the LU number described in the configuration definition file
- **M**: Shows the MU# described in the configuration definition file
- **SERIAL**: Shows the production (serial#) number (Equipment ID) of the 9500V subsystem
- **LDEV**: Shows the LDEV# within the 9500V subsystem. LDEV indicates LU for Synchronous TrueCopy 9500V/ShadowImage 9500V operations.

```
# cat /etc/hormperm.conf | raidscan –find conf 0 –g ORA
HORCM_DEV
#dev_group         dev_name  port#    TargetID      LU#        MU#
#/dev/rdsk/c23t0d0    SER = 6145  LDEV =   2 [FIBRE FCTBL = 4]
ORA                 ORA_000       CL2-A        0     0      0
#/dev/rdsk/c23t0d1    SER = 6145  LDEV =   3 [FIBRE FCTBL = 4]
ORA                 ORA_001       CL2-A        0     1      0
#/dev/rdsk/c23t0d2    SER = 6145  LDEV =   4 [FIBRE FCTBL = 4]
ORA                 ORA_002       CL2-A        0     2      0
#/dev/rdsk/c23t0d3    SER = 6145  LDEV =   5 [FIBRE FCTBL = 4]
ORA                 ORA_003       CL2-A        0     3      0
#ERROR  [CMDDEV]  /dev/rdsk/c23t0d7  SER = 6145  LDV = 9 [DF600F-CM]
```

Figure 4.33 Raidscan Command –find conf Option

Note:

- If a command device is included in the STDIN device, a comment is displayed as follows and the target device will not be included.

```
#ERROR  [CMDDEV]  /dev/rdsk/c23t0d7  SER = 6145  LDV = 9 [DF600F-CM]
```

- If an STDIN device is shared by multiple device files and is displayed as target device, a comment is displayed as follows and the target device will not be included.

```
#ERROR [LDEV LINK] /dev/rdsk/c24t0d3 SER = 6145 LDV = 5 [FIBRE FCTBL = 4]
```

- If the STDIN device does not have a proper MU#, a comment is displayed as follows and the target device will not be included.

```
#ERROR [INVALID MUN(2<1)] /dev/rdsk/c24t0d3 SER = 6145 LDV = 5 [DF600F  ]
```

- If the STDIN device is mixed between the subsystems that have different mirroring control, a comment is displayed as follows and the target device will be not be included.

```
#ERROR [MIXING RAID TYPE] /dev/rdsk/c24t0d3 SER = 6145 LDV = 5 [DF600F  ]
```

The following is an example of synchronizing (flushing) the system buffer that corresponds to group ORB in the configuration definition file.

```
raidscan -pi $Volume -find sync -g ORB
[SYNC] : ORB ORB_000[-] -> \Dmt1\Dsk1 : Volume{bf48a395-0ef6-11d5-8d69-00c00d003b1e}
[SYNC] : ORB ORB_001[-] -> \Dmt1\Dsk2 : Volume{bf48a395-0ef6-11d5-8d69-00c00d003b1e}
[SYNC] : ORB ORB_002[-] -> \Dmt1\Dsk3 : Volume{bf48a395-0ef6-11d5-8d69-00c00d003b1e}
```

Figure 4.34 Raidscan Command –find sync Option Example (1)

The following is an example of synchronizing (flushing) the system buffer that corresponds to all groups in the CCI local instance.

```
[SYNC] : ORA ORA_000[-] -> \Vol44\Dsk0: Volume{56e4954a-28d5-4824-a408-3ff9a6521e5d}
[SYNC] : ORA ORA_000[-] -> \Vol45\Dsk0: Volume{56e4954a-28d5-4824-a408-3ff9a6521e5e}
[SYNC] : ORB ORB_000[-] -> \Dmt1\Dsk1 : Volume{bf48a395-0ef6-11d5-8d69-00c00d003b1e}
[SYNC] : ORB ORB_001[-] -> \Dmt1\Dsk2 : Volume{bf48a395-0ef6-11d5-8d69-00c00d003b1e}
[SYNC] : ORB ORB_002[-] -> \Dmt1\Dsk3 : Volume{bf48a395-0ef6-11d5-8d69-00c00d003b1e}
```

Figure 4.35 Raidscan Command –find sync Option Example (2)

4.12.2 Raidar Command

The **raidar** command displays configuration, status, and I/O activity information for the specified 9500V port/TID(s)/device(s) at the specified time interval. The configuration information is acquired directly from the 9500V subsystem (not from the configuration definition file). Table 4.16 lists and describes the raidar command parameters. Figure 4.36 shows an example of the raidar command and its output.

Note:

- The I/O activity of a ShadowImage 9500V S-VOL in the COPY or PAIR state includes only host-requested I/Os. The I/O activity of a P-VOL or simplex volume includes only host-requested I/Os. If state changed into SMPL in S-VOL (COPY, PAIR) I/O actively, and then I/O activity of the between is reported in the SMPL state.
- The subsystem has I/O activity information to be displayed for each controller. Therefore, if you want to display the I/O information using this command, when creating the configuration definition file, you must set the command device that is specified by the controller that accepts the host I/O.

Table 4.16 Raidar Command Parameters

Parameter	Value
Command Name	raidar
Format	raidar { -h | -q | -z | -p <port> <targ> <lun> <mun> | -pd <raw_device | -s [interval] [count] }
Options	**-h**: Displays Help/Usage and version information. **-q**: Terminates the interactive mode and exits the command. **-z** or **-zx**: Makes the raidar command enter the interactive mode. The **-zx** option guards performing of the HORCM in the interactive mode. When this option detects a HORCM shut down, interactive mode terminates. **-p <port> <targ> <lun> <mun>....**: Monitors one or more (up to 16) devices at a time. **<port>**: Specifies the port to be reported: CL1-A to CL1-B and CL2-A to CL2-B. This option must be specified. **<targ>**: Specifies the SCSI/Fibre target ID (0 to 15) of the specified port (see Appendix D for fibre-to-SCSI address conversion information). **<lun>**: Specifies the LUN (0 to 7) on the specified TID. **<mun>**: Specifies the MU number of the specified LUN. **pd <raw_device>**: Allows designation of an LDEV by raw device file name. **-s [interval]** or **-sm [interval]**: Designates the time interval in seconds. **-s**: Interprets the time interval as seconds. **-sm**: Interprets the time interval as minutes. **[interval]**: Designates the time interval value (1 to 60). If the interval is not specified, the default interval (3) is used. **[count]**: Designates number of repeats. When omitted, this command repeats until CNTL-C.

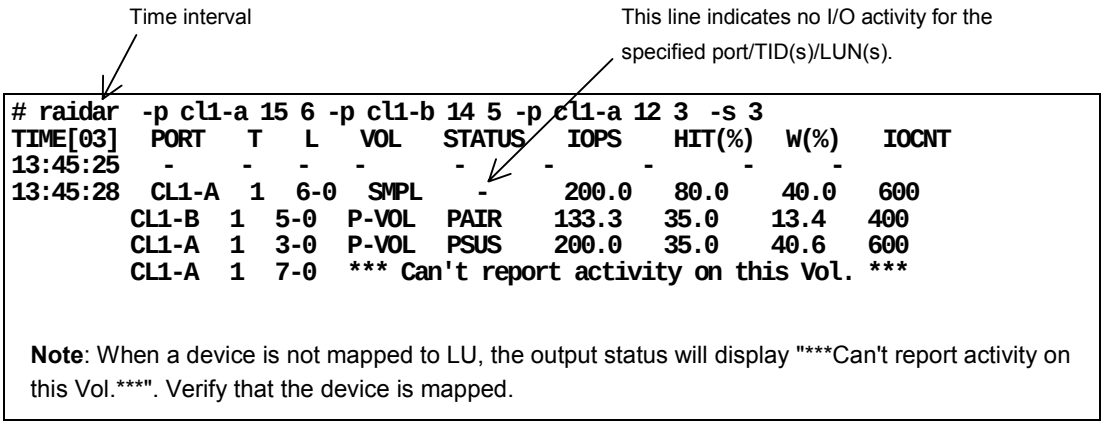

Figure 4.36 Raidar Command Example

The output of the raidar command includes:

- **TIME [ ]:** Shows the interval time
- **PORT:** Shows the port name of the 9500
- **T:** Shows the port ID
- **L:** Shows the LU number in the target ID of the 9500
- **VOL:** Shows the volume attribute (P-VOL, S-VOL, SMPL)
- **STATUS:** Shows the pair status of the paired volume
- **IOPS:** Shows the number of I/Os (read/write) per second (total I/O rate)
- **HIT(%):** Shows the hit rate for read I/Os (read hit rate)
- **W(%):** Shows the ratio of write I/Os to total I/Os (percent writes)
- **IOCNT:** Shows the number of times of write and read

4.12.3 Raidqry Command

The raidqry command (RAID query) displays the configuration of the connected host and 9500V subsystem. Table 4.17 lists and describes the raidqry command parameters. Figure 4.37 shows an example of the raidqry command output.

Table 4.17 Raidqry Command Parameters

Parameter	Value
Command Name	raidqry
Format	raidqry { -h | -q | -z | -I | -r <group> | [-f] }
Options	**-h**: Displays Help/Usage and version information. **-q**: Terminates the interactive mode and exits the command. **-z** or **-zx**: Makes the raidqry command enter the interactive mode. The **-zx** option guards performing of the HORCM in the interactive mode. When this option detects a HORCM shut down, interactive mode terminates. **-I**: Displays the configuration information for the local host and the local 9500V. **-r <group>**: Displays the configuration information for the remote host and the remote 9500V that contains the specified group. **-f**: Displays the hostname (ip_address) as specified in the configuration definition file. Use this option if "floatable IP address" is used for the hostname (ip_address) in the configuration file.

```
# raidqry –l
No Group    Hostname         HORCM_ver    Uid  Serial#    Micro_ver    Cache(MB)
 1  ---     HOSTA          01-10-03/02     0     3005   06-51-00/00      1024

# raidqry -r oradb
No Group    Hostname         HORCM_ver    Uid  Serial#    Micro_ver    Cache(MB)
 1 oradb    HOSTB          01-10-03/02     0     3005   06-51-00/00      1024

# raidqry -l  -f
No Group    Floatable Host   HORCM_ver    Uid  Serial#    Micro_ver    Cache(MB)
 1  ---    xxx.xxx.xxx.xxx  01-10-03/02     0     3005   06-51-00/00      1024
```

Figure 4.37 Raidqry Command

The output of the raidqry command includes:

- **No**: Shows the order when the group name (dev_group) described in the configuration definition file has multiple remote hosts
- **Group**: When the **-r** option is used, this column shows the group name (dev_group) described in the configuration definition file.
- **Hostname**: When the **-l** option is used, this column shows the host name of the local host. When the **-r** option is used, this column shows the remote host name for the group (dev_group) which is described in the configuration definition file. The host name with more than or equal to 30 characters cannot be displayed.
- **Floatable Host**: When the **-f** option is used, this column shows the host name (ip_address) described in the configuration definition file. Up to 30 host names can be displayed. The **-f** option interprets the host name as utilizing floatable IP for a host.
- **HORCM_ver**: Shows the version of the HORC Manager on the local or remote host. The **-l** option specifies local host. The **-r** option specifies remote host.
- **Uid Serial# Micro_ver**: Shows the unitID, serial number (Equipment ID), and microcode version of the 9500V connected to the local or remote host. The **-l** option specifies local host. The **-r** option specifies remote host.
- **Cache(MB)**: Shows the logical cache capacity (in MB) of the 9500V connected to the local or remote host. The **-l** option specifies local host. The **-r** option specifies remote host.

4.13 Controlling CCI Activity

4.13.1 Horcmstart Command

The **horcmstart** command is a shell script that starts the HORCM application (/etc/horcmgr). This shell script also sets the environment variables for HORCM as needed (e.g., HORCM_CONF, HORCM_LOG, HORCM_LOGS). Table 4.18 lists and describes the horcmstart command parameters.

Table 4.18 Horcmstart Command Parameters

Parameter	Value
Command Name	horcmstart
Format	horcmstart.sh { inst ... } horcmstart.exe { inst ... }
Options	**Inst:** Specifies the HORCM instance number (numerical value). When this option is specified, the horcmstart shell script sets the environment variables (HORCMINST, HORCM_CONF, HORCM_LOG, HORCM_LOGS) that correspond to the instance number, and starts the specified HORCM instance. (Environment variables set by the user become invalid.) When this option is not specified, the horcmstart shell script starts 1 HORCM and uses the environment variables set by the user. If you have designated full environment variables, you should use horcmstart.sh without any arguments. If you did not designate environment variables (HORCM_CONF, HORCM_LOG, HORCM_LOGS), then this shell script sets the environment variables as follows: For UNIX®-based platforms: If HORCMINST is specified: HORCM_CONF = /etc/horcm*.conf (* is instance number) HORCM_LOG = /HORCM/log*/curlog HORCM_LOGS = /HORCM/log*/tmplog If no HORCMINST is specified: HORCM_CONF = /etc/horcm.conf HORCM_LOG = /HORCM/log/curlog HORCM_LOGS = /HORCM/log/tmplog For Windows® NT/2000 platform: If HORCMINST is specified: HORCM_CONF = \WINNT\horcm*.conf (* is instance number) HORCM_LOG = \HORCM\log*\curlog HORCM_LOGS = \HORCM\log*\tmplog If no HORCMINST is specified: HORCM_CONF = \WINNT\horcm.conf HORCM_LOG = \HORCM\log\curlog HORCM_LOGS = \HORCM\log\tmplog ***Note:*** The **HORCM_LOGS** environment variable is used to specify the log file directory for automatic storing. When HORCM starts up, the log files created in the operation are stored automatically in the HORCM_LOGS directory. This log directory must give an equality class with HORCM_LOG.

4.13.2 Horcmshutdown Command

The **horcmshutdown** command is a shell script for stopping the HORCM application (/etc/horcmgr). Table 4.19 describes the shutdown command parameters.

Table 4.19 Horcmshutdown Command Parameters

Parameter	Value
Command Name	horcmshutdown
Format	horcmshutdown.sh {inst...} horcmshutdown.exe {inst...}
Option	**Inst**: Specifies the HORCM (CCI) instance number (numerical value). When this option is specified, the command stops the specified HORCM instance. When this option is not specified, the command refers to the instance (environment variable HORCMINST) of the execution environment of this shell script and stops the following the HORCM instance. When HORCMINST is specified, this command stops the HORCM instance of the execution environment of this shell script. When HORCMINST is not specified, this command stops the HORCM having no instance setting.

4.13.3 Horcctl Command

The HORCM software has logs that identify the cause of software and/or hardware errors as well as a tracing function for investigating such errors. The location of the log files depends on the user's command execution environment and the HORC Manager's execution environment. The command trace file and core file reside together under the directory specified in the HORC Manager's execution environment. See Appendix B for log file and log directory information.

The **horcctl** command can be used for both maintenance and troubleshooting. The horcctl command allows you to change and display the internal trace control parameters (e.g., level, type, buffer size) of the HORC Manager commands. If a new value for a parameter is not specified, the current trace control parameter is displayed. Table 4.20 lists and describes the horcctl command parameters.

Caution: Do not change the trace level unless directed to do so by a Hitachi Data Systems representative. Level 4 is the normal trace level setting. Levels 0-3 are for troubleshooting. Setting a trace level other than 4 may impact problem resolution. If you request a change of the trace level using the **horcctl -l <level>** command, a warning message is displayed, and this command enters interactive mode.

Table 4.20 Horcctl Command Parameters

Parameter	Value
Command Name	horcctl
Format	horcctl { -h | -q | -z | -d | -c | -l <level> | -d <y/n> | -s <size(KB)> | -t <type> | -S | -D | -C | [-u <-unitid> | -ND | -NC | -g <group>}
Options	**-h**: Displays Help/Usage and version information. **-q**: Terminates the interactive mode and exits the command. **-z** or **-zx**: Makes the horcctl command enter the interactive mode. The **-zx** option guards performing of the HORCM in the interactive mode. When this option detects a HORCM shut down, interactive mode terminates. **-d**: Interprets the control options following this option (-l <level>, -b <y/n>, -s <size(KB)>, and -t <type>) as the parameters of the CCI commands. **-c**: Interprets the control options following this option (-l <level>, -b <y/n> and -t <type>) as the parameters of the HORC Manager (HORCM). **-l <level>**: Sets the trace level (range = 0 to 15). If a negative value is specified, the trace mode is canceled. A negative value "n" must be specified as "--n". ***Caution:*** Do not change the trace level unless directed to do so by a Hitachi Data Systems representative. Level 4 is the normal trace level setting. Levels 0-3 are for troubleshooting. Setting a trace level other than 4 may impact problem resolution. If you request a change of the trace level using the **horcctl -l <level>** command, a warning message is displayed, and this command enters interactive mode. **-b <y/n>**: Sets the trace writing mode: Y = buffer mode, N = synchronous mode. **-t <type>**: Sets the trace type (range = 0 to 511). When this option is used, only traces of the specified type are output. One or more values can be specified. **-s <size(KB)>**: Changes the default trace buffer size, which is 1 MB, in units of 1024 bytes. **-S**: Shuts down HORCM. **-D**: Displays the command device name currently used by HORCM. If the command device is blocked due to online maintenance (microcode replacement) of the 9500, you can check the command device name in advance using this option. **-C**: Changes the command device name being used by HORCM and displays the new command device name. If the command device is blocked due to online maintenance (microcode replacement) of the 9500V, you can change the command device in advance using this option. **-u <unitid>**: Used to specify the unit ID of a command device as the target. This option is effective when the **-D** or **-C** option is specified. If this option is not specified, the unit ID is 0. **-ND -g <group>**: Displays the network address and port name being used by HORCM. The **-g <group>** option is used to specify the group name defined in the configuration definition file. **-NC -g <group>**: Changes the network address and port name being used by HORCM and displays the new network address name. The **-g <group>** option specifies the group name defined in the configuration definition file.

The following is an example of changing the trace level to 15.

```
C:\HORCM\etc>horcctl -d -l 15
/****************************** WARNING ***********************************/
/*    This is an option for maintenance, and used for troubleshootings.       */
/*    When it is issued, the internal trace control parameters of the HORC  */
/*    manager and HORC commands are changed and displayed.                  */
/*    These trace control parameters should not be changed unless directed  */
/*    by a CS&S service representative.                                     */
/*    For cancel   -> Enter '-q' option                                     */
/*    For continue -> Re-enter '-c ...' or '-d ...'options                  */
/****************************** WARNING ***********************************/
horcctl[HOMRCF]: -q
C:\HORCM\etc>
```

Figure 4.38 Horcctl command –d –l Option

```
C:\HORCM\etc>horcctl –d
logdir = C:\HORCM\log1\curlog
[Client]:
trace = ON
level = 4
mode  = Buffer
size  = 1024 KB
type  =  0, 1, 2, 3, 4, 5, 6, 7, 8, 9,10,11,12,13,14,15,16,17,18,19...Full
```

Figure 4.39 Horcctl command –d Option

```
C:\HORCM\etc>horcctl –c
logdir = C:\HORCM\log1\curlog
[HORCM]:
trace = ON
level = 4
mode  = Buffer
size  = 1024 KB
type  =  0, 1, 2, 3, 4, 5, 6, 7, 8, 9,10,11,12,13,14,15,16,17,18,19...Full
```

Figure 4.40 Horcctl command –c Option

4.14 Windows NT® and Windows® 2000 Subcommands

The CCI software provides subcommands for the Windows NT® and Windows® 2000 platforms which are executed as options (-x <command> <arg>) of another command. When you specify a subcommand as the only option of a command, you do not need to start HORCM. If another option of the command and the subcommand is specified on the same command line, place the other option after the subcommand.

4.14.1 Findcmddev Subcommand

The **findcmddev** subcommand (find command device) searches for command devices within the specified range of disk drive numbers. If it is found, the command device is displayed in the same format as in the configuration definition file. This subcommand is used when the command device name is not known. Table 4.21 lists and describes the findcmddev subcommand parameters. Figure 4.41 shows an example of the findcmddev subcommand used as an option of the raidscan command output.

Caution: The findcmddev subcommand must be used when HORCM is not running.

Note: The findcmddev subcommand searches for the physical and logical devices associated with the command device. If the command device is indicated as a logical device, you must delete the drive letter assigned to the command device to prevent utilization by general users. The physical drive number may change at every reboot. If the number changes, use Volume{guid} for which the same name is kept. Volume{guid} is created when you make a partition by using the Windows' Disk Management. Do not format.

Table 4.21 Findcmddev Subcommand Parameters

Parameter	Value
Command Name	findcmddev
Format	-x findcmddev drive#(0-N)
Argument	**drive#(0-N)**: Specifies the range of disk drive numbers on the Windows® NT/2000 system.

```
raidscan  -x  findcmddev   hdisk0, 20
cmddev of ser#   3001 = \\.\PhysicalDrive0
cmddev of ser#   3001 = \\.\E:
cmddev of ser#   3001 = \\.\Volume{b9b31c79-240a-11d5-a37f-00c00d003b1e}
```

Note: This example searches for command devices in the range of disk drive numbers 0 to 20.

Figure 4.41 Findcmddev Subcommand

4.14.2 Drivescan Subcommand

The **drivescan** subcommand displays the relationship between the disk numbers assigned by the Windows NT®/2000 system and the LDEVs on the 9500V, and also displays attribute and status information for each LDEV. Table 4.22 lists and describes the drivescan subcommand parameters. Figure 4.42 shows an example of the drivescan subcommand used as an option of the raidscan command and its output.

Table 4.22 Drivescan Subcommand Parameters

Parameter	Value
Command Name	drivescan
Format	-x drivescan drive#(0-N)
Argument	**drive#(0-N)**: Specifies the range of disk drive numbers on the Windows® NT/2000 system.

```
raidscan -x drivescan  harddisk0,20
Harddisk 0... Port[ 1] PhId[ 0] TId[  0] Lun[  0]  [HITACHI] [DK328H-43WS]
Harddisk 1... Port[ 2] PhId[ 4] TId[ 29] Lun[  0] [HITACHI] [DF600F]
              Port[CL1-A] Ser#[   3005] LDEV#[   9(0x009)]
              HORC = NONE  HOMRCF[MU#0 = SMPL MU#1 = NONE MU#2 = NONE]
              RAID5[Group  2- 1] SSID = 0x0008 CTGID = 3
Harddisk 2... Port[ 2] PhId[ 4] TId[ 29] Lun[  1] [HITACHI] [DF600F]
              Port[CL1-A] Ser#[   3005] LDEV#[  10(0x00A)]
              HORC = NONE  HOMRCF[MU#0 = SMPL MU#1 = NONE MU#2 = NONE]
              RAID5[Group  2- 1] SSID = 0x0004 CTGID = 3
Harddisk 3... Port[ 2] PhId[ 4] TId[ 29] Lun[  6] [HITACHI] [DF600F-CM]
              Port[CL1-A] Ser#[   3005] LDEV#[  15(0x00F)]
```

Note: This example displays the devices for the range of disk drive numbers from 0 to 20.

Figure 4.42 Drivescan Subcommand Example

The output of the drivescan subcommand includes:

- **Harddisk #**: Shows the hard disk recognized by the Windows® NT/2000 system
- **Port**: Shows the port number on the device adapter recognized by the NT/2000 system
- **PhId**: Shows the bus number on the device adapter port recognized by NT/2000 system
- **TId**: Shows the target ID of the hard disk(s) on the specified port and bus. For further information on fibre-to-SCSI address conversion, see Appendix D.
- **LUN**: Shows the LU number of the hard disk on the specified port, bus, and TID
- **Port[CLX-Y]**: Shows the port number on the 9500V subsystem
- **Ser#**: Shows the production number (Equipment ID) of the 9500V subsystem
- **LDEV#**: Shows the LDEV ID (hexadecimal) of the specified volume on the 9500. LDEV indicates LU for Synchronous TrueCopy 9500V/ShadowImage 9500V operations.
- **HORC**: Shows the Synchronous TrueCopy 9500V attribute (P-VOL, S-VOL, SMPL)
- **HOMRCF**: Shows the ShadowImage 9500V attribute (P-VOL, S-VOL, SMPL) and MU number (0-2) of the specified volume

- **RAIDX[Group]**: Shows the physical location (frame number-parity group number) of the specified volume and the RAID level of this parity group
- **SSID**: Shows the SSID of the specified volume

4.14.3 Portscan Subcommand

The **portscan** subcommand displays the devices on the specified port(s). Table 4.23 lists and describes the portscan subcommand parameters. Figure 4.43 shows an example of the portscan subcommand used as an option of the raidscan command and its output.

Table 4.23 Portscan Subcommand Parameters

Parameter	Value
Command Name	portscan
Format	-x portscan port#(0-N)
Argument	**port#(0-N)**: Specifies the range of port numbers on the Windows® NT/2000 system

```
raidscan  -x  portscan  port0,20
PORT[ 0] IID [ 7] SCSI Devices
         PhId[ 0] TId[  3] Lun[  0]  [MATSHIT] [CD-ROM CR-508  ] ...Claimed
         PhId[ 0] TId[  4] Lun[  0]  [HP     ] [C1537A         ] ...Claimed
PORT[ 1] IID [ 7] SCSI Devices
         PhId[ 0] TId[  0] Lun[  0]  [HITACHI ] [DK328H-43WS   ] ...Claimed
PORT[ 2] IID [ 7] SCSI Devices
         PhId[ 0] TId[  5] Lun[  0]  [HITACHI ] [DF600F        ] ...Claimed
         PhId[ 0] TId[  5] Lun[  1]  [HITACHI ] [DF600F        ] ...Claimed
         PhId[ 0] TId[  5] Lun[  2]  [HITACHI ] [DF600F        ] ...Claimed
         PhId[ 0] TId[  6] Lun[  0]  [HITACHI ] [DF600F        ] ...Claimed
```

Note: This example displays the devices for the range of ports from 0 to 20.

Figure 4.43 Portscan Subcommand

The output of the portscan subcommand includes:

- **Port**: Shows the port number on the device adapter recognized by the Windows® system
- **IID**: Shows the initiator ID on the specified device adapter port
- **PhId**: Shows the BUS number on the specified device adapter port
- **TId**: Shows the target ID of the hard disk(s) on the specified adapter port and bus. For further information on Fibre-to-SCSI address conversion, see Appendix D.
- **LUN**: Shows the LU number of each hard disk on the specified device adapter port/bus. This item shows LDEV# of the partner who becomes a pair in or among the 9500V.

4.14.4 Sync Subcommand

The **sync** (synchronization) subcommand sends unwritten data remaining on the Windows NT®/2000 server to the specified device(s) to synchronize the pair(s) before the CCI command is executed. Table 4.24 lists and describes the sync subcommand parameters.

Caution: The logical and physical devices to be synchronized must be offline to all other applications. The Sync does not propagate to a specified drive which has a directory mount on the Windows® 2000 system.

Table 4.24 Sync Subcommand Parameters (Continued on next page)

Parameter	Value
Command Name	Sync
Format	-x sync A: B: C: ... -x sync all -x sync drive#(0-N) ... -x sync Volume#(0-N) ... (Windows® 2000 systems only) -x sync D:\directory or \directory pattern ... (Windows® 2000 systems only)

Argument	**A: B: C:[\directory or \directory pattern] ...** : Specifies the logical devices that you want to synchronize. The data is flushed to the specified logical device and to the physical device that corresponds to the specified logical device.
	
	If a directory-mounted volume exists in the specified logical device, the data will be flushed including the directory-mounted volume as follows:
	`pairsplit –x sync D:`
	`[SYNC] D: HarddiskVolume2`
	`[SYNC] D: \hd1 HarddiskVolume8`
	`[SYNC] D: \hd2 HarddiskVolume9`
	
	[\directory or \directory pattern] (Windows 2000 systems only)
	Specifies the directory or the directory pattern for searching the directory mount point in the logical device.
	If directory is specified: The applicable directory-mounted volume will be flushed.
	`pairsplit –x sync D:\hd1`
	`[SYNC] D:\hd1 HarddiskVolume8`
	
	If directory patter is specified: The directory-mounted volume that matches the specified pattern will be flushed.
	`pairsplit –x sync D:\h`
	`[SYNC] D:\hd1 HarddiskVolume8`
	`[SYNC] D:\hd2 HarddiskVolume9`
	
	all: Synchronizes all logical devices. The logical device on which the CCI software is installed and the logical device containing the Windows® NT/2000 directory are excluded.
	If a directory-mounted volume exists in the specified logical device, the data will be flushed including the directory-mounted volume as follows:
	`pairsplit –x sync all`
	`[SYNC] C: HarddiskVolume1`
	`[SYNC] D: \hd1 HarddiskVolume8`
	`[SYNC] D: \hd2 HarddiskVolume9`
	`[SYNC] G: HarddiskVolume10`
	
	drive#(0-N) ...: Specifies the range of devices on the Windows® NT/2000 system.

Table 4.24 Sync Subcommand Parameters (continued)

Parameter	Value
Argument	**Volume#(0-N) ...**: For Windows® 2000 systems only. Synchronizes the data to the specified LDM (logical device manager) volume. The LDM volume specifies the following device objects: \Vol#, \Dms#, \Dmt#, or \Dmr#. For Windows® 2000 LDM volumes, see Appendix E.

Note	The sync command executes the following procedures depending on the execution condition.
	
	If the logical device is closed from the application, the system buffer is flushed and changes the logical device to "Dismount" status.
	If the logical device is opened from the application, the sync option flushes the system buffer only. In this case, [WARNING] will be displayed.
	`pairsplit –x sync –C:`
	`WARNING: Only flushed to [\\.\C] drive due to be opening.`
	`[SYNC] C: HarddiskVolume3`

The following examples show the sync subcommand used as an option of the pairsplit command. For the example in Figure 4.44, the data remaining on logical devices C: and D: is written to disk, all pairs in the specified group are split (status = PSUS), and read/write access is enabled for all S-VOLs in the specified group.

```
pairsplit -x sync  C: D:  -g  oradb
```

Figure 4.44 Sync Subcommand Example 1

For the example in Figure 4.45, the data remaining on physical devices harddisk2 and harddisk3 is written to disk, all pairs in the specified group are deleted (status = SMPL), which enables read/write access for all secondary volumes.

```
pairsplit -x sync  hdisk2  hdisk3 -g oradb -S
```

Figure 4.45 Sync Subcommand Example 2

4.14.5 Mount Subcommand

The **mount** subcommand mounts the specified drive to the specified partition on the specified hard disk drive. If the mount subcommand is executed without specifying an argument, all drives that are currently mounted are displayed. Table 4.25 lists and describes the mount subcommand parameters. Figure 4.46 and Figure 4.47 show examples of the mount subcommand used as an option of the pairsplit command output.

Caution: The partition on the specified disk drive (hard disk) must be recognized on the Windows NT®/2000 system.

Table 4.25 Mount Subcommand Parameters

Parameter	Value
Command Name	mount
Format	-x mount -x mount drive: hdisk# partition# ... (for Windows NT) -x mount drive: Volume#(0-N) ... (for Windows 2000) -x mount drive: [\directory] Volume#(0-N) ... (for Windows 2000)

Arguments	**drive: hdisk# partition#**: For Windows NT. The drive specifies the logical drive to be mounted. The hdisk# specifies the hard disk drive (number) to be mounted. The partition # specifies the partition number to be mounted. If the partition# is not specified, the drive is mounted as HarddiskVolume# for Windows® 2000 systems. If this command is executed without any argument, the device that is already mounted will be displayed.
	
	drive: [\directory] Volume#: For Windows 2000. The drive specifies the logical device to be mounted. The hdisk# specifies the following device object names of the LDM volume to be mounted: hdisk#, \Vol#, \Dms#, \Dmt#, or \Dmr#.
	[\directory]
	Specifies the directory for specifying the directory mount point in the logical device.
	`pairsplit –x mount D:\hd1 \Vol8`
	`D:\hd1 <+> HarddiskVolume8`
	`pairsplit –x mount D:\hd2 \Vol9`
	`D:\hd2 <+> HarddiskVolume9`
	If this command is executed without any argument, the device that is already mounted will be displayed. For Windows® 2000 LDM volumes, see Appendix E.

If the command is executed without an argument, a mounted device that contains the directory-mounted volume will be displayed. If the mounted volume is an LDM volume, the physical drive (hard disk) number that configures the LDM volume will be displayed also.

```
pairsplit -x mount F: hdisk2 p1   -x mount G: hdisk1 p1
pairsplit -x mount

Drive   FS_name  VOL_name   Device      Partition   ...  Port  PathID  Targ  Lun
C:      FAT      Null       Harddisk0   Partition1  ...    1      0      0    0
F:      FAT      Null       Harddisk2   Partition1  ...    2      0      5    1
G:      NTFS     Null       Harddisk1   Partition1  ...    2      0      5    0
Z:      CDFS     Null       CdRom0                  ...  Unknown
```

This example mounts the F: drive to partition 1 on disk drive 2, and mounts the G: drive to partition 1 on disk drive 1.

Figure 4.46 Mount Subcommand (Windows NT®)

```
pairsplit -x mount  F: hdisk2
pairsplit  -x mount
Drive  FS_name  VOL_name   Device          Partition  ...   Port  PathID  Targ  Lun
C:     NTFS     Null       Harddiskvolume1           ...   Harddisk0
F:     NTFS     Null       Harddiskvolume2           ...   Harddisk1
D:     NTFS     Null       Harddiskvolume3           ...   Harddisk2
D:\hd1 NTFS     Null       Harddiskvolume4           ...   Harddisk3
D:\hd2 NTFS     Null       Harddiskvolume5           ...   Harddisk4
G:     NTFS     Null       HarddiskDmVolumes\        ...   \Volume1 ... Harddisk5[3]
```

This example executes mount from the pairsplit command option, mounting the **F:** drive to the harddiskvolume2, after the mounted devices are displayed.

Figure 4.47 Mount Subcommand (Windows® 2000)

The output of the mount subcommand includes:

- **Drive**: Shows the logical device recognized by the Windows NT®/2000 system
- **FS_name**: Shows the name of the file system formatted on the specified drive
- **VOL_name**: Shows the volume label name for the specified drive
- **Device, Partition**: Shows the device name and partition for the specified drive
- **Port, PathID, Targ, Lun**: Shows the port number, path ID (bus), target ID, and LUN for the specified drive. For further information on Fibre-to-SCSI address conversion, see Appendix D.

Note: For the Windows 2000 system, if you specified both hdisk# and partition# arguments for the mount subcommand, the drive characters will not be displayed when the Disk Management is activated in the Control Panel → Administrative Tools → Computer Management → Storage folder. Even if this sub command is executed, a drive letter may not be assigned. In this case, the target volume may be used. Execute the command after making the target volume to unused status. Alternatively, a drive letter can be assigned by activating the Disk Management in the Control Panel – Administrative Tools – Computer Management – Storage folder.

4.14.6 Umount Subcommand

The **umount** subcommand unmounts the specified logical drive and deletes the drive letter. Before deleting the drive letter, this subcommand executes sync internally for the specified logical drive and flushes unwritten data. Table 4.26 lists and describes the umount subcommand parameters. Figure 4.48 shows an example of the umount subcommand used as an option of the pairsplit command.

Caution: The logical drive to be unmounted and the corresponding physical drive must be closed to all applications.

Table 4.26 Umount Subcommand Parameters

Parameter	Value
Command Name	umount
Format	-x umount drive: -x umount drive:[\directory] ... (for Windows 2000)

Argument	drive: Specifies the mounted logical device
	drive:[\directory]: For Windows 2000. Specify the mounted logical device for **drive**.
	[\directory]: Specifies the directory for specifying the directory mount point in the logical device.
	`pairsplit –x umount D:\hd1`
	`D:\hd1 <-> HarddiskVolume8`
	`pairsplit –x umount D:\hd2`
	`D:\hd2 <-> HarddiskVolume9`

```
pairsplit -x umount F:  -x umount G: -g oradb
pairsplit -x mount

Drive   FS_name   VOL_name   Device       Partition   ...  Port PathID  Targ  Lun
C:      FAT       Null       Harddisk0    Partition1  ...    1    0      0    0
Z:      Unknown   Unknown    CdRom0                   ...  Unknown
```

This example unmounts the F: and G: drives, splits all pairs in the specified group (status = PSUS), enables read/write access to all secondary volumes in the specified group, and then displays all mounted drives.

Figure 4.48 Umount Subcommand

The output of the umount subcommand includes:

- **Drive**: Shows the logical drive recognized by the Windows NT®/2000 system
- **FS_name**: Shows the name of the file system formatted on the specified drive
- **VOL_name**: Shows the volume label name for the specified drive
- **Device, Partition**: Shows the device name and partition for the specified drive
- **Port, PathID, Targ, Lun**: Shows the port number, path ID (bus), target ID, and LUN for the specified drive. For further information on Fibre-to-SCSI address conversion, see Appendix D.

Note: For Windows® 2000 system, Even if this sub command is executed, a drive letter may not be deleted. In this case, the target volume may be used. Execute the command after making the target volume to unused status. Alternatively, a drive letter can be deleted by activating the Disk Management in the Control Panel – Administrative Tools – Computer Management – Storage folder.

4.14.7 Environment Variable Subcommands

If no environment variables are set in the execution environment, the environment variable subcommand sets or cancels an environment variable within the CCI command. The setenv subcommand sets the specified environment variable(s). The usetenv subcommand deletes the specified environment variable(s). The env subcommand command displays the environment variable(s). The sleep subcommand causes CCI to wait for the specified time. Table 4.27 lists and describes the environment variable subcommands and their parameters.

Caution: The environment variables must be set before connecting to HORCM, and must be specified during interactive mode (**-z** option). Changing an environment variable after a CCI command execution error is invalid.

Table 4.27 Environment Variable Subcommand Parameters

Parameter	Value
Command Name	setenv usetenv env sleep
Format	-x setenv vaname value -x usetenv vaname -x env -x sleep time
Argument	**Vaname**: Specifies the environment variable to be set or canceled **Value**: Specifies the value or character string of the environment variable to be set **Time**: Specifies the sleep time in seconds

Figure 4.49 shows an example of the setenv subcommand used as an option of the raidscan command. This example changes from "Synchronous TrueCopy 9500V" to "ShadowImage 9500V" an execution environment of the raidscan command which makes a dialog mode, because of establishing "HORCC_MRCF" as an environment variable.

Important: Always set HORCC_MRCF 1 for the ShadowImage 9500V operation.

```
raidscan[HORC]:  -x setenv  HORCC_MRCF 1
raidscan[MRCF]:

raidscan[MRCF]:  -x usetenv  HORCC_MRCF
raidscan[HORC]:
```

Figure 4.49 Environment Variable Subcommand

4.15 Command Tools

4.15.1 Inqraid Command Tool

The **inqraid** command tool confirms the drive connection between the 9500V and the host system. The inqraid command displays the relation between special file(s) on the HP-UX® system and actual physical drive of the 9500V. Table 4.28 lists and describes the inqraid command and parameters. Figure 4.50 to Figure 4.57 show examples of using inqraid and system commands to display the connection between the STDIN special file and the actual physical drive of 9500V.

Table 4.28 Inqraid Command Parameters (Continued on next page)

Parameter	Value
Command Name	inqraid
Format	/HORCM/usr/bin/inqraid [-h | quit | -inqdump | -fx | -find[c] | <special file> | –CLI | -CLIWP | -CLIWN | sort | -CM] \HORCM\etc\inqraid [-h | quit | -inqdump | -fx | -find[c] | <special file> | –CLI | -CLIWP | -CLIWN | sort | -CM | -gvinf | -svinf]
Options	**-h**: This option displays Help/Usage. **quit**: This option terminates from waiting STDIN and exits this command. **-inqdump**: This option displays information for standard inquiry with Dump Image of hexadecimal. **-fx**: This option displays the LDEV number with hexadecimal. **-find [c]**: This option searches a group on the configuration definition file (local instance) from <special file> of STDIN by using pairdisplay command, and uses the following options of the pairdisplay command to display its state. The **–find** option executes the following command (see Figure 4.58 also). pairdisplay –d <Seq#><LDEV> 0 1 2 -l [-fx] [-CLI] 2>/dev/null The **–find[c]** option executes the following command and then displays the result edited to CLI format (see Figure 4.59 also). pairdisplay –d <Seq#><LDEV><MU#> -fd –CLI 2>/dev/null **<special file>**: This option is used to specify the special file name as argument of command. **If no argument**, this command makes mode that wait for STDIN without argument. **-CLI**: Displays the CLI. This option displays the CLI using one header and divides the column with space or with "-" (see Figure 4.60 also). **-CLIWP** and **-CLIWN**: These options display the WWN of the host adaptor in CLI format. PWWN or NWWN will be displayed for WWN (see Figure 4.62 also). **-sort [-CM]**: Sorts and displays the order of the Thunder 9500™ V subsystem product number and the volume management number. The **[-CM]** option searches for command device from the specified special file (raw device file) provided via STDIN or argument input, and then displays only the command device in the image of the configuration definition file (see Figure 4.65 also).

Table 4.28 Inqraid Command Parameters (continued)

Parameter	Value
Options (continued)	**-gvinf**: For Windows NT/2000 systems only. This option obtains the signature and the volume layout information from the device that is given by a standard input or by an argument, and saves (shelters) the information under the system device in the following format). File format: \WindowsDirectory\VOL*ssss_llll*.ini (*ssss* indicates the serial number (Equipment ID) of the subsystem, and *llll* indicates the LDEV number) Usually, you do not need to be aware of this file, since the file is used by the Windows' Disk Management at the beginning when setting the S-VOL signature and the volume information. **-svinf**: For Windows NT/2000 systems only. This option sets the saved (sheltered) signature and the system volume layout information of the system device into the device that is given by a standard input or by an argument. When setting the signature and the volume layout information to the device, the host issues the SCSI inquiry command, obtains the serial number (Equipment ID) of the subsystem and the LDEV number, and then reads the applicable VOLssss_llll.ini file. Therefore, even if the hard disk number is changed due to changing the configuration, the signature and the volume layout information will be set correctly since it is managed by the subsystem serial number (Equipment ID) and the LDEV number. For [=PTN] after –svinf option, specify a pattern for selecting the character string that is given as a device by a standard input or by an argument). **-fv**: For Windows 2000 systems only. This option is used with "$Volume specification". Volume{guid} of the appropriate volume is displayed in the wide format. Example: # inqraid -CLI $Vol -fv DEVICE_FILE PORT SERIAL LDEV CTG H/M/12 SSID R:Group PRODUCT_ID Volume{cec25efe-d3b8-11d4-aead-00c00d003b1e}\Vol3\Dsk0 CL1-B 2496 56 - - - - DF600F-CM
Returned values	The –svinf option returns the following values to distinguish the execution result from the user program: Normal termination: 0. Abnormal termination: 1 (when the execution to the specified device did not end normally).
Restriction	The special file of STDIN or Argument must be specified following name: HP-UX® Solaris®: /dev/rdsk/* Solaris®: /dev/rdsk/*s2, c*s2 AIX®: /dev/rhdisk, /dev/hdisk.hdisk... Linux®: /dev/sd... or /dev/rd... Tru64 UNIX®: /dev/rdisk/dsk*c IRIX®: /dev/rdsk/*vol, /dev/rdsk/node_wwn/*vol/*, /dev/dsk/*vol, /dev/dsk/node_wwn/*vol/* Windows® NT systems: hd0-10, harddisk0...(numbers indicate the drive number) $LETALL, $Phys, D:\DskX\pY, \DskX\pY Windows® 2000 systems: hd0-10, harddisk0...(numbers indicate the drive number) $LETALL, $Volume, $Phys, D:\Vol(Dms,Dmt, Dmr)X\DskY, \Vol(Dms,Dmt, Dmr)X\DskY For Windows® 2000 LDM volumes, see Appendix E.

```
ioscan -fun | grep rdsk | ./inqraid
/dev/rdsk/c0t2d1 -> [ST] CL2-A Ser =   3005 LDEV =   9 [HITACH  ] [DF600F          ]
                    HORC = NONE HOMRCF[MU#0 = SMPL MU#1 = NONE MU#2 = NONE]
                    RAID5[Group  2- 0] SSID = 0x0000
/dev/rdsk/c0t4d0 ->[ST] CL2-A Ser =   3005 LDEV =  14 [HITACH ] [DF600F-CM       ]
```

Figure 4.50 Inqraid Command Example (HP-UX®)

```
ls /dev/rdsk/* | ./inqraid
/dev/rdsk/c0t2d1s2 -> [ST] CL2-A Ser =  3005 LDEV =  9 [HITACH  ] [DF600F          ]
                      HORC = NONE HOMRCF[MU#0 = SMPL MU#1 = NONE MU#2 = NONE]
                      RAID5[Group  2- 0] SSID = 0x0000
/dev/rdsk/c0t4d0s2 ->[ST] CL2-A Ser =   3005 LDEV =  14 [HITACH ] [DF600F-CM       ]
```

Figure 4.51 Inqraid Command (Solaris®)

```
lsdev -C -c disk | ./inqraid
hdisk10 -> [ST] CL2-A Ser =   3005 LDEV =   9 [HITACH  ] [DF600F          ]
           HORC = NONE HOMRCF[MU#0 = SMPL MU#1 = NONE MU#2 = NONE]
           RAID5[Group  2- 0] SSID = 0x0000
hdisk11 -> [ST] CL2-A Ser =   3005 LDEV =  14 [HITACH ] [DF600F-CM       ]
```

Figure 4.52 Inqraid Command (AIX®)

```
ls /dev/sd* | inqraid
/dev/sdh -> CHNO =   0 TID =   1 LUN =   7
              [ST] CL2-A Ser =   3005 LDEV =   9 [HITACH  ] [DF600F          ]
              HORC = NONE HOMRCF[MU#0 = SMPL MU#1 = NONE MU#2 = NONE]
              RAID5[Group  2- 0] SSID = 0x0000
/dev/sdi -> CHNO =   0 TID =   4 LUN =   0
              [ST] CL2-A Ser =   3005 LDEV =  14 [HITACH ] [DF600F-CM       ]
```

Figure 4.53 Inqraid Command (Linux®)

```
ls /dev/rdisk/dsk* | ./inqraid
/dev/rdisk/dsk10c -> [ST] CL2-A Ser =  3005 LDEV =   9 [HITACH  ] [DF600F          ]
             HORC = NONE HOMRCF[MU#0 = SMPL MU#1 = NONE MU#2 = NONE]
             RAID5[Group  2- 0] SSID = 0x0000
/dev/rdisk/dsk11c -> [ST] CL2-A Ser =   3005 LDEV =  14 [HITACH ] [DF600F-CM       ]
```

Figure 4.54 Inqraid Command (Tru64™ UNIX®)

```
echo hd10-11 | ./inqraid
Harddisk10 -> [ST] CL2-A Ser =   3005 LDEV =   9 [HITACH  ] [DF600F          ]
              HORC = NONE HOMRCF[MU#0 = SMPL MU#1 = NONE MU#2 = NONE]
              RAID5[Group  2- 0] SSID = 0x0000
Harddisk11 -> [ST] CL2-A Ser =   3005 LDEV =  14 [HITACH ] [DF600F-CM       ]
```

Figure 4.55 Inqraid Command (Windows NT®/2000)

```
ls /dev/rdsk/*vol | ./inqraid
/dev/rdsk/dks1d6vol ->[ST] CL2-A Ser =   3005 LDEV =   9 [HITACH  ] [DF600F      ]
                           HORC = NONE  HOMRCF[MU#0 = SMPL MU#1 = NONE MU#2 = NONE]
                           RAID5[Group  2- 0] SSID = 0x0000
/dev/rdsk/dks1d7vol ->[ST] CL2-A Ser =   3005 LDEV =  14 [HITACH ] [DF600F-CM    ]
```

Figure 4.56 Inqraid Command (IRIX® FC_AL)

```
ls /dev/rdsk/*/*vol/* | ./inqraid
/dev/rdsk/50060e8000100262/lun3vol/c8p0 ->[ST] CL2-A Ser =   3005 LDEV =  9 [HITACH  ]
[DF600F     ]
                            HORC = NONE  HOMRCF[MU#0 = SMPL MU#1 = NONE MU#2 = NONE]
                            RAID5[Group  2- 0] SSID = 0x0000
/dev/rdsk/50060e8000100262/lun4vol/c8p0 ->[ST] CL2-A Ser =   3005 LDEV =  14 [HITACH ]
[DF600F-CM       ]
```

Figure 4.57 Inqraid Command (IRIX® Fabric Fibre)

The output of the inqraid command includes:

- **CL*X*-*Y***: Shows the port number
- **Ser**: Shows the serial number (Equipment ID)
- **LDEV**: Shows the LDEV ID. LDEV indicates LU for ShadowImage 9500V operations.
- **HOMRCF**: Shows ShadowImage 9500V attributes (PVOL/SVOL/SMPL) of the volume for MU#0 (zero fixed)
- **Group**: Shows the array (parity) group ID (physical position of the volume in the 9500). LDEV indicates LU for Synchronous TrueCopy 9500V/ShadowImage 9500V operations.
- **SSID**: Shows the subsystem ID of the volume. LDEV indicates LU for Synchronous TrueCopy 9500V/ShadowImage 9500V operations.
- **CHNO**: Channel number on the device adapter that recognizes the volume. Displayed only for Linux systems.
- **TID**: Target ID of the volume. Displayed only for Linux systems. For further information on Fibre-to-SCSI address conversion, see Appendix D.
- **LUN**: Logical unit number of the volume. Displayed only for Linux systems.

Note: The display of HOMRCF, **Group**, and **SSID** depends on the 9500V microcode level.

```
# echo /dev/rdsk/c23t0d0  /dev/rdsk/c23t2d3 | ./inqraid –find
Group  PairVol(L/R)   (Port#,TID,LU-M), Seq#,LDEV#, P/S, Status, Seq#, P-LDEV#  M
horcl  dev00(L)        (CL2-A,  0, 0-0)6145     0  S-VOL  SSUS, -----       9   -
->/dev/rdsk/c23t0d0
Group  PairVol (L/R)   (Port#,TID,LU-M), Seq#,LDEV#, P/S, Status, Seq#, P-LDEV#  M
horcl  dev10(L)        (CL2-A,  2, 3-0)6145     3  S-VOL  SSUS, -----       6   -
->/dev/rdsk/c23t2d3
```

Figure 4.58 Inqraid Command –find Option (HP-UX®)

```
# echo /dev/rdsk/c23t0d0  /dev/rdsk/c23t2d3 | ./inqraid –findc
DEVICE_FILE    M Group    PairVol  P/S   Stat  R_DEVICE  M  P/S   Stat  LK
C23t0d0  0 horcl    dev00    S-VOL SSUS  c23t0d1   0  P-VOL PSUS  OK
/dev/rdsl/c23t0d0[1]   -> No such on the group
/dev/rdsl/c23t0d0[2]   -> No such on the group
DEVICE_FILE    M Group    PairVol  P/S   Stat  R_DEVICE  M  P/S   Stat  LK
C23t2d3  0 horcl    dev10    S-VOL SSUS  c23t2d2   0  P-VOL PSUS  OK
/dev/rdsl/c23t2d3[1]   -> No such on the group
/dev/rdsl/c23t2d3[2]   -> No such on the group

# echo /dev/rdsk/c23t0d0  /dev/rdsk/c23t2d3 | ./inqraid -findc –CLI
DEVICE_FILE    M Group    PairVol  P/S   Stat  R_DEVICE  M  P/S   Stat  LK
C23t0d0  0 horcl    dev00    S-VOL SSUS  c23t0d1   0  P-VOL PSUS  OK
C23t2d3  0 horcl    dev10    S-VOL SSUS  c23t2d2   0  P-VOL PSUS  OK
```

Figure 4.59 Inqraid Command –findc Option (HP-UX®)

The output of the inqraid command with –find (and -findc) includes:

- **DEVICE_FILE**: Shows only the device file name
- **M**: Shows the MU# (zero fixed)
- **Group**: Shows the group name (dev_group) described in the configuration definition file
- **PairVol**: Shows the paired volume name (dev_name) of the group described in the configuration definition file
- **P/S**: Shows the volume attribute (P-VOL, S-VOL, SMPL)
- **Stat**: Shows the status of the paired volume
- **R_DEVICE**: Shows the device file name of the remote site
- **LK**: Shows the configuration check result for the paired volume connection path (physical link of the paired volumes between the servers)

```
# ls /dev/sd* | ./inqraid –CLI
DEVICE_FILE     PORT    SERIAL  LDEV  H/M/12  SSID R:Group PRODUCT_ID
sdh             CL2-B   3005    23   -/P/--  0004 5:02-01 DF600F
sdi             CL2-B   3005    14     -       -      -    DF600F-CM
sdj             -       -        -      -       -      -    -
```

Figure 4.60 Inqraid Command -CLI Option (Linux®)

The output of the inqraid command with -CLI includes:

- **DEVICE_FILE**: Shows only the device file name
- **PORT**: Shows the port name of the 9500
- **SERIAL**: Shows the product number (Equipment ID) of the 9500
- **LDEV**: Shows the volume management number in the 9500. LDEV indicates LU for Synchronous TrueCopy 9500V/ShadowImage 9500V operations.
- **H/M/12**: Shows the volume attribute of Synchronous TrueCopy 9500V/ShadowImage 9500V (P-VOL is indicated as P, S-VOL is indicated as S, and SMPL is indicated as s)
- **SSID**: Shows the subsystem ID where LDEV is allocated. LDEV indicates LU for Synchronous TrueCopy 9500V/ShadowImage 9500V operations.
- **R:Group**: Shows the position of the physical CCI group mapped in LDEV. LDEV indicates LU for Synchronous TrueCopy 9500V/ShadowImage 9500V operations.
- **PRODUCT_ID**: Shows the product ID in the standard inquiry page

```
# echo  /dev/rdsk/c23t0d0  /dev/rdsk/c23t0d1  |  ./inqraid –CLIWP
DEVICE_FILE    PWWN             AL PORT  LUN    SERIAL   LDEV  PRODUCT_ID
c23t0d0        500060e802f01018  -  CL2-A   -    6145      12   DF600F
c23t0d1        500060e802f01018  -  CL2-A   -    6145      12   DF600F
```

Figure 4.61 Inqraid Command –CLIWP Option Example (HP-UX®)

The output of the inquaid command with –CLIWP includes:

- **DEVICE_FILE**: Shows only the device file name
- **WWN (PWWN or NWWN)**: If –CLIWP option is specified, this option shows the Port_WWN of the host adaptor for the specified device. If –CLIWN option is specified, this option shows the Node_WWN of the host adaptor.
- **AL** and **LUN**: Shows a hyphen (-) all the time
- **PORT**: Shows the port name of the 9500V
- **SERIAL**: Shows the product number (Equipment ID) of the 9500V
- **LDEV**: Shows the volume management number in the 9500V. LDEV indicates LU for Synchronous TrueCopy 9500V/ShadowImage 9500V operations.
- **PRODUCT_ID:** Shows the product ID in the standard inquiry page

```
# ioscan  -fun  |  grep  rdsk  |  ./inqraid  -sort  -CM  -CLI
HORCM_CMD
# dev_name            dev_name               dev_name
#UnitID 0  (Serial# 3001)
/dev/rdsk/c0t3d0        /dev/rdsk/c1t2d1
#UnitID 1  (Serial# 3002)
/dev/rdsk/c2t3d0
```

Figure 4.62 Inqraid Command –sort[-CM] Option (HP-UX®)

Note: The unit ID is added in the order of the 9500V subsystem's product number (Equipment ID). If multiple command devices exist within the subsystem, the device file that is used to share between the subsystem port would be chosen in prior, and is used as alternate command device.

The following is an example of using the **-gvinf** option of the inqraid command. The information in all physical drives will be saved (sheltered) by giving **$Phy**.

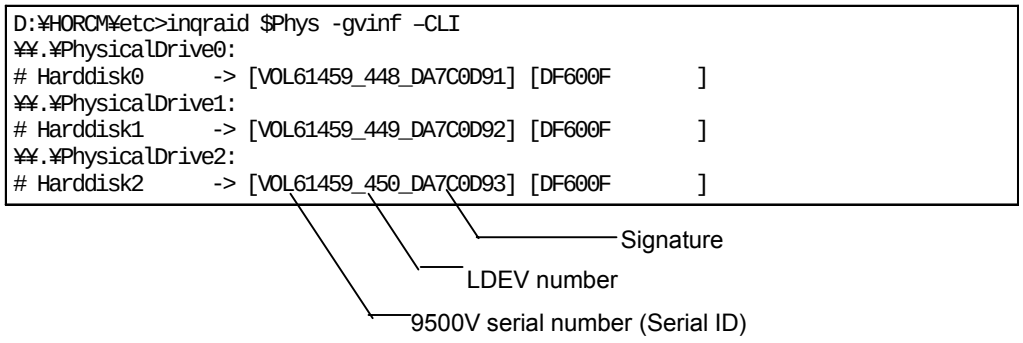

Figure 4.63 Inqraid Command –gvinf Option Example

The following is an example of using the **-svinf** option of the inqraid command. The information is set to the hard disk number indicated by the pairdisplay command that is issued from S-VOL instance.

```
D:\HORCM\etc>pairdisplay -l -fd -g URA
Group   PairVol(L/R) Device_File   M ,Seq#,LDEV#.P/S, Status, Seq#,P-LDEV# M
URA     URA_000(L)    Harddisk3     0  6145  51..S-VOL SSUS,-----     48  -
URA     URA_001(L)    Harddisk4     0  6145  52..S-VOL SSUS,-----     49  -
URA     URA_002(L)    Harddisk5     0  6145  53..S-VOL SSUS,-----     50  -

D:\HORCM\etc>pairdisplay -l -fd -g URA | inqraid -svinf=Harddisk
[VOL6145_51_5296A763] -> Harddisk3        [DF600F          ]
[VOL6145_52_5296A760] -> Harddisk4        [DF600F          ]
[VOL6145_53_5296A761] -> Harddisk5        [DF600F          ]
```

Figure 4.64 Inqraid Command –svinf Option Example

Important: If you have created a pair using a noread option, the device object (\Device\HarddiskVolume#) and the Volume{guid} of the S-VOL will not be generated at the time when the Windows is booted in Noread status. However, executing the pairsplit command, and then executing the inqraid **-svinf** command results in generation of the device object (\Device\HarddiskVolume#) and the Volume{guid}.

The hard disk volume number of the device object is generated in the order of the command that the Windows issued. Therefore, the same number will be maintained unless the configuration is changed. However, if you want the hard disk volume number to be absolute, add the –sort option, as shown below, so that the serial number of the subsystem and the LDEV number will be sorted in numbers.

Example:

```
D:\HORCM\etc>echo hd5 hd4 hd3 | inqraid -svinf –sort
[VOL6145_51_5296A763] -> Harddisk3        [DF600F          ]
[VOL6145_52_5296A760] -> Harddisk4        [DF600F          ]
[VOL6145_53_5296A761] -> Harddisk5        [DF600F          ]
```

4.15.2 Mkconf Command Tool

The **mkconf** command executes the following operation and creates a configuration definition file automatically from the special file via STDIN. The user must edit the created configuration definition file as needed.

1. The mkconf command executes **inqraid –sort –CM –CLI** and then creates a configuration definition file only for HORCM_CMD.
2. The mkconf command starts the HORCM instance using the created definition file.
3. The mkconf command executes **raidscan –find conf** by using the special file via STDIN, and then creates the definition file that contains HORCM_DEV and HORCM_INST.
4. The mkconf command starts the HORCM instance again to verify the created definition file.
5. The mkconf command executes **raidscan –find verify** and displays the correspondence between the special files via STDIN with the definition file.

Table 4.29 Mkconf Command Parameters

Parameter	Value
Command Name	mkconf
Format	/HORCM/usr/bin/mkconf.sh [-g <group> [-m <mu#> [-i <inst#>] [-s <service>] [-a]] \HORCM\etc\ mkconf.exe [-g <group> [-m <mu#> [-i <inst#>] [-s <service>] [-a] [-c <drive>]
Options	**If no argument**, this command creates a mode that waits for STDIN without argument. **-g <group>**: This option specifies the group described in the configuration definition file. If omitted, "VG" is used for the group name. **-m <MU#>**: This option specifies the mirror descriptor MU# (zero fixed). **-i <inst>**: This option shows the instance number. **-s <service>**: This option specifies the service name (port number) described in the configuration definition file. If omitted, "52323" is used for the port number. **-a**: This option adds a group to an already created configuration definition file. **-c <drive>**: (for Windows NT/200 only) This option specifies the range of the command device to be searched. If omitted, "$PhysicalDrive" is used for <drive>.
Note	The configuration definition file and the log file are created in the current directory as horcm*.conf or log* (* as an instance number). The user must edit the created configuration definition file as necessary (such as ip_addresses, and services).

WARNING: If you have created the configuration definition file using the mkconf command tool, remember to change the value of the poll(10ms) manually. For details on calculating the poll(10ms) value, see section 2.5.3. Setting the value incorrectly may cause a conflict between CCI and the subsystem; the internal processing of the subsystem may suspend temporarily. Processing may not proceed.

```
# cd /tmp/test
# cat /etc/horcm.conf | /HORCM/usr/bin/mkconf.sh -g ORA -i 9 -m 0
starting HORCM inst 9
HORCM inst 9 starts successfully.
HORCM Shutdown inst 9 !!!
A CONFIG file was successfully completed.
starting HORCM inst 9
HORCM inst 9 starts successfully.
DEVICE_FILE             Group    PairVol     PORT   TARG  LUN M   SERIAL  LDEV
/dev/rdsk/c23t0d0       ORA      ORA_000     CL2-A     0    0 0    6145      0
/dev/rdsk/c23t0d1       ORA      ORA_001     CL2-A     0    1 0    6145      1
/dev/rdsk/c23t0d2       ORA      ORA_002     CL2-A     0    2 0    6145      2
/dev/rdsk/c23t0d3       ORA      ORA_003     CL2-A     0    3 0    6145      3
/dev/rdsk/c23t0d4       ORA      ORA_004     CL2-A     0    4 0    6145      4
/dev/rdsk/c23t0d5       ORA      ORA_005     CL2-A     0    5 0    6145      5
/dev/rdsk/c23t0d6       ORA      ORA_006     CL2-A     0    6 0    6145      6
/dev/rdsk/c23t0d7       -         -            -       -    - 0    6145      -
HORCM Shutdown inst 9 !!!
Please check '/tmp/test/horcm9.conf','/tmp/test/log9/curlog/horcm_*.log', and modify
ip_address & service'.

# ls
horcm9.conf  log9
# vi horcm9.conf

# Created by mkconf.sh on Mon Jan 22 17:59:11 JST 2001

HORCM_MON
#ip_address        service         poll(10ms)     timeout(10ms)
localhost          52323           1000[Note]             3000

HORCM_CMD
#dev_name               dev_name               dev_name
#UnitID 0 (Serial# 6145)
/dev/rdsk/c23t3d0

HORCM_DEV
#dev_group         dev_name        port#     TargetID     LU#     MU#
# /dev/rdsk/c23t0d0    SER =    6145  LDEV =    0  [ FIBRE FCTBL = 4 ]
ORA                ORA_000          CL2-A             0        0       0
# /dev/rdsk/c23t0d1    SER =    6145  LDEV =    1 [ FIBRE FCTBL = 4 ]
ORA                ORA_001          CL2-A             0        1       0
# /dev/rdsk/c23t0d2    SER =    6145  LDEV =    2 [ FIBRE FCTBL = 4 ]
ORA                ORA_002          CL2-A             0        2       0
# /dev/rdsk/c23t0d3    SER =    6145  LDEV =    3 [ FIBRE FCTBL = 4 ]
ORA                ORA_003          CL2-A             0        3       0
# /dev/rdsk/c23t0d4    SER =    6145  LDEV =    4 [ FIBRE FCTBL = 4 ]
ORA                ORA_004          CL2-A             0        4       0
# /dev/rdsk/c23t0d5    SER =    6145  LDEV =    5 [ FIBRE FCTBL = 4 ]
ORA                ORA_005          CL2-A             0        5       0
# /dev/rdsk/c23t0d6    SER =    6145  LDEV =    6 [ FIBRE FCTBL = 4 ]
ORA                ORA_006          CL2-A             0        6       0
# ERROR [CMDDEV] /dev/rdsk/c23t0d7    SER =    6145  LDEV =    7 [DF600F-CM       ]

HORCM_INST
#dev_group        ip_address      service
ORA                localhost       5232
```

Figure 4.65 Mkconf Command Example (HP-UX®)

Note: Remember to change the value of poll(10ms) parameter using the equation described in section 2.5.3.

Notes:

- Unit IDs are added in the order of the subsystem product number. If multiple command devices exist in the subsystem, the device file that was shared between the ports of the subsystem will be selected, and will be handled as an alternative command device.
- If the standard input device includes a command device, the target device will be displayed as the comment shown below and will be omitted.

Example:

```
# ERROR [CMDDEV] /dev/rdsk/c23t0d7  SER = 6145  LDEV = 7 [DF600F-CM      ]
```

- If the standard input device is shared between several command devices and is already displayed as a target device, the target device will be displayed as the comment shown below and will be omitted.

Example:

```
# ERROR [LDEV LINK] /dev/rdsk/c24t0d3 SER = 6145  LDEV = 3 [FIBRE FCTBL = 4]
```

- If the standard input device does not have an appropriate mirror descriptor (MU#), the target device will be displayed as the comment shown below and will be omitted.

Example:

```
# ERROR [LDEV MUN (2<1)] /dev/rdsk/c24t0d3  SER = 6145  LDEV = 3 [DF600F  ]
```

- If the device from the standard input co-reside between subsystems with different mirror control, the target device will be displayed as the comment shown below and will be omitted.

Example:

```
# ERROR [MIXING RAID TYPE] /dev/rdsk/c24t0d3 SER = 6145 LDEV = 3 [DF600F  ]
```

Chapter 5 Troubleshooting

This chapter includes the following troubleshooting information:

- ShadowImage 9500V Troubleshooting
- Synchronous Remote Copy (Synchronous TrueCopy 9500V) Troubleshooting
- General Troubleshooting

5.1 ShadowImage 9500V Troubleshooting

If a hardware error occurs while you are operating ShadowImage 9500V, both of the following are necessary:

- A CCI user intervention
- Assistance from a Hitachi Data Systems Customer Service representative

For example, when formatting is needed to resolve an LU error and that LU is used for ShadowImage 9500V, the pair must be released by the user (CCI operation by the user) before the LU can be formatted. Therefore, please contact Hitachi personnel because maintenance requires the user to issue CCI commands. Note that the Hitachi personnel can only remove errors which result from hardware. An operation such as recovering a ShadowImage 9500V pair status (e.g. resynchronizing) must be done by the user. Figure 5.1 shows the flow of action when the PSUE error occurs. Table 5.1 shows the share of action to be taken by the user.

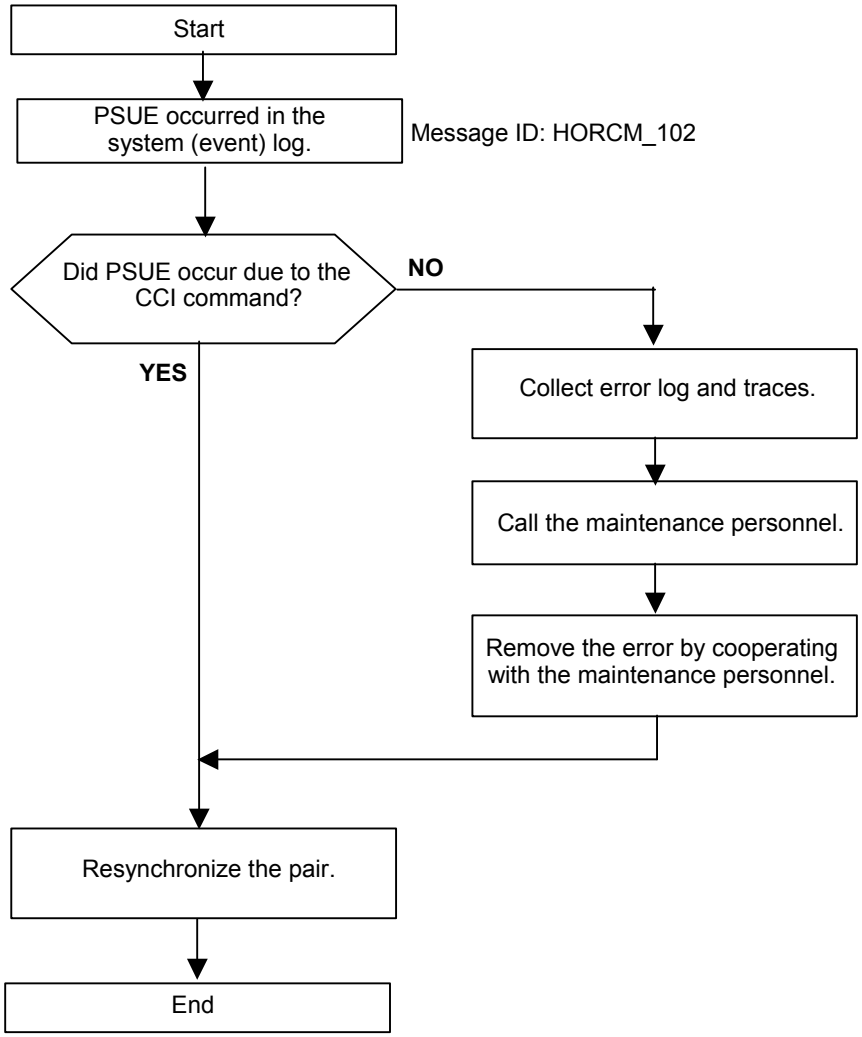

Figure 5.1 Pair Status Information Example

Table 5.1 Operational Notes for ShadowImage Operations

Action	Action taken by whom
Confirm the message (syslog).	User
Verify that PSUE is caused by the user operation.	User
Verify the status of the subsystem.	User
Call maintenance personnel when the subsystem malfunctions.	User
For other reasons, call the Hitachi support center.	User (only for users that are registered in order to receive a support)
Hardware maintenance.	Hitachi Customer Service
Reconfigure and recover the pair.	User

5.2 Synchronous TrueCopy 9500V Troubleshooting

If a hardware or Fibre path error occurs while you are operating the Synchronous TrueCopy 9500V, both of the following are necessary:

- A CCI user intervention
- Assistance from a Hitachi Data Systems Customer Service representative

For example, when formatting is needed to resolve an LU error and that LU is used for Synchronous TrueCopy 9500V, the pair must be released by the user (CCI operation by the user) before the LU can be formatted. Therefore, please contact Hitachi personnel because maintenance requires the user to issue CCI commands. Note that the Hitachi personnel can only remove errors which result from hardware. An operation such as recovering the Synchronous TrueCopy 9500V pair status (e.g. resynchronizing) must be done by the user. Figure 5.2 shows the flow of action when the PSUE error occurs. Table 5.2 shows the share of action to be taken by the user.

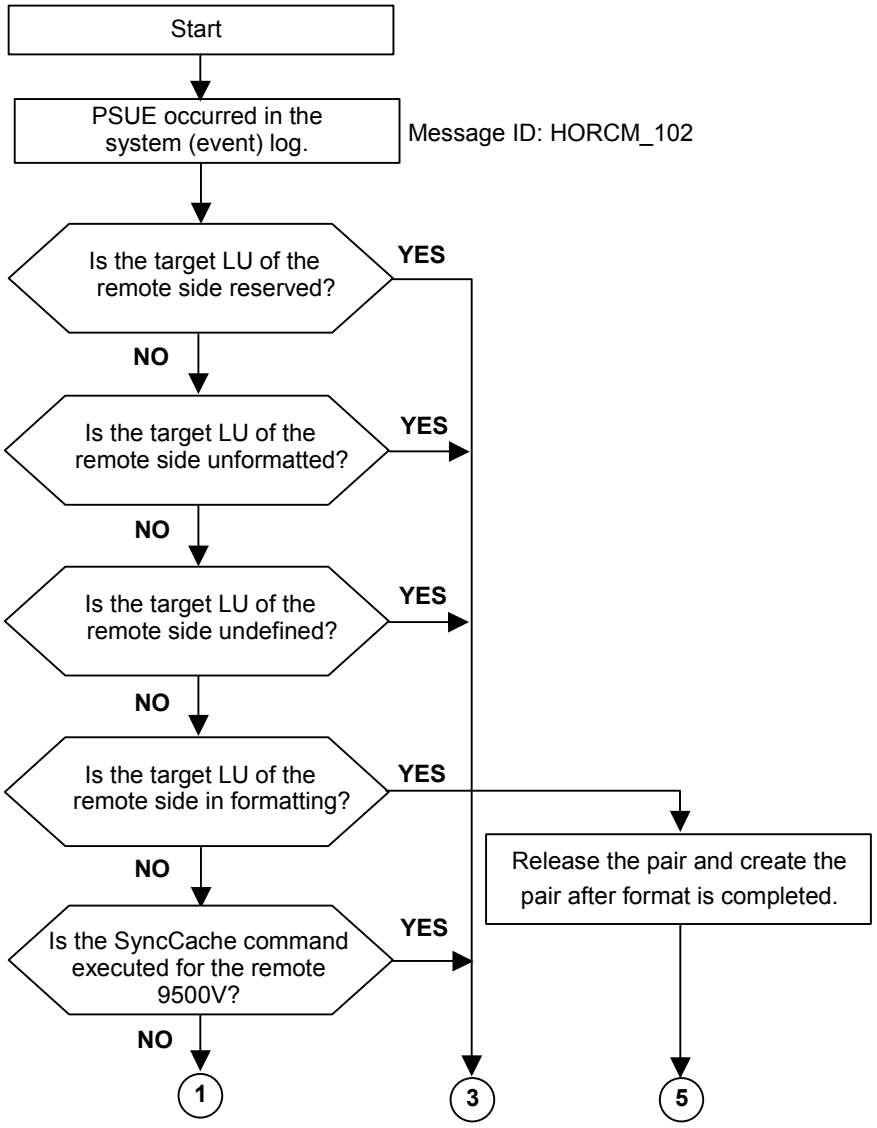

Figure 5.2 Pair Status Information Example Using Synchronous TrueCopy 9500V

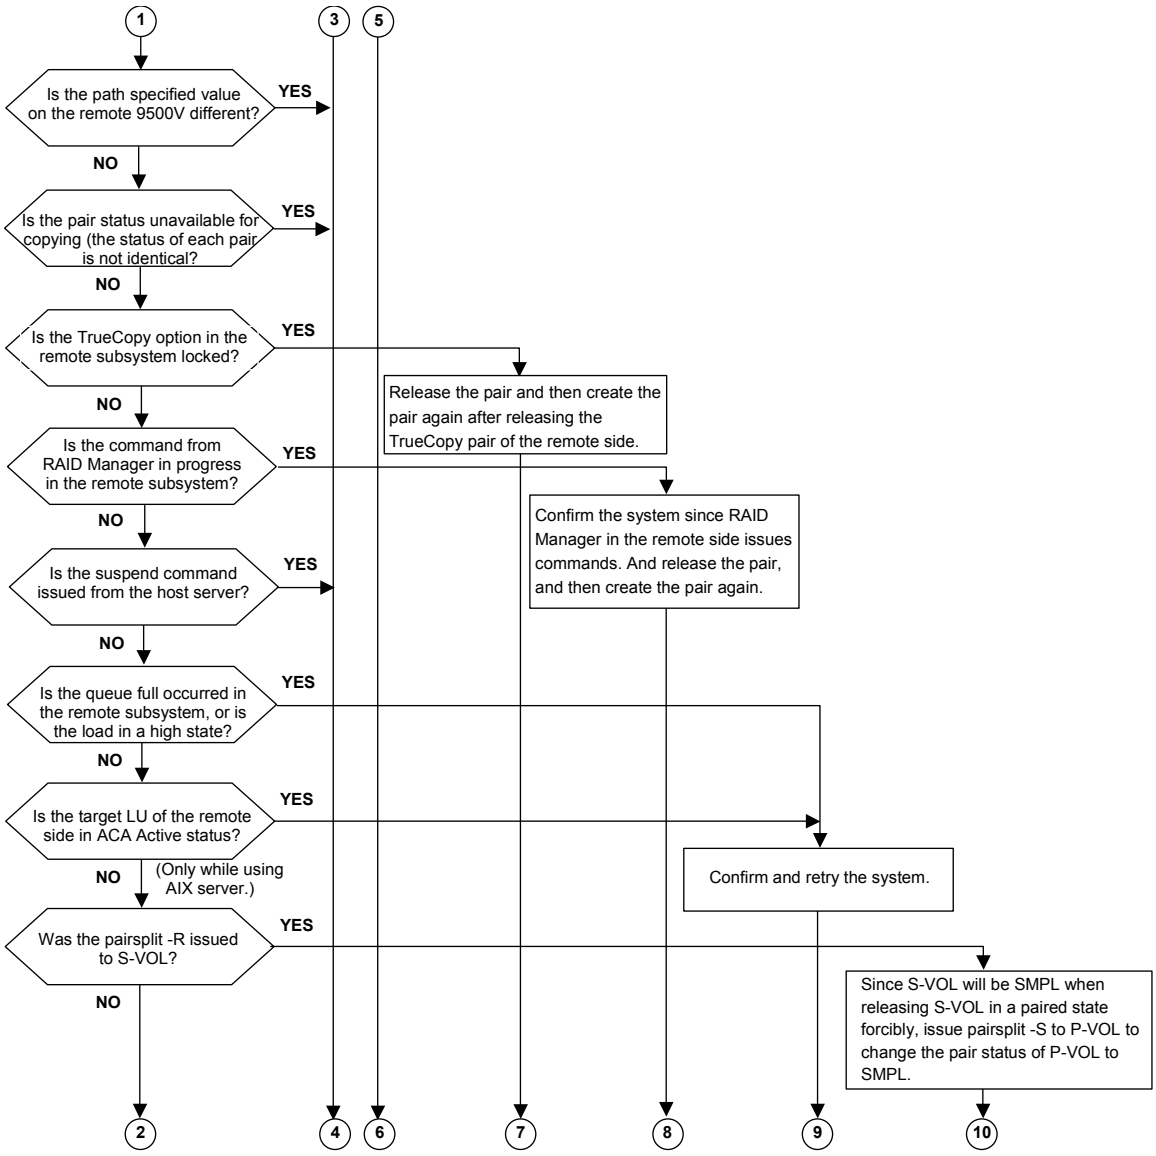

Figure 5.2 Pair Status Information Example Using Synchronous TrueCopy 9500V (continued)

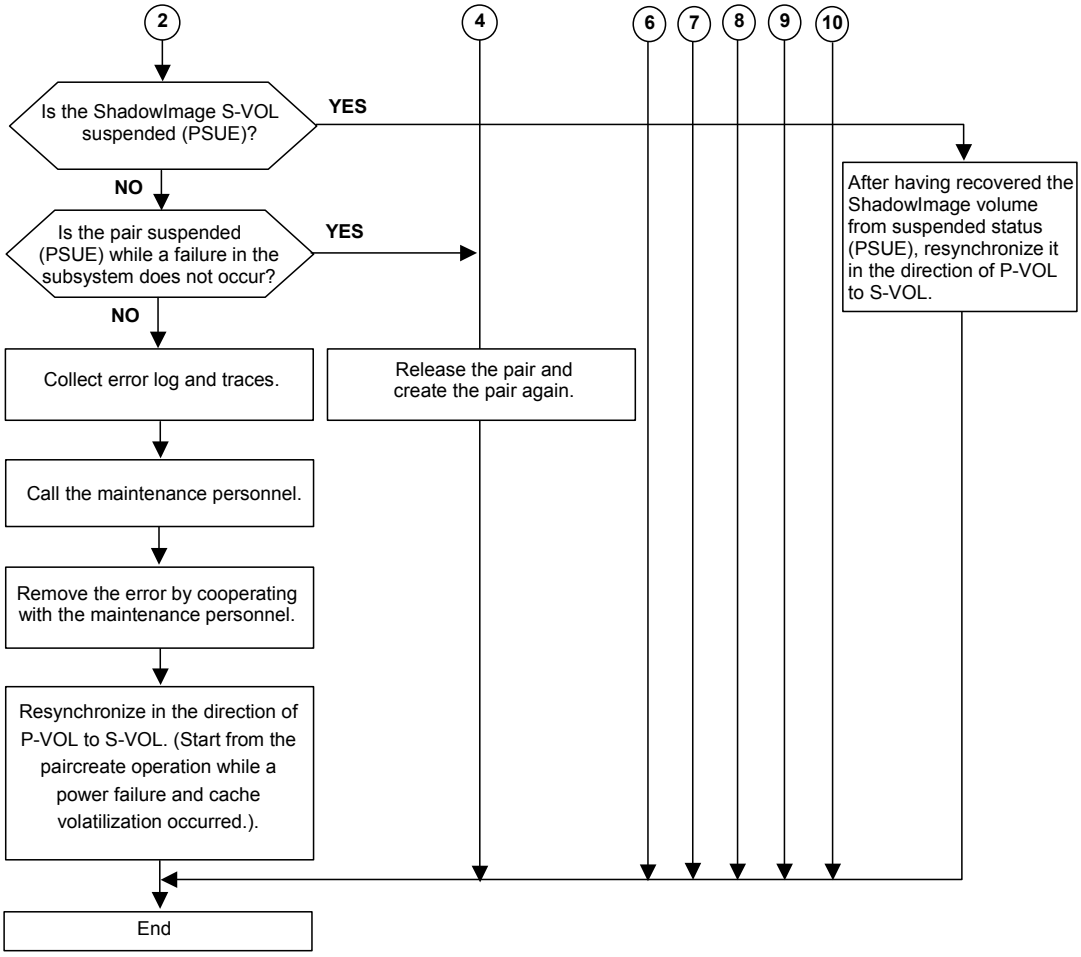

Figure 5.2 Pair Status Information Example Using Synchronous TrueCopy 9500V (continued)

When an error occurs in the Fibre route between subsystems while operating Synchronous TrueCopy 9500V, the path is blocked. The system maintenance by the Hitachi personnel is necessary. Figure 5.3 shows the flow of the failure isolation while the path blockage occurred.

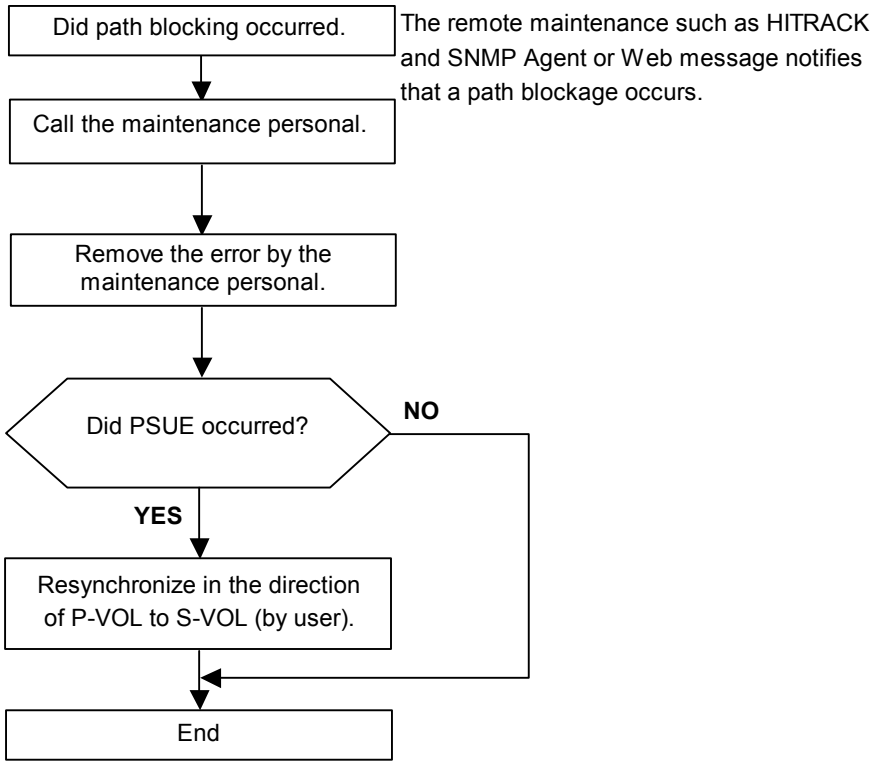

Figure 5.3 Path Status Information Example Using Synchronous TrueCopy 9500V

Table 5.2 Operational Notes for Synchronous TrueCopy 9500V Operations

Action	Action taken by whom
Confirm the message (syslog).	User
Confirm whether PSUE is caused by the user operation.	User
Confirm the status of the subsystem.	User
Call the maintenance personnel when the subsystem malfunctions.	User
For other reasons, call the Hitachi support center.	User (only for users that are registered in order to receive a support)
Hardware maintenance (including path blockage).	Hitachi maintenance personnel
Reconfigure and recover the pair.	User

5.3 General Troubleshooting

If you have a problem with the CCI software, first verify that the problem is not being caused by the UNIX®/PC server hardware or software, and try restarting the server. Table 5.3 provides operational notes and restrictions for CCI operations.

For maintenance of Synchronous TrueCopy 9500V/ShadowImage 9500V volumes, if a failure occurs, it is important to find the failure in the paired volumes, recover the volumes, and continue operation in the original system.

Table 5.3 Operational Notes for CCI Operations (Continued on next page)

Condition	Recommended Action
Startup/shutdown restrictions	When the server starts, the primary volume server may update the secondary volume. The secondary volume must not be mounted automatically in the startup sequence. If the secondary volume is used by the LVM, the volume group of the LVM must be deactivated. The secondary volume must be mounted in the split state or in the simplex mode. When the server starts, the secondary volume can be activated without confirmation, once it can be guaranteed that the secondary volume has been PSUS (R/W enable) or it is in the SMPL state during the server shutdown sequence.
Hot standby operations	CCI commands cannot execute hot standby operations between the primary and secondary volumes. Use the takeover command intended for the HA configuration to execute the hot standby operation. In hot standby operation, two servers are used, and the active (primary) and standby (secondary) server programs are run alternately in each server in case one server fails. Follow these precautions: **Operation across volumes**. Since each CCI command causes the server software to handle the volume by volume, a single volume should not be partitioned to prevent it from being used by selected servers. **Using LVM and paired volume together**. When constructing the LVM on the paired volume in the mutual hot standby configuration, the LVM logical volumes must be constructed in units of volume to prevent the volumes from being mounted by the LVM.
Coexistence of LVM mirror and Synchronous TrueCopy 9500V	When the LVM mirror and Synchronous TrueCopy 9500V volumes are used together, the LVM mirror handles write errors and changes the volumes. Thus, the fence level of the volumes used by the LVM mirror must be set to **data**.
Using paired volume in a single host	When constructing a paired volume in a single host, it is necessary to activate two or more CCI instances. To activate two or more CCI instances, instance numbers must be assigned using the environmental variable **HORCMINST**. The HORCM and CCI commands must possess this environmental variable. A configuration definition file and a log directory must be set for each instance. The command device described in the configuration definition file must be established to follow every instance. If using a command device between different instances on the same port, the maximum number of instances per command device is 16. If this number is exceeded, use a different path for each instance.

Table 5.3 Operational Notes for CCI Operations (continued)

Condition	Recommended Action
Sharing volumes in a hot standby configuration	When a paired volume is used for the disk shared by the hosts in a hot standby configuration using HA software, use the primary volume as the shared disk and describe the corresponding hosts using the paired volume in the configuration definition file as shown below. In the HA configuration, if a CCI command issued by host C fails in host B (because host B has gone down and/or there is a command device IO_ERROR), host A is connected and command execution is retried. 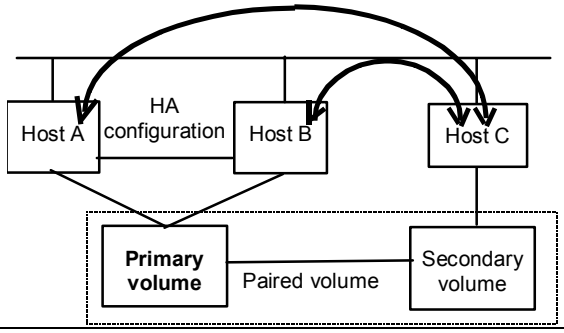
Linkage with HA software	The HORCM must not be an object of the process monitored by the HA software (cluster manager), because HORCM should run at the same level as the cluster manager. ***Note:*** Do not use the pair volume for the cluster lock disk that HA software uses for election.
Maintenance	A HORCM restart is required if the 9500V configuration is changed (e.g., microcode exchange, cache memory install/uninstall).
Command device	Each CCI command is executed by issuing a command to the command device. This command is read or written from/into the specific block area of the command device. Therefore, the user cannot use the command device. In addition, this device must not belong to an LVM volume group. For Windows® systems, do not assign a drive letter to the command device to prevent utilization by general users.

Table 5.3 Operational Notes for CCI Operations (continued)

Condition	Recommended Action
Alternate path restrictions	If the P-VOL and S-VOL are on the same server, an alternate path from P-VOL to S-VOL cannot be used. Use of alternate path to a volume pair is limited to primary (secondary) volumes.
HORCM failure to activate	After a new system has been constructed, failure to activate HORCM may occur due to an improper environmental setting and/or configuration definition by the user. Refer to the HORCM activation log, and correct the setting(s).
Abnormal termination of command	Refer to the command log file and HORCM log file to identify the cause of the error. If a command terminates abnormally because of a remote server failure, recover the server from the failure, then re-execute the command. If HORCM has shut down, restart HORCM. If an unrecoverable error occurs, obtain the log files and contact the Hitachi Data Systems Support Center.
Restriction for formatting the volumes	Do not execute Synchronous TrueCopy 9500V/ShadowImage 9500V operations while formatting the volume. Formatting takes priority and Synchronous TrueCopy 9500V/ShadowImage 9500V operations will be suspended.
The poll(10ms) parameter in the configuration definition file	Always set the poll(10ms) parameter with a value more than or equal to 6,000. For details on calculating the poll(10ms) parameter, see section 2.5.3. If you have created the configuration definition file using the mkconf command tool, set this value manually. If the value is set incorrectly, it may cause a conflict between CCI and the subsystem; the internal processing of the subsystem suspends temporarily. The process may not proceed.
Internal process conflict between the CCI and the subsystem	When an internal process conflict occurs between the CCI and the subsystem, the process of the subsystem suspends temporarily. If the conflict continues, the internal process may not proceed. Therefore, when monitoring (polling) the status of the subsystem (by creating a script using the CCI commands), set the information-display-based commands (e.g. pairdisplay, raidscan, raidar, and raidqry) to be issued to more than or equal to a minute.
Differential Management on Synchronous TrueCopy 9500V/ShadowImage 9500V	Synchronous TrueCopy 9500V/ShadowImage 9500V can manage a differential data capacity up to 144 GB×100 HDDs. Therefore, for an LU that will be set as a Synchronous TrueCopy 9500V/ShadowImage 9500V pair, a pair with a capacity more than 144 GB×100 HDDs cannot be created while using 72 GB HDD.
Differential Management while Cascading Synchronous TrueCopy 9500V and ShadowImage 9500V	Up to 964 blocks of differential data can be managed regarding 16,384 bits as a block. The volume capacity per block depending on the width of RAID Group shows below. Set the volume capacity referring to the example for fear of exceeding 964 blocks.
Restriction for Synchronous TrueCopy 9500V	Since S-VOL is reserved while using S-VOL of Synchronous TrueCopy 9500V, the resources of MSCS, P-VOL cannot be copied to S-VOL. Therefore, S-VOL of Synchronous TrueCopy 9500V must not be used as the resources of MSCS

Table 5.4 Width of RAID Group and Volume Capacity for One Block of the Difference

Width of RAID group	Volume capacity for one block of the difference (Unit: GB)
ShadowImage 9500V (2D ~ 15D + 1P)	16
Synchronous TrueCopy 9500V (2D + 1P), (4D + 1P), (8D + 1P)	16
Synchronous TrueCopy 9500V (3D + 1P), (6D + 1P), (12D + 1P)	24
Synchronous TrueCopy 9500V (5D + 1P), (10D + 1P)	20
Synchronous TrueCopy 9500V (7D + 1P), (14D + 1P)	28
Synchronous TrueCopy 9500V (9D + 1P)	18
Synchronous TrueCopy 9500V (11D + 1P)	22
Synchronous TrueCopy 9500V (13D + 1P)	26
Synchronous TrueCopy 9500V (15D + 1P)	30

Note:

Calculation method of the differential blocks number (for cascading Synchronous TrueCopy 9500V and ShadowImage 9500V)

Differential blocks number=
| LU size/LU capacity for one block of the difference for each width of RAID group | ↑ × 2 +
| LU size/LU capacity for one block of the difference for each width of RAID group | ↑
(ShadowImage 9500V is necessary for the duplicate differential area)

Example: The total number of the required blocks when the 600 GB LUs are cascaded Synchronous TrueCopy 9500V and ShadowImage 9500V :

(RAID5 14D+1P)

Differential blocks number=| 600/16| ↑ × 2 + |600/28| ↑=98

5.4 Error Reporting

Table 5.5 lists and describes the HORCM system log messages and provides guidelines for resolving the error conditions.

Table 5.6 lists and describes the command error messages and their return values and also provides guidelines for resolving the error conditions. Table 5.7 and Table 5.8 list and describe the common error messages for each command. Table 5.9 lists the unique error messages.

Table 5.5 System Log Messages

Message ID	Condition	Cause	Recommended Action
HORCM_001	The HORCM log file cannot be opened.	The file cannot be created in the HORCM directory.	Create space on the disk where the root directory resides.
HORCM_002	The HORCM trace file cannot be opened.	The file cannot be created in the HORCM directory.	Create space on the disk where the root directory resides.
HORCM_003	The HORCM daemon process cannot create a child process due to an error.	HORCM daemon attempted to create more processes than the maximum allowable number.	Terminate unnecessary programs or daemon processes running simultaneously.
HORCM_004	HORCM assertion failed, resulting in a fatal internal error in the HORCM.	An internal error occurred which could not be identified by the HORCM.	Restart the system, collect all HORCM log and trace information, and call the Hitachi Data Systems support center.
HORCM_005	The CCI software failed to create the end point for remote communication.	HORCM failed to create a socket, or an error exists in the HORCM configuration file ($HORCM_CONF).	Refer to the HORCM startup log to identify the cause of the error and reset the parameters.
HORCM_006	HORCM memory allocation failed.	HORCM memory could not be secured.	Increase the system virtual memory, or close any unnecessary programs.
HORCM_007	An error exists in the HORCM setup file.	An error exists in the HORCM setup file.	Refer to the startup log and reset the parameters.
HORCM_008	HORCM configuration file parameters could not be read.	An error exists in the format or parameters of the HORCM configuration file ($HORCM_CONF).	Refer to the HORCM startup log to identify the cause of the error.
HORCM_009	ShadowImage 9500V connection to the CCI software failed.	System devices are improperly connected, or an error exists in the HORCM configuration file.	Refer to the HORCM startup log to identify the cause of the error.
HORCM_101	ShadowImage 9500V and the CCI software communication fails.	A system I/O error occurred or an error exists in the HORCM configuration file ($HORCM_CONF).	Refer to the HORCM startup log to identify the cause of the error.
HORCM_102	The volume is suspended in code 0006.	The pair status was suspended due to code 0006.	(For ShadowImage 9500V) Determine how the pair was suspended. If the pair was forcibly suspended, resynchronize the pair. If the pair was not suspended forcibly, the pair was suspended due to an error in the subsystem. Collect all HORCM log and trace information, and call the Hitachi Data Systems support center. After maintenance personnel has recovered the subsystem, refer to section 5.5 to recover the data. (For Synchronous TrueCopy 9500V) Refer to section 5.6 to recover the data.

Table 5.6 Command Error Messages (Continued on next page)

Message ID	Error Message	Condition	Recommended Action	Return Value
EX_COMERR	Cannot Communicate with HORC Manager	This command failed to communicate with the CCI software.	Verify that HORCM is running by using UNIX® commands [ps - ef | grep horcm].	255
EX_REQARG	Required Arg list	An option or arguments of an option are not sufficient.	Choose the correct option using the -h option.	254
EX_INVARG	Invalid argument	An option or arguments of an option are incorrect.	Choose the correct option using the -h option.	253
EX_UNWOPT	Unknown option	Designated an unknown option.	Choose the correct option using the -h option.	252
EX_ATTHOR	Cannot attach to HORC Manager	Could not connect with HORCM.	Verify that HORCM is running and/or that HORCMINST is set correctly.	251
EX_ATTDBG	Cannot attach to a Debug layer	Failed to communicate with HORCM, or cannot make a log directory file.	Verify that HORCM is running by using UNIX® commands [ps - ef | grep horcm].	250
EX_INVNAM	Invalid name of option	The name specified in an argument of an option is not appropriate.	Choose the correct option using the -h option.	249
EX_OPTINV	A specified option is invalid	Detected contradiction in information that RAID reported.	Call the Hitachi Data Systems Support Center.	248
EX_ENOENT	No such device or group	The designated device or group name does not exist in the configuration file.	Verify the device or group name and add it to the configuration file of the remote and local hosts.	247
EX_ENODEV	No such device	The designated device name does not exist in the configuration file.	Verify the device name and add it to the configuration file of the remote and local hosts.	246
EX_ENOUNT	No such RAID unit	The designated RAID unit ID does not exist in the configuration file.	Verify the RAID unit ID and add it to the configuration file of the remote and local hosts.	219
EX_ENQSER	Unmatched Serial# vs RAID unit ID	The group designated by ShadowImage 9500V paircreate does not have the same RAID unit, or the unitID is not identical to the unit ID in the same RAID serial# (Equipment ID).	Verify the serial# (Equipment ID) using the pairdisplay command, or confirm serial# (Equipment ID) of the RAID using the raidqry -r command	218
EX_ENOMEM	Not enough core	Insufficient memory exists.	Increase the virtual memory capacity of the system, or close any unnecessary programs and/or daemon processes.	245
EX_ERANGE	Result too large	Tried to use arguments for an option beyond the maximum allowed, or a result beyond the maximum was created.	Refer to the error message, and designate an appropriate value.	244
EX_ENAMLG	File name too long	Undefined error.	-	243
EX_ENORMT	No remote host alive for remote commands or remote HORCM might be blocked (sleeping) on an existing I/O	A timeout occurred on remote communication, and HORC Manager failed to re-execute.	Confirm that the HORC Manager in the remote host is running, and then increase the value of the timeout in the configuration file. If you want to execute within the local server, execute the command that supports the local option (-l).	242
EX_INVMOD	Invalid RAID command mode	Detected a contradiction for a command.	Collect all HORCM log information, and call the Hitachi Data Systems Support Center.	241

Table 5.6 Command Error Messages (continued)

Message ID	Error Message	Condition	Recommended Action	Return Value
EX_INVCMD	Invalid RAID command	Detected a command error.	Collect all HORCM log information, and call the Hitachi Data Systems Support Center.	240
EX_ENOGRP	No such group	The designated device or group name does not exist in the configuration file, or the network address for remote communication does not exist.	Verify the device or group name and add it to the configuration file of the remote and local hosts.	239
EX_UNWCOD	Unknown function code	Detected a command error.	Collect all HORCM log information, and call the Hitachi Data Systems Support Center.	238
EX_CMDIOE	Control command I/O error	A read/write to the command device failed with an I/O error.	Refer to the host syslog file, and investigate the cause of the error. If the problem persists, collect all HORCM log information, and call the Hitachi Data Systems Support Center.	237
EX_CMDRJE	An order to the control/command device was rejected	The request to the command device failed or was rejected. ***Note:*** This error code is sometimes caused by the operating system and reported as EX_CMDIOE instead of EX_CMDRJE (see next row).	Verify Synchronous TrueCopy 9500V/ShadowImage 9500V functions are installed. Verify that the target volume is available. If the problem persists, collect all HORCM log information, and call the Hitachi Data Systems Support Center.	221
EX_CMDIOE	Control command I/O error or rejected	A read/write to the command device failed with an I/O error or was rejected.	Refer to the host syslog file, and investigate the cause of the error. If the cause is “Illegal Request(0x05)” Sense Key, please confirm the following items. If the problem persists, collect all HORCM log information, and call the Hitachi Data Systems Support Center. Verify Synchronous TrueCopy 9500V/ShadowImage 9500V functions are installed. Verify that the target volume is available.	237
EX_ENQVOL	Unmatched volume status within the group	The volume attribute or the fence level within a group is not identical.	Verify status using the pairdisplay command. Make sure all volumes in the group have the same fence level and volume attributes.	236
EX_EVOLCE	Pair Volume combination error	Combination of a volume is unsuitable between the remote and local host.	Verify volume status using the pairdisplay command, and change the combination of volumes properly.	235
EX_EWSUSE	Pair suspended at WAIT state	Detected a suspended status (PSUE) for the paired volume, before it achieves the designated status.	Determine how the pair was suspended. If the pair was forcibly suspended, resynchronize the pair. If the pair was not suspended forcibly, the pair was suspended due to an error in the subsystem. Collect all HORCM log and trace information, and call the Hitachi Data Systems support center. After maintenance personnel has recovered the subsystem, refer to section 5.5 to recover the data.	234
EX_EWSTOT	Timeout waiting for specified status	Detected a timeout before it achieved the designated status.	Increase the value of the timeout using the -t option.	233

Table 5.6 Command Error Messages (continued)

Message ID	Error Message	Condition	Recommended Action	Return Value
EX_EWSLTO	Timeout waiting for specified status on the local host	Timeout error because the remote did not note expected status.	Verify that HORC Manager on the remote host is running.	232
EX_ESTMON	HORCM Monitor stopped	HORC Manager monitoring was refused.	Verify the value of "poll" in the configuration file.	231
EX_UNWCMD	Unknown command	An unknown command was attempted.	Verify the command name.	230
EX_INCSTG	Inconsistent status in group	The pair status of a volume within a group is not identical to the status of the other volumes in the group.	Verify the pair status using the pairdisplay command.	229
EX_INVSTP	Invalid pair status	The pair status of the target volume is not appropriate.	Verify the pair status using the pairdisplay command.	228
EX_INVVOL	Invalid volume status	The volume status of the target volume is not appropriate.	Verify the pair status using the pairdisplay -l command.	222
EX_INVMUN	Invalid mu# with HORC or HOMRCF	The MU# of the volume to be operated is not appropriate.	Verify the MU# for the specified group using the pairdisplay command. MU #1/2 cannot be used for ShadowImage 9500V. MU # must be P-VOL and zero fixed for ShadowImage 9500V. Also verify whether the command execution environment is set as HORCC_MRCF 1.	220
EX_ENLDEV	No such LDEV within the RAID	A device defined in the configuration file does not have a mapping to a real LUN and Target ID within the RAID.	Please confirm that the Port, Target ID, LUN is defined correctly under HORCM_DEV in the configuration file.	227
EX_INVRCD	Invalid return code	Wrong return code.	Collect all HORCM log information, and call the Hitachi Data Systems Support Center.	226
EX_VOLCUR	S-Vol currency error	Currency check error for S-VOL. Cannot guarantee identical data on S-VOL.	Check the volume list to see if an operation was directed to the wrong S-VOL.	225
EX_VOLCUE	Local volume currency error	The volume specified with the SVOL-takeover command is not the same as the P-VOL.	Verify the pair status of the local volume.	224
EX_VOLCRE	Local and remote volume currency error	The combination of the volumes specified with Swap-takeover is unsuitable.	Verify the pair status of remote and local volumes using the pairdisplay command.	223
EX_UNWERR	Unknown error code	Wrong error code.	Collect all HORCM log information, and call the Hitachi Data Systems Support Center.	--
EN_PERM	Permission denied with the LDEV	The device described in the configuration definition file is not permitted for pair operation.	Execute the pairdisplay command or the raidscan – use the verify command to verify that the device is permitted for the pair operation.	213

Table 5.6 Command Error Messages (continued)

Message ID	Error Message	Condition	Recommended Action	Return Value
EX_ENQSIZ	Unmatched volume size for pairing	1. The number or the size of the volume between the local and remote volume is not appropriate. 2. The product ID on the local subsystem does not correspond with the product ID on the remote subsystem.	1. Check the volume size of the target pair using raidscan –f option, and then set the same size for local and remote volumes for the target pair. 2. Set the identical product ID in the local and remote subsystem.	212
EX_ERPERM	Permission denied with RAID	The volume on the configuration file is a volume that is not permitted to operate a pair in the subsystem.	Set HITACHI for the vendor ID.	211

The following list includes common error messages for the **horctakeover, paircurchk, paircreate**, **pairsplit**, **pairresync**, **pairevtwait**, **pairvolchk** and **pairdisplay** commands.

Table 5.7 Common Error Messages (1)

Type	Message ID	Error Message	Return Value
Syntax for Argument Unrecoverable	EX_REQARG	Required Arg list	254
	EX_INVARG	Invalid argument	253
	EX_INVNAM	Invalid name of option	249
	EX_UNWOPT	Unknown option	252
	EX_UNWCOD	Unknown function code	238
	EX_UNWCMD	Unknown command	230
	EX_ERANGE	Result too large	244
	EX_ENAMLG	File name too long	243
	EX_INVRCD	Invalid return code	226
Configuration Unrecoverable	EX_ENOGRP	No such group	239
	EX_ENOENT	No such device or group	247
	EX_ENODEV	No such device	246
	EX_ENLDEV	No such LDEV within the RAID	227
	EX_ENOUNT	No such RAID unit	219
	EX_INVMUN	Invalid mu# with HORC or HOMRCF	220
	EX_ENQSER	Unmatched Serial# vs RAID unit ID	218
	EX_ENPERM	Permission denied with the LDEV	213
	EX_ERPERM	Permission denied with RAID	211
Command I/O to RAID Recoverable	EX_CMDRJE	An order to the control/command device was rejected	221
	EX_CMDIOE	Control command I/O error Control command I/O error, or rejected	237
	EX_OPTINV	A specified option is invalid	248
	EX_INVMOD	Invalid RAID command mode	241
	EX_INVCMD	Invalid RAID command	240
Communication for HORCM Recoverable	EX_ATTHOR	Cannot attach to HORC Manager	251
	EX_ATTDBG	Cannot attach to a Debug layer	250
	EX_COMERR	Cannot Communicate with HORC Manager	255
	EX_ENORMT	No remote host alive for remote commands or Remote RAID Manager might be blocked (sleeping) on an existing I/O.	242
Resource Unrecoverable	EX_ENOMEM	Not enough core	245

Note: Unrecoverable indicates errors that cannot be recovered by executing the command again. Recoverable indicates errors that can be recovered by executing the command again.

The following list includes common error messages for the **raidscan**, **raidar**, **raidqry** and **horcctl** commands.

Table 5.8 Common Error Messages (2)

Type	Message ID	Error Message	Return Value
Syntax for Argument Unrecoverable	EX_REQARG	Required Arg list	254
	EX_INVARG	Invalid argument	253
	EX_INVNAM	Invalid name of option	249
	EX_UNWOPT	Unknown option	252
	EX_UNWCOD	Unknown function code	238
	EX_UNWCMD	Unknown command	230
	EX_ERANGE	Result too large	244
	EX_ENAMLG	File name too long	243
	EX_INVRCD	Invalid return code	226
Configuration	EX_ENOGRP	No such group	239
	EX_ENOENT	No such device or group	247
	EX_ENODEV	No such device	246
	EX_ENLDEV	No such LDEV within the RAID	227
	EX_ENOUNT	No such RAID unit	219
	EX_INVMUN	Invalid mu# with HORC or HOMRCF	220
	EX_ENQSER	Unmatched Serial# vs RAID unitID	218
	EN_PERM	Permission denied with the LDEV	213
	EX_ERPERM	Permission denied with RAID	211
Command I/O to RAID Recoverable	EX_CMDIOE	Control command I/O error Control command I/O error, or rejected	237
	EX_OPTINV	A specified option is invalid	248
	EX_INVMOD	Invalid RAID command mode	241
	EX_INVCMD	Invalid RAID command	240
Communication for HORCM Recoverable	EX_ATTHOR	Cannot attach to HORC Manager	251
	EX_ATTDBG	Cannot attach to a Debug layer	250
	EX_COMERR	Cannot Communicate with HORC Manager	255
	EX_ENORMT	No remote host alive for remote commands or RAID Manager might be blocked (sleeping) on an existing I/O	242
Resource Unrecoverable	EX_ENOMEM	Not enough core	245

Note: Unrecoverable indicates errors that cannot be recovered by executing the command again. Recoverable indicates errors that can be recovered by executing the command again.

The following list includes unique error messages for the **paircreate**, **pairsplit**, **pairresync**, **pairevtwait**, and **pairvolchk** commands.

Table 5.9 Unique Error Messages

Command	Type	Message ID	Error Message	Return Value
paircreate, pairsplit, pairresync, pairevtwait, pairvolchk, horctakeover	Volume status Unrecoverable	EX_ENQVOL	Unmatched volume status within the group	236
paircreate, pairsplit, pairresync, pairevtwait, horctakeover		EX_INCSTG	Inconsistent status in group	229
paircreate, pairsplit, pairresync, pairevtwait		EX_INVVOL	Invalid volume status	222
pairsplit, pairevtwait, pairvolchk, horctakeover		EX_EVOLCE	Pair Volume combination error	235
paircreate, pairsplit, pairresync		EX_INVSTP	Invalid pair status	228
paircurchk, horctakeover		EX_VOLCUR	S-VOL currency error	225
horctakeover		EX_VOLCUE	Local Volume currency error	224
horctakeover		EX_VOLCRE	Local and Remote Volume currency error	223
pairsplit, pairevtwait		EX_EWSUSE	Pair suspended at WAIT state	234
paircreate		EX_ENQSIZ	Unmatched volume size for pairing	212
pairevtwait, horctakeover	Timer Recoverable	EX_EWSTOT	Timeout waiting for specified status	233
pairevtwait		EX_EWSLTO	Timeout waiting for specified status on the local host	232

Note: Unrecoverable indicates errors that cannot be recovered by executing the command again. Recoverable indicates errors that can be recovered by executing the command again. However, EX_EWSTOT of the horctakeover command is excluded. When the command completed with an error, the error is recorded in the CCI command log ($HORCC_LOG) (see Appendix B).

5.5 Data Recovery Procedure for Suspended (PSUE) Status

When a pair was suspended (PSUE) due to a failure in the subsystem, collect the CCI system log files and call the Hitachi Data Systems Customer Support (see section 5.5.3).

Maintenance personnel will remove the source of error in the subsystem. After the source of the error has been removed, the system administrator must recover the pair using the CCI commands.

5.5.1 Recovering the Pair

Make sure that the system administrator has verified that the PSUE occurred while resynchronizing from S-VOL to P-VOL (pairresync –restore). Table 5.10 shows data assurance and the method for recovering the pair.

Table 5.10 Data Assurance and the Method for Recovering the Pair

State before PSUE	Data Assurance	Action to be taken after PSUE
Other than RCPY	P-VOL: Assured S-VOL: Not assured	Resynchronize (pairresync) in the direction of P-VOL to S-VOL. Note that the pair may have been split due to the drive's double-malfunction in either or both volumes. In such case, confirm that the data exists in the P-VOL, and then create a pair (paircreate).
RCPY	P-VOL: Not assured S-VOL: Not assured	Split the pair (pairsplit –S), restore the backup data to P-VOL, and then create a pair (paircreate). Note that the pair may have been split due to the drive's double-malfunction in either or both volumes. In such case, confirm that the backup data restoration has been completed to the P-VOL, and then create a pair.

5.5.2 Correspondence of Internal LU in the Subsystem and the Device Recognized by the Server

To verify the correspondence of the internal LU number in the subsystem and the device name recognized by the server, use the inqraid command tool or the raidscan command. The following example illustrates using the inqraid command tool in the HP-UX® system.

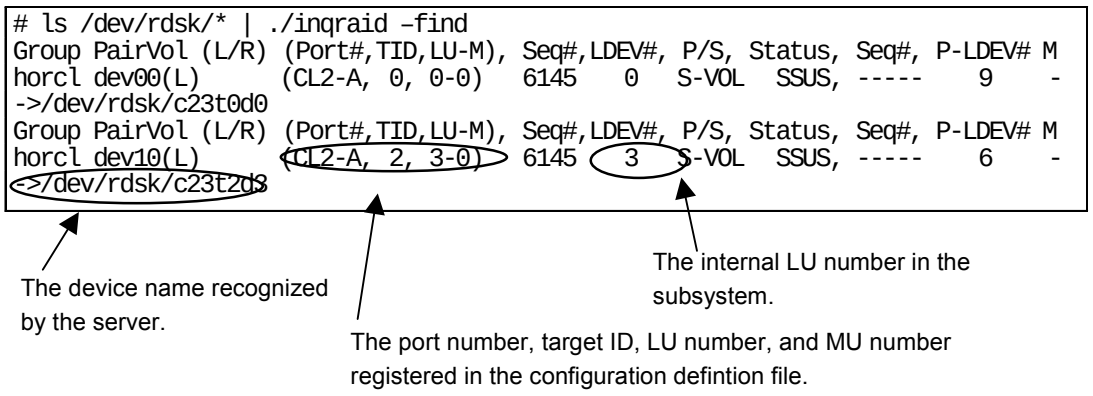

Figure 5.4 Correspondence between the Internal LU and the Device Recognized by the Server

For details on the inqraid command tool, refer to section 4.15.1. For details on raidscan command, refer to section 4.12.1.

5.5.3 Notes on Splitting the Pair before LU Format or during Reverse Pairresync (status change to PSUE)

When the host cannot recognize an LU while the protection function is ON, the following message is displayed and operations to the pair of the unrecognized LU cannot be performed.

```
[EX_ENPERM] Permission denied with the LDEV
```

Figure 5.5 Message Displayed when Host Cannot Recognize an LU

Issue the pairdisplay command and verify the pair status. An example is shown below.

```
C:¥horcm¥etc>pairdisplay –g vg01 -fc
Group PairVol (L/R) (Port#,TID,LU-M),Seq#,LDEV#, P/S, Status, %, P-LDEV# M
vg01   oradb1(L)    (CL1-A ,1, 1-0)  9876   1.. P-VOL PSUE,  0      2  -
vg01   oradb1(R)    (CL1-A ,1, 2-0)  9876 ****.. ---- ----,-----  ---- -
vg01   oradb2(L)    (CL1-A ,1, 3-0)  9876   3.. P-VOL PSUE,  0      2  -
vg01   oradb2(R)    (CL1-A ,1, 4-0)  9876 ****.. ---- ----,-----  ---- -
vg01   oradb3(L)    (CL1-A ,1, 5-0)  9876   5.. P-VOL PSUE,  0      2  -
vg01   oradb3(R)    (CL1-A ,1, 6-0)  9876 ****.. ---- ----,-----  ---- -
```

LUs not recognized by the host.

Figure 5.6 Pair Status Verification

After the pair status is verified, turn the protection function OFF, restart HORCM, and then split the pair (to turn the protection function OFF, delete the environment variable $HORCMPROMOD). For details on the environment variable, see section 4.2).

After splitting the pair and when the drive maintenance is done by the maintenance personnel, restore the backup data, make the host recognize the LU, and create the pair. And then turn the protection function ON and start HORCM.

5.6 Data Recovery Procedure for Suspended (PSUE) Status Using Synchronous TrueCopy 9500V

When a pair was suspended (PSUE) due to a failure in the subsystem, collect the CCI system log files and then call the Hitachi maintenance personnel The maintenance personnel will remove the error cause of the subsystem. After the error cause has been removed, the system administrator must recover the pair using the CCI commands.

5.6.1 Recovering the Pair

Resynchronization in the direction of P-VOL to S-VOL enables to recover a Synchronous TrueCopy 9500V pair. (The resynchronization cannot be performed in the direction of S-VOL to P-VOL.) However, when a pair was suspended (status = PSUE) without a failure in the system, once release the pair and create the pair again. Please recover the pair following the procedure for the recovery in the flow of Figure 5.2.

5.7 Calling the Hitachi Data Systems Support Center

If you need to call the Hitachi Data Systems Support Center, provide as much information about the problem as possible. Please include the circumstances surrounding the error or failure and the exact content of any error messages and/or codes displayed on the Resource Manager 9500V and/or logged at the host.

The worldwide Hitachi Data Systems Support Centers are:

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

Appendix A Acronyms and Abbreviations

CM	Cluster Manager
CVS	Custom Volume Size (also called Virtual LUN)
HA	High Availability
HORCM	HORC Manager
I/O	Input / Output
LDEV	Logical Device
LU	Logical Unit
LUN	Logical Unit Number
LVM	Logical Volume Manager
MU	Mirrored Unit
P-VOL	Primary Volume
SCSI	Small Computer System Interface
Synchronous TrueCopy 9500V	Synchronous Remote Copy
S-VOL	Secondary Volume
TID	Target ID

Appendix B Maintenance Logs and Tracing Functions

B.1 Log Files

The CCI software (HORCM) and Synchronous TrueCopy 9500V/ShadowImage 9500V commands maintain internal logs and traces which can be used to identify the source of errors and keep records of the status transition history of paired volumes. Figure B.1 displays the CCI logs and traces.

HORCM logs are classified into start-up logs and execution logs. The start-up logs contain data on errors that occur before the HORCM is ready to provide services. Thus, if the HORCM fails to start up due to improper environment settings, refer to the start-up logs to resolve the problem. The HORCM execution logs (error log, trace, and core files) contain data on errors that are caused by software or hardware problems. These logs contain internal error data, which does not apply to any user settings; therefore, users do not need to refer to the HORCM execution logs. When an error occurs during the execution of a command, data on the error is collected in the command log file. Users may refer to the command log file if a command execution error occurs.

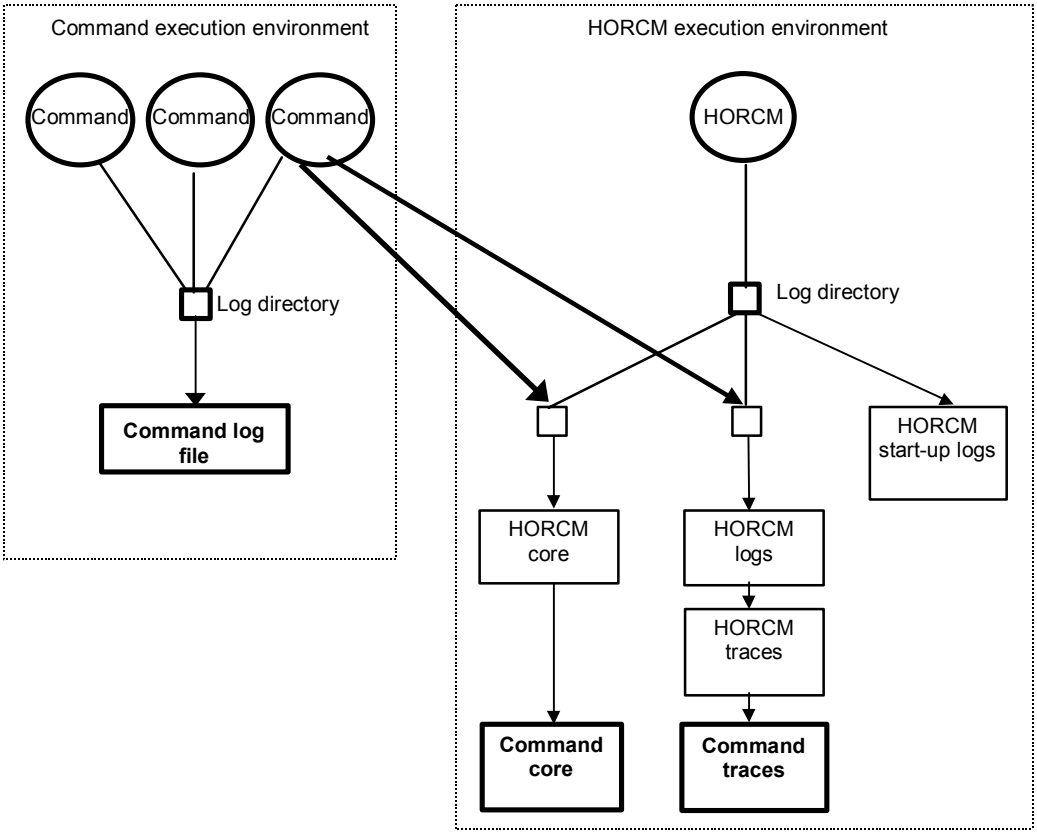

Figure B.1 Logs and Traces

The start-up log, error log, trace, and core files are stored as shown in Table B.1. Specify the directories for the HORCM and command log files using the HORCM_LOG and HORCC_LOG environmental variables as shown in Table B.2. If it is not possible to create the log files, or if an error occurs before the log files are created, the error logs are output in the system log file. If the HORCM activation fails, the system administrator should check the system log file, identify the error cause, and take the proper action. For details on message output to the system log file and the recommended actions for resolving the error conditions, see Chapter 5. The system log file for UNIX®-based systems is the syslog file. The system log file for Windows®-based systems is the event log file.

Table B.1 Log Files

File	UNIX®-based Systems	Windows®-based Systems
Start-up log	HORCM start-up log: $HORCM_LOG/horcm_HOST.log Command log: $HORCC_LOG/horcc_HOST.log	HORCM start-up log: $HORCM_LOG\horcm_HOST_log.txt Command log: $HORCC_LOG\horcc_HOST_log.txt
Error log	HORCM error log: $HORCM_LOG/horcmlog_HOST/horcm.log	HORCM error log: $HORCM_LOG\horcmlog_HOST\horcm_log.txt
Trace	HORCM trace: $HORCM_LOG/horcmlog_HOST/horcm_PID.trc Command trace: $HORCM_LOG/horcmlog_HOST/horcc_PID.trc	HORCM trace: $HORCM_LOG\horcmlog_HOST\horcm_PID_trc.txt Command trace: $HORCM_LOG\horcmlog_HOST\horcc_PID_trc.txt
Core	HORCM core: $HORCM_LOG/core_HOST_PID/core Command core: $HORCM_LOG/core_HOST_PID/core	HORCM core: $HORCM_LOG\core_HOST_PID\core Command core: $HORCM_LOG\core_HOST_PID\core

Note: HOST denotes the host name of the corresponding machine. PID denotes the process ID of that machine.

The location of the directory that contains the log file depends on the user's command execution environment and the HORCM execution environment. The command trace file and core file reside together under the directory specified in the HORCM execution environment. A directory specified using the environmental variable HORCM_LOG is used as the log directory in the HORCM execution environment. If no directory is specified, directory **/tmp** is used. A directory specified using the environmental variable HORCC_LOG is used as the log directory in the command execution environment. If no directory is specified, the directory **/HORCM/log*** is used (* = instance number). A nonexistent directory may be specified as a log directory using the environmental variable.

Table B.2 Log Directories

Directory Name	Definition
$HORCM LOG	A directory specified using the environmental variable **HORCM_LOG**. The HORCM log file, trace file, and core file as well as the command trace file and core file are stored in this directory. If no environmental variable is specified, **/HORCM/log/curlog** is used.
$HORCC LOG	A directory specified using the environmental variable **HORCC_LOG**. The command log file is stored in this directory. If no environmental variable is specified, the directory **/HORCM/log*** is used (* is the instance number). While the HORCM is running, the log files are stored in the **$HORCM_LOG** directory shown in (a). When the HORCM starts up, the log files created in the operation are stored automatically in the **$HORCM_LOGS** directory shown in (b). a. HORCM log file directory in operation $HORCM_LOG = /HORCM/log*/curlog (* is instance number) b. HORCM log file directory for automatic storing $HORCM_LOGS = /HORCM/log*/tmplog (* is instance number)

B.2 Trace Files

The command trace file is used for maintenance troubleshooting maintenance. It is not created normally. If the source of an error cannot be identified by means of the log file, the environmental variables or trace control commands with trace control parameters are issued to start tracing and the trace file is created. The trace control parameters consist of the trace level, file size, mode, etc. Increasing the trace level enables more detailed tracing. Tracing is made in wraparound within the range of the file size. The HORCM creates the trace file according to the trace level specified in the HORCM start-up shell script set to activate the HORCM.

B.3 Trace Control Command

The trace control command (one of the HORCM control commands) sets or changes the trace control parameters. This command is used for troubleshooting and maintenance. If no trace control parameters can be specified using the environmental variables in the user's command execution environment, it is possible to change the trace control parameters into the global parameters using this command. Table B.3 lists and describes the parameters of the trace control command.

Table B.3 Trace Command Parameters

Parameter	Function
Trace level parameter	Specifies the trace level, range = 0 to 15
Trace size parameter	Specifies the trace file size in kB
Trace mode parameter	Specifies the buffer mode or non-buffer mode for writing data in the trace file
Trace type parameter	Specifies the trace type defined internally
Trace change instruction	Specifies either the command or the HORCM (CCI instance) for which the trace control parameters are changed

Appendix C Uninstalling and Upgrading CCI

C.1 Uninstalling UNIX® CCI Software

After verifying that the CCI software is not running, you can uninstall the CCI software. If the CCI software is still running when you want to uninstall, shut down the CCI software using the **horcmshutdown.sh** command to ensure a normal end to all Synchronous TrueCopy 9500V/ShadowImage 9500V functions.

Caution: Before uninstalling the CCI software, make sure that all device pairs are in simplex status.

To uninstall the CCI software from a root directory (see Figure C.1): issue the uninstall command, go to the root directory, and delete the HORCM directory.

To uninstall the CCI software from a non-root directory (see Figure C.2): issue the uninstall command, go to the root directory, delete the HORCM link, and delete the HORCM directory.

```
#/HORCM/horcmuninstall.sh    ← Issue the uninstall command.
#cd  /                       ← Change directories.
#rm -rf /HORCM               ← Delete the CCI directory.
```

Figure C.1 Uninstalling the CCI Software from a Root Directory

```
#/HORCM/horcmuninstall.sh    ← Issue the uninstall command.
#cd  /          ← Change directories.
#rm    /HORCM   ← Delete the CCI link.
#rm  -rf   /non-root_directory_name/HORCM    ← Delete the CCI directory.
```

Figure C.2 Uninstalling the CCI Software from a Non-Root Directory

C.2 Upgrading Unix® CCI Software

After verifying that CCI is not running, you can upgrade the CCI software. If CCI is still running when you want to upgrade software versions, shut down the CCI software using the **horcmshutdown.sh** command to ensure a normal end to all Synchronous TrueCopy 9500V/ShadowImage 9500V functions. For instructions on upgrading the CCI software in a UNIX® environment, see Chapter 3.

C.3 Uninstalling Windows® NT/2000 CCI Software

After verifying that the CCI software is not running, you can uninstall the CCI software. If the CCI software is still running when you want to uninstall, shut down the CCI software using the horcmshutdown command to ensure a normal end to all Synchronous TrueCopy 9500V/ShadowImage 9500V functions.

Caution: Before uninstalling the CCI software, make sure that all device pairs are in simplex mode.

To uninstall the CCI software:

1. On the **Control** panel, select the **Add/Remove** programs option.
2. When the Add/Remove Program Properties panel opens, choose the **Install/Uninstall** tab and select **CCI** from the program products list.
3. Select the **Add/Remove** button to remove the CCI software.

C.4 Upgrading Windows NT®/2000 Software

After verifying that the CCI software is not running, you can upgrade the CCI software. If the CCI software is still running when you want to upgrade software versions, shut down the CCI software using the **horcmshutdown** command to ensure a normal end to all Synchronous TrueCopy 9500V/ShadowImage 9500V functions. To upgrade the CCI software:

1. On the Control panel, select the Add/Remove programs option.
2. When the Add/Remove Program Properties panel opens, choose the **Install/Uninstall** tab and select **CCI** from the program products list.
3. Select the **Add/Remove** button to remove the CCI software.
4. Insert the installation medium in the proper I/O device, and from the **Start** menu choose the **Run** command.
5. When the Run window opens, enter **x:\Setup.exe** (where x: is a CD drive) in the **Open** pull down list box.
6. An InstallShield will open. Follow the on screen instructions to install the CCI software.
7. Reboot the Windows® NT/2000 server and verify that the correct version of the CCI software is running on your system by executing the **raidqry -h** command.

Appendix D Fibre-to-SCSI Address Conversion

Fibre channel physical addresses are converted to SCSI target Ids (TIDs) using a conversion table. Table D.1 displays the current limits for SCSI TIDs on various operating systems.

Table D.1 Limits for Target IDs

	HP-UX® and other Systems		Solaris®, IRIX® Systems		Windows® Systems	
Port	TID	LUN	TID	LUN	TID	LUN
Fibre	0 to 15	0 to 511	0 to 125	0 to 511	0 to 31	0 to 511

An example of using the **raidscan** command to display the TID and LUN of Harddisk6 (HP® system) is displayed in Figure D.1.

Note: You must start HORCM without descriptions of HORCM_DEV or HORCM_INST in the configuration definition file because of the unknown target IDs and LUNs.

```
C:\>raidscan -pd hd6 -x drivescan hd6
Harddisk 6... Port[ 2] PhId[ 4] TId[ 3] Lun[  5] [HITACHI     ] [DF600F         ]
              Port[CL1-A] Ser#[   3005] LDEV#[  14(0x00E)]
              HORC = NONE  HOMRCF[MU#0 = SMPL MU#1 = NONE MU#2 = NONE]
              RAID5[Group  1- 0] SSID = 0x0000
PORT# /ALPA/C,TID#,LU#.Num(LDEV#....)...P/S, Status,Fence,LDEV#,P-Seq#,P-LDEV#
CL1-A / e2/ 4,  29,  0.1(9).............SMPL  ----  ------ ----, -----  ----
CL1-A / e2/ 4,  29,  1.1(10)............SMPL  ----  ------ ----, -----  ----
CL1-A / e2/ 4,  29,  2.1(11)............SMPL  ----  ------ ----, -----  ----
CL1-A / e2/ 4,  29,  3.1(12)............SMPL  ----  ------ ----, -----  ----
CL1-A / e2/ 4,  29,  4.1(13)............SMPL  ----  ------ ----, -----  ----
CL1-A / e2/ 4,  29,  5.1(14)............SMPL  ----  ------ ----, -----  ----
CL1-A / e2/ 4,  29,  6.1(15)............SMPL  ----  ------ ----, -----  ----
Specified device is LDEV# 0014
```

Figure D.1 Using Raidscan to Display TID and LUN for Fibre-Channel Devices

In this case, the target ID indicated by the raidscan command must be used in the configuration definition file. This can be done using either of the following two methods:

- **Using default conversion table**. The TID# and LU# indicated by the raidscan command are used in the HORCM configuration definition file.
- **Changing default conversion table**. The HORCMFCTBL environmental variable enables you to change the default conversion table as shown in Figure D.2.

```
C:\> set HORCMFCTBL=X                          'X' is fibre conversion table number.
C:\> horcmstart ...                            Start of  HORCM.
:
:
Result of "set HORCMFCTBL=X" command:
C:\>raidscan -pd hd6 -x drivescan hd6
Harddisk 6... Port[ 2] PhId[ 4] TId[ 3] Lun[  5] [HITACHI     ] [DF600F        ]
              Port[CL1-A] Ser#[   3005] LDEV#[  14(0x00E)]
              HORC = NONE  HOMRCF[MU#0 = SMPL MU#1 = NONE MU#2 = NONE]
              RAID5[Group  1- 0] SSID = 0x0000
PORT# /ALPA/C,TID#,LU#.Num(LDEV#....)...P/S,Status,Fence,LDEV#,P-Seq#,P-LDEV#
CL1-A / e2/ 0,   3,  0.1(9).............SMPL  ----  ------ ----, -----  ----
CL1-A / e2/ 0,   3,  1.1(10)............SMPL  ----  ------ ----, -----  ----
CL1-A / e2/ 0,   3,  2.1(11)............SMPL  ----  ------ ----, -----  ----
CL1-A / e2/ 0,   3,  3.1(12)............SMPL  ----  ------ ----, -----  ----
CL1-A / e2/ 0,   3,  4.1(13)............SMPL  ----  ------ ----, -----  ----
CL1-A / e2/ 0,   3,  5.1(14)............SMPL  ----  ------ ----, -----  ----
CL1-A / e2/ 0,   3,  6.1(15)............SMPL  ----  ------ ----, -----  ----
Specified device is LDEV# 0014
```

Figure D.2 Using HORCMFCTBL to Change the Default Fibre Conversion Table

Table D.2, Table D.3, and Table D.4 show the fibre address conversion tables:

- Table number 0 = HP-UX® systems
- Table number 1 = Solaris® systems
- Table number 2 = Windows NT®/2000 systems

Note 1: If the TID displayed on the system is different than the TID indicated in the fibre address conversion table, you must use the TID (and LU#) returned by the **raidscan** command to specify the device(s).

Note 2: The conversion table for Windows NT®/2000 is based on the Emulex driver. If a different fibre-channel adapter is used, the target ID indicated by the **raidscan** command may be different than the target ID indicated by the Windows NT®/2000 system. In such case, for the configuration definition file, use the target ID that is displayed (obtained) by the raidscan –find command.

Note 3: The conversion table for Native Fibre is used when the FC_AL conversion for the host is unknown, or when the FC_AL conversion is the device file displayed in LUN as Fabric mode. Only LUN is displayed and the target ID (displayed as zero) will not be used. Therefore, there is no table for Native Fibre since there is no conversion to target ID.

Table D.2 Fibre Address Conversion Table for HP-UX® Systems (Table 0)

C0		C1		C2		C3		C4		C5		C6		C7	
AL-PA	TID	AL-PA	TID	AL-PA	TID	AL-PA	TID	AL-PA	TID	AL-PA	TID	AL-PA	TID	AL-PA	TID
EF	0	CD	0	B2	0	98	0	72	0	55	0	3A	0	25	0
E8	1	CC	1	B1	1	97	1	71	1	54	1	39	1	23	1
E4	2	CB	2	AE	2	90	2	6E	2	53	2	36	2	1F	2
E2	3	CA	3	AD	3	8F	3	6D	3	52	3	35	3	1E	3
E1	4	C9	4	AC	4	88	4	6C	4	51	4	34	4	1D	4
E0	5	C7	5	AB	5	84	5	6B	5	4E	5	33	5	1B	5
DC	6	C6	6	AA	6	82	6	6A	6	4D	6	32	6	18	6
DA	7	C5	7	A9	7	81	7	69	7	4C	7	31	7	17	7
D9	8	C3	8	A7	8	80	8	67	8	4B	8	2E	8	10	8
D6	9	BC	9	A6	9	7C	9	66	9	4A	9	2D	9	0F	9
D5	10	BA	10	A5	10	7A	10	65	10	49	10	2C	10	08	10
D4	11	B9	11	A3	11	79	11	63	11	47	11	2B	11	04	11
D3	12	B6	12	9F	12	76	12	5C	12	46	12	2A	12	02	12
D2	13	B5	13	9E	13	75	13	5A	13	45	13	29	13	01	13
D1	14	B4	14	9D	14	74	14	59	14	43	14	27	14		
CE	15	B3	15	9B	15	73	15	56	15	3C	15	26	15		

Note: AL-PA is an abbreviation for Arbitrated Loop Physical Address, and indicates the physical address for Fibre. TID indicates the target ID.

Table D.3 Fibre Address Conversion Table for Solaris® and IRIX® Systems (Table 1)

C0		C1		C2		C3		C4		C5		C6		C7	
AL-PA	TID	AL-PA	TID	AL-PA	TID	AL-PA	TID	AL-PA	TID	AL-PA	TID	AL-PA	TID	AL-PA	TID
EF	0	CD	16	B2	32	98	48	72	64	55	80	3A	96	25	112
E8	1	CC	17	B1	33	97	49	71	65	54	81	39	97	23	113
E4	2	CB	18	AE	34	90	50	6E	66	53	82	36	98	1F	114
E2	3	CA	19	AD	35	8F	51	6D	67	52	83	35	99	1E	115
E1	4	C9	20	AC	36	88	52	6C	68	51	84	34	100	1D	116
E0	5	C7	21	AB	37	84	53	6B	69	4E	85	33	101	1B	117
DC	6	C6	22	AA	38	82	54	6A	70	4D	86	32	101	18	118
DA	7	C5	23	A9	39	81	55	69	71	4C	87	31	103	17	119
D9	8	C3	24	A7	40	80	56	67	72	4B	88	2E	104	10	120
D6	9	BC	25	A6	41	7C	57	66	73	4A	89	2D	105	0F	121
D5	10	BA	26	A5	42	7A	58	65	74	49	90	2C	106	08	122
D4	11	B9	27	A3	43	79	59	63	75	47	91	2B	107	04	123
D3	12	B6	28	9F	44	76	60	5C	76	46	92	2A	108	02	124
D2	13	B5	29	9E	45	75	61	5A	77	45	93	29	109	01	125
D1	14	B4	30	9D	46	74	62	59	78	43	94	27	110		
CE	15	B3	31	9B	47	73	63	56	79	3C	95	26	111		

Note: AL-PA is an abbreviation for Arbitrated Loop Physical Address, and indicates the physical address for Fibre. TID indicates the target ID.

Table D.4 Fibre Address Conversion Table for Windows NT®/2000 Systems (Table 2)

C5(Phld5)				C4(Phld4)				C3(Phld3)				C2(Phld2)				C1(Phld1)			
AL-PA	TID	AL-PA	TID	AL-PA	TID	AL-PA	TID	AL-PA	TID	AL-PA	TID	AL-PA	TID	AL-PA	TID	AL-PA	TID	AL-PA	TID
						CC	15			98	15			56	15			27	15
				E4	30	CB	14	B1	30	97	14	72	30	55	14	3C	30	26	14
				E2	29	CA	13	AE	29	90	13	71	29	54	13	3A	29	25	13
				E1	28	C9	12	AD	28	8F	12	6E	28	53	12	39	28	23	12
				E0	27	C7	11	AC	27	88	11	6D	27	52	11	36	27	1F	11
				DC	26	C6	10	AB	26	84	10	6C	26	51	10	35	26	1E	10
				DA	25	C5	9	AA	25	82	9	6B	25	4E	9	34	25	1D	9
				D9	24	C3	8	A9	24	81	8	6A	24	4D	8	33	24	1B	8
				D6	23	BC	7	A7	23	80	7	69	23	4C	7	32	23	18	7
				D5	22	BA	6	A6	22	7C	6	67	22	4B	6	31	22	17	6
				D4	21	B9	5	A5	21	7A	5	66	21	4A	5	2E	21	10	5
				D3	20	B6	4	A3	20	79	4	65	20	49	4	2D	20	0F	4
				D2	19	B5	3	9F	19	76	3	63	19	47	3	2C	19	08	3
				D1	18	B4	2	9E	18	75	2	5C	18	46	2	2B	18	04	2
		EF	1	CE	17	B3	1	9D	17	74	1	5A	17	45	1	2A	17	02	1
		E8	0	CD	16	B2	0	9B	16	73	0	59	16	43	0	29	16	01	0

Note: AL-PA is an abbreviation for Arbitrated Loop Physical Address, and indicates the physical address for Fibre. TID indicates the target ID.

Appendix E CCI Operation on Windows NT®/2000 Systems

E.1 LDM Volume Search and Flush

Windows® 2000 supports a logical device manager called LDM. A logical drive is usually linked to LDM volumes (e.g. \Device\HarddiskVolumeX). Therefore, it is not easy to link the physical volume in the subsystem with LDM volume. Linking physical volumes in the subsystem with LDM volume is necessary when creating the configuration definition file, as shown in Figure E.1.

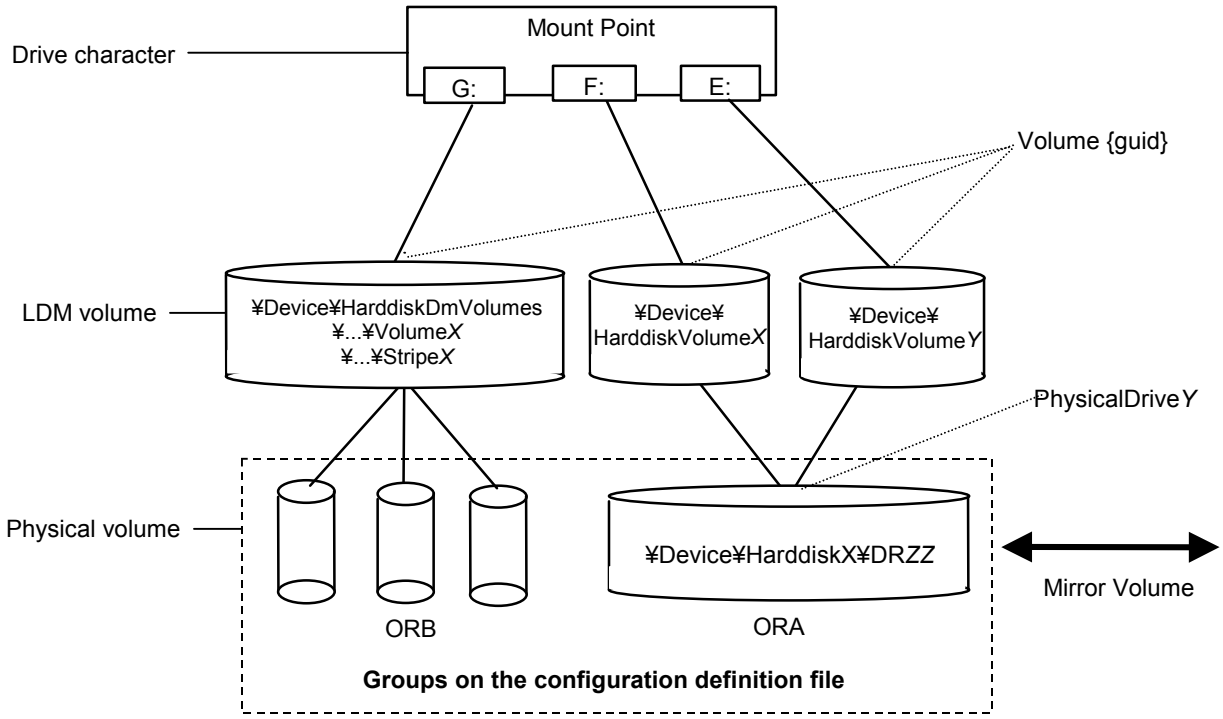

Figure E.1 LDM Volume Configuration

E.1.1 LDM Volume Search

CCI provides you with a search function that shows the relationship between the physical volumes in the subsystem and the LDM volumes. There are three types of volume searches:

- Physical

  Using $Physical as a search condition, the CCI displays the relation between the physical drive and the LDEV configuration of the subsystem.

- LDM Volume

  Using $Volume as a search condition, the CCI displays the relation between the LDM volume/physical drive and the LDEV configuration of the subsystem. Note that Windows NT does not support LDM volumes. For Windows NT systems, use $LETALL for the search condition instead of $Volume.

- Logical Device (drive character)

Using $LETALL as a search condition, the CCI displays the relation between the logical device and directory-mounted volume/LDM volume/physical drive and the LDEV configuration of the subsystem.

Note: The search condition ($Physical, #Volume, #LETALL) can be used for raidscan – **find** command, **inqraid** command, and **mkconf** command.

Also, the search condition ($LETALL) for searching the directory-mounted volumes can be used for raidscan —find command and mkconf command.

In the Windows NT/2000 systems, the DOS devices (e.g. C: Volume{guid}) are related to the device object name (\Device\...). The CCI changes the long name of these device objects to a shorter name shown below.

- **For Windows NT**
  - Device object name for Windows NT partition
    \Device\Harddisk*X*\Partition*Y* → **\Dsk*X*\p*Y***
  - Device object name for Windows NT physical drive
    \Device\Harddisk*X*\Partition0 → **Harddisk*X***

  Use the search condition ($Physical, $LETALL) for the **inqraid** command as shown below. You can see the relation between the logical device/LDM volume/physical drive and the LDEV configuration of the subsystem.

```
inqraid $Phy -CLI
DEVICE_FILE     PORT    SERIAL  LDEV CTG  H/M/12  SSID R:Group PRODUCT_ID
Harddisk0       -         -       -   -     -       -     -    DDRS-34560D
Harddisk1       CL2-A   6145      7   -   s/s/ss  0000 1:01-00 DF600F
```

Figure E.2 Inqraid Command using $Physical for Search Condition (Windows NT®)

```
inqraid $LETALL –CLI
DEVICE_FILE     PORT    SERIAL  LDEV CTG  H/M/12  SSID R:Group PRODUCT_ID
D:\Dsk0\p1      -         -       -   -     -       -     -    DDRS-34560D
E:\Dsk1\p1      CL2-A   6145      7   -   s/s/ss  0000 1:01-00 DF600F
F:\Dsk1\p2      CL2-A   6145      7   -   s/s/ss  0000 1:01-00 DF600F
```

Figure E.3 Inqraid Command using $LETALL for Search Condition (Windows NT®)

- **For Windows® 2000**
  - LDM device object name for Windows® 2000
    \Device\HarddiskVolume*X* (partition volume) → **\Vol*X*\Dsk*Y***
    Dsk*Y* indicates that Vol*X* is configured by Harddisk*Y*.
    \Device\HarddiskDmVolumes\...\Volume (span volume) → **\Dms*X*\Dsk*Y*s**
    \Device\HarddiskDmVolumes\...\Stripe*X* (stripe volume) → **\Dmt*X*\Dsk*Y*s**
    \Device\HarddiskDmVolumes\…\Raid*X* (RAID 5 volume) → **\Dmt*X*\Dsk*Y*s**
    Dsk*Y*s indicates that Dms*X* (Dmt*X*, Dmr*X*) is configured by several Harddisk*Y1*, *Y2*....
  - Device object name for Windows 2000 physical drive
    \Device\Harddisk*X*\DRZZ → **Harddisk*X***

  Use the search condition (&Physical, $Volume, $LETALL) for the **inqraid** command as shown below. You can see the relation between the logical device and directory-mounted volume/LDM volume/physical drive and the LDEV configuration of the subsystem.

```
inqraid $Phy -CLI
DEVICE_FILE     PORT    SERIAL  LDEV CTG  H/M/12  SSID R:Group PRODUCT_ID
Harddisk0       CL2-A    6145     7    -  s/s/ss  0000 1:01-00 DF600F
Harddisk1       CL2-A    6145     8    -  s/s/ss  0000 1:01-00 DF600F
Harddisk2       CL2-A    6145     9    -  s/s/ss  0000 1:01-00 DF600F
Harddisk3       CL2-A    6145    10    -  s/s/ss  0000 1:01-00 DF600F
```

Figure E.4 Inqraid Command using $Physical for Search Condition (Windows® 2000)

```
inqraid $Volume -CLI
DEVICE_FILE     PORT    SERIAL  LDEV CTG  H/M/12  SSID R:Group PRODUCT_ID
\Vol44\Dsk0     CL2-A    6145     7    -  s/s/ss  0000 1:01-00 DF600F
\Vol45\Dsk0     CL2-A    6145     7    -  s/s/ss  0000 1:01-00 DF600F
\Dmt1\Dsk1      CL2-A    6145     8    -  s/s/ss  0000 1:01-00 DF600F
\Dmt1\Dsk2      CL2-A    6145     9    -  s/s/ss  0000 1:01-00 DF600F
\Dmt1\Dsk3      CL2-A    6145    10    -  s/s/ss  0000 1:01-00 DF600F
```

Figure E.5 Inqraid Command using $Volume for Search Condition (Windows® 2000)

```
inqraid $LETALL –CLI
DEVICE_FILE     PORT    SERIAL  LDEV CTG  H/M/12  SSID R:Group PRODUCT_ID
E:\Vol44\Dsk0   CL2-A    6145     7   -  s/s/ss  0000 1:01-00 DF600F
F:\Vol45\Dsk0   CL2-A    6145     7   -  s/s/ss  0000 1:01-00 DF600F
G:\Dmt1\Dsk1    CL2-A    6145     8   -  s/s/ss  0000 1:01-00 DF600F
G:\Dmt1\Dsk2    CL2-A    6145     9   -  s/s/ss  0000 1:01-00 DF600F
G:\Dmt1\Dsk3    CL2-A    6145    10   -  s/s/ss  0000 1:01-00 DF600F
```

Figure E.6 Inqraid Command using $LETALL for Search Condition (Windows® 2000)

If you want to know the relationship between the LDM volumes and the groups in the configuration definition file, use the search conditions ($Physical, $Volume $LETALL) in the raidscan –find verify command. You can see the relation between the logical drive/LDM volume/physical drive and the groups.

```
raidscan -pi $LETALL -find verify
DEVICE_FILE          Group    PairVol     PORT   TARG  LUN M   SERIAL  LDEV
E:\Vol44\Dsk0        ORA      ORA_000     CL2-A     7    2 -    6145   7
F:\Vol45\Dsk0        ORA      ORA_000     CL2-A     7    2 -    6145   7
G:\Dmt1\Dsk1         ORB      ORB_000     CL2-A     7    4 -    6145   8
G:\Dmt1\Dsk2         ORB      ORB_001     CL2-A     7    5 -    6145   9
G:\Dmt1\Dsk3         ORB      ORB_002     CL2-A     7    6 -    6145   10
```

Figure E.7 Raidscan –find verify Command using $LETALL for Search Condition

```
raidscan -pi $LETALL –find
DEVICE_FILE         UID  S/F PORT   TARG  LUN    SERIAL  LDEV  PRODUCT_ID
E:\Vol44\Dsk0         0   F  CL2-A    7    2      6145    7    DF600F
F:\Vol45\Dsk0         0   F  CL2-A    7    2      6145    7    DF600F
G:\Dmt1\Dsk1          0   F  CL2-A    7    4      6145    8    DF600F
G:\Dmt1\Dsk2          0   F  CL2-A    7    5      6145    9    DF600F
G:\Dmt1\Dsk3          0   F  CL2-A    7    5      6145   10    DF600F
```

Figure E.8 Raidscan –find Command using $LETALL for Search Condition

The following is an example of LDM volume that is mounted by using a directory (hd1, hd2) that exists in D drive.

```
D:\HORCM\etc>inqraid $LETALL –CLI
DEVICE_FILE     PORT    SERIAL  LDEV CTG  H/M/12  SSID R:Group PRODUCT_ID
D:\Vol2\Dsk7     -         -      -   -     -      -     -     DDRS-34560D
D:\hd1\Vol8\Dsk0 CL2-B   6145    48   -  s/s/ss  0000 1:01-00 DF600
D:\hd2\Vol9\Dsk1 CL2-B   6145    49   -  s/s/ss  0000 1:01-00 DF600
G:\Dms1\Dsk2    CL2-A    6145    56   -  s/s/ss  0000 1:01-00 DF600
G:\Dms1\Dsk3    CL2-A    6145    57   -  s/s/ss  0000 1:01-00 DF600
G:\Dms1\Dsk4    CL2-A    6145    58   -  s/s/ss  0000 1:01-00 DF600
```

Figure E.9 Inqraid Command using $LETALL for Search Condition (Windows® 2000)

The directory-mounted volume can be operated using –x sync subcommand, -x mount subcommand, and –x umount subcommand that are embedded in CCI commands.

E.1.2 Mountvol Command on Windows® 2000

Note that the **mountvol** command attached on Windows® 2000 does not flush the NT file system buffer to the corresponding specified drive. Therefore, you cannot flush an unwritten data to P-VOL nor browse S-VOL by using this command.

The mountvol command displays the mounted volume in "\\?\Volume{*XXXX*}\" format.

```
mountvol
Creates, deletes, or lists a volume mount point.
.
.
MOUNTVOL [drive:]path VolumeName
MOUNTVOL [drive:]path /D
MOUNTVOL [drive:]path /L

    \\?\Volume{56e4954a-28d5-4824-a408-3ff9a6521e5d}\
        G:\
    \\?\Volume{bf48a395-0ef6-11d5-8d69-00c00d003b1e}\
        F:\
```

Figure E.10 Mountvol Command Displaying the Mounted Volumes

You can use the inqraid command or the raidscan command to see the relationship between the device object name and the physical drive of the \\?\Volume{*XXXX*}\.

```
inqraid $Volume{bf48a395-0ef6-11d5-8d69-00c00d003b1e} –CLI
DEVICE_FILE     PORT    SERIAL  LDEV CTG  H/M/12  SSID R:Group PRODUCT_ID
\Vol46\Dsk1     CL2-A    6145   6     -   S/s/ss  0000 1:01-00 DF600F
```

Figure E.11 Inqraid Command Displaying the Mounted Volumes

```
raidscan -pi $Volume{bf48a395-0ef6-11d5-8d69-00c00d003b1e} –find
DEVICE_FILE     UID  S/F PORT   TARG  LUN    SERIAL  LDEV  PRODUCT_ID
\Vol46\Dsk1      0    F  CL2-A   7     1     6145       6  DF600F
```

Figure E.12 Raidscan Command Displaying the Mounted Volumes

E.1.3 Flushing System Buffer

There are two methods for sending (flushing) unwritten data remaining on Windows NT® system buffer files to logical drive (physical disk).

First method: Use the –x sync option to directly specify the logical drive. You must know the logical drive that corresponds to the groups on the configuration definition file before executing the –x sync option. If the logical drive is mounted on each directory, you need to know the mounted volume names also.

Second method: Search the logical drive that corresponds to the groups on the configuration definition file, then send (flush) the data on the system buffer file to the logical drive (physical disk). This method simplifies the first method, and is provided by the raidscan –find sync command. The volumes can be flushed to the directory since it does not rely on the mount point.

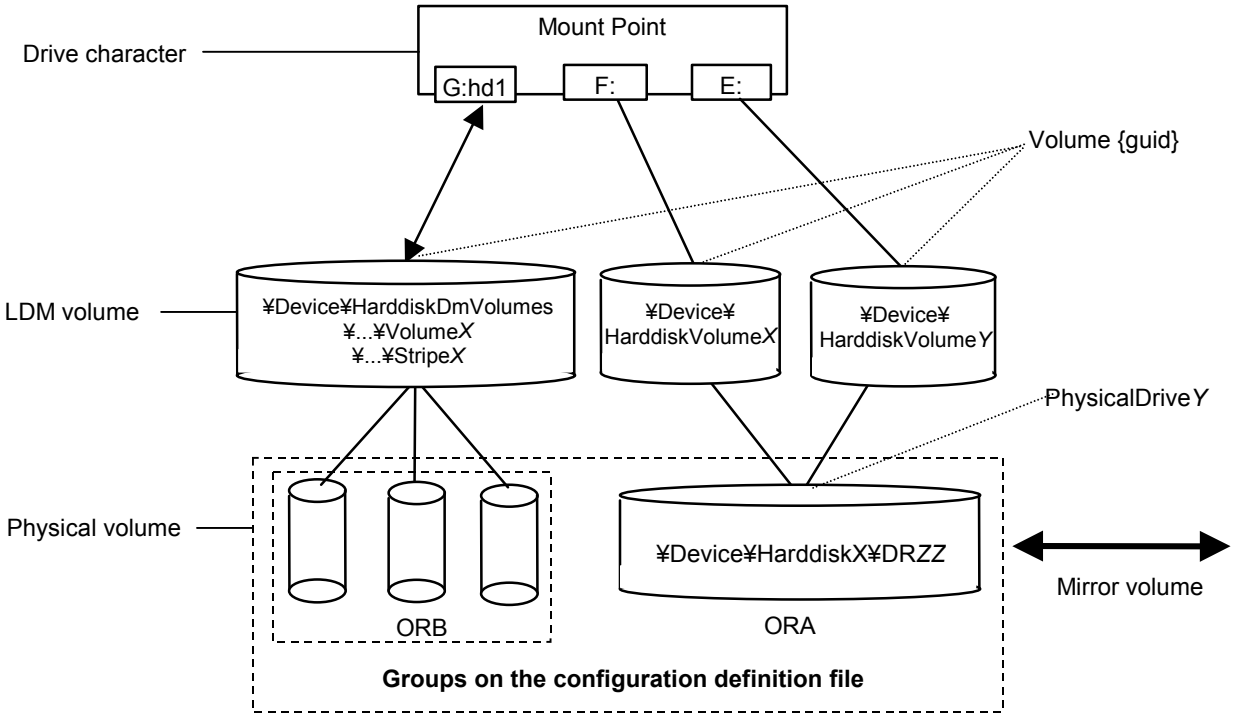

Figure E.13 LDM Volume Flush

The following is an example of flushing the system buffer that corresponds to group ORB in the configuration definition file.

```
raidscan -pi $Volume -find sync -g ORB
[SYNC] : ORB ORB_000[-] -> \Dmt1\Dsk1 : Volume{bf48a395-0ef6-11d5-8d69-00c00d003b1e}
[SYNC] : ORB ORB_001[-] -> \Dmt1\Dsk2 : Volume{bf48a395-0ef6-11d5-8d69-00c00d003b1e}
[SYNC] : ORB ORB_002[-] -> \Dmt1\Dsk3 : Volume{bf48a395-0ef6-11d5-8d69-00c00d003b1e}
```

Figure E.14 Flushing System Buffer Example (1)

The following is an example of flushing the system buffer that corresponds to all groups in the CCI local instance.

```
raidscan -pi $Volume -find sync
[SYNC] : ORA ORA_000[-] -> \Vol44\Dsk0 : Volume{56e4954a-28d5-4824-a408-3ff9a6521e5d}
[SYNC] : ORA ORA_000[-] -> \Vol45\Dsk0 : Volume{56e4954a-28d5-4824-a408-3ff9a6521e5e}
[SYNC] : ORB ORB_000[-] -> \Dmt1\Dsk1  : Volume{bf48a395-0ef6-11d5-8d69-00c00d003b1e}
[SYNC] : ORB ORB_001[-] -> \Dmt1\Dsk2  : Volume{bf48a395-0ef6-11d5-8d69-00c00d003b1e}
[SYNC] : ORB ORB_002[-] -> \Dmt1\Dsk3  : Volume{bf48a395-0ef6-11d5-8d69-00c00d003b1e}
```

Figure E.15 Flushing System Buffer Example (2)

Note: The Windows NT system does not support LDM volumes. Use $LETALL for search conditions, instead of $Volume.

- Instantaneous Offline Backup on Windows NT

The following shows an example of flushing the system buffer that corresponds to the logical device of group ORB in the configuration definition file using raidscan –find sync command (rather than –x mount or –x umount).

Table E.1 Instantaneous Offline Backup on Windows NT®

Step	P-VOL Side	S-VOL Side
1	The application closes all logical devices on P-VOL.	—
2	Flushes the system buffer that corresponds to the ORB (P-VOL) raidscan –pi $LETALL –find sync –g ORB	—
3	Splits the pair with Read/Write enabled pairsplit –g ORB	—
4	The application opens all logical drives on P-VOL, and restarts.	—
5	—	Executes the S-VOL backup
6		Flushes the system buffer that corresponds to ORB (updated S-VOL) after the backup is completed, or just before the paired volume is resynchronized raidscan –pi $LETALL –find sync –g ORB
7	Resynchronizes the paired volume pairresync –g ORB	

Note: The logical device of group ORB in S-VOL must be closed.

- Instantaneous Offline Backup on Windows 2000

The following shows an example of flushing the system buffer that corresponds to Volume{guid} of group ORB in the configuration definition file using raidscan -find sync command (rather than -x mount or -x umount).

Table E.2 Instantaneous Offline Backup on Windows® 2000

Step	P-VOL Side	S-VOL Side
1	The application closes all logical devices on P-VOL.	—
2	Flushes the system buffer that corresponds to the ORB (P-VOL) raidscan –pi $Volume –find sync –g ORB	—
3	Splits the pair with Read/Write enabled pairsplit –g ORB	—
4	The application opens all logical devices on P-VOL, and restarts.	—
5	—	Flushes the system buffer that corresponds to ORB (updated S-VOL) when splitting the paired volume completes and when the backup starts raidscan –pi $Volume –find sync –g ORB
6	—	Executes the S-VOL backup
7	—	Flushes the system buffer that corresponds to ORB (updated S-VOL) after the backup is completed, or just before the paired volume is resynchronized raidscan –pi $Volume –find sync –g ORB
8	Resynchronizes the paired volume pairresync –g ORB	—

Note: The logical device of group ORB in S-VOL must be closed.

- Online Backup on Windows NT®

The following shows an example of flushing the system buffer that corresponds to the logical device of group ORB in the configuration definition file using raidscan -find sync command (rather than -x mount or -x umount).

Table E.3 Online Backup on Windows NT®

Step	P-VOL Side	S-VOL Side
1	The application stops all Write accesses on P-Vol and freezes the database.	—
2	Flushes the system buffer that corresponds to the ORB (P-VOL) raidscan –pi $LETALL –find sync –g ORB	—
3	Splits the pair with Read/Write enabled pairsplit –g ORB	—
4	The application de-freezes the database and enables the Write access.	—
5	—	Executes the S-VOL backup
6		Flushes the system buffer that corresponds to ORB (updated S-VOL) after the backup is completed, or just before the paired volume is resynchronized raidscan –pi $LETALL –find sync –g ORB
7	Resynchronizes the paired volume pairresync –g ORB	

Note: For the logical device of group ORB in P-VOL, all the write access on P-VOL must be stopped before issuing the raidscan –find sync command. The logical device of group ORB in S-VOL must be closed before issuing the raidscan –find sync command.

- Online Backup on Windows® 2000

The following shows an example of flushing the system buffer that corresponds to Volume{guid} of group ORB in the configuration definition file using raidscan -find sync command (rather than -x mount or -x umount).

Table E.4 Online Backup on Windows® 2000

Step	P-VOL Side	S-VOL Side
1	The application stops all Write accesses on P-Vol and freezes the database.	—
2	Flushes the system buffer that corresponds to the ORB (P-VOL) raidscan –pi $Volume –find sync –g ORB	—
3	Splits the pair with Read/Write enabled pairsplit –g ORB	—
4	The application opens all logical devices on P-VOL, and restarts.	—
5	—	Flushes the system buffer that corresponds to ORB (updated S-VOL) when splitting the paired volume completes and when the backup starts raidscan –pi $Volume –find sync –g ORB
6	—	Executes the S-VOL backup
7	—	Flushes the system buffer that corresponds to ORB (updated S-VOL) after the backup is completed, or just before the paired volume is resynchronized raidscan –pi $Volume –find sync –g ORB
8	Resynchronizes the paired volume pairresync –g ORB	—

Note: For the logical device of group ORB in P-VOL, all the write access on P-VOL must be stopped before issuing the raidscan –find sync command. The logical device of group ORB in S-VOL must be closed before issuing the raidscan –find sync command.

E.2 Signature Changing Feature for Windows® 2000 Systems

The following configuration has two nodes, which MSCS on P-VOL is shared by node 1 and node 2, and the S-VOL is used for backing up node2. If the node 2 is rebooted while the status of node 1 is active, the MSCS on node 2 will incorrectly recognize the S-VOL as a volume under MSCS. Therefore, as a result, a phenomenon in which the backup application may not be able to use the S-VOL occurs.

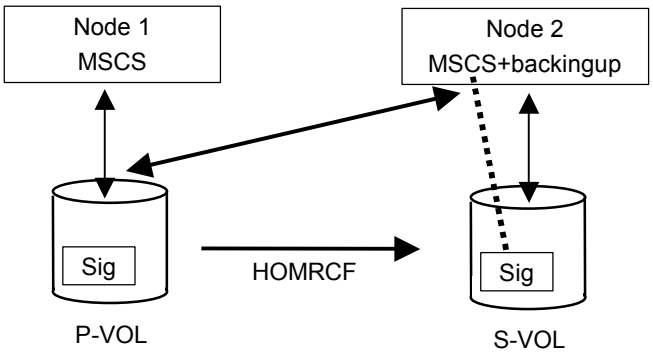

Figure E.14 Signature Changing Feature for Windows® 2000 Systems

This phenomenon occurs because when copying is processed from P-VOL to S-VOL, the original signature of S-VOL will change to the same signature of P-VOL, which makes the MSCS incorrectly recognize the volume under MSCS (MSCS tend to determine the volume by its signature).

Windows® 2000 can re-write the signature when the same signature is detected at the time when Windows® 2000 is booted. However, when MSCS is activated at the time when the system is being booted, the signature cannot be re-written.

Therefore, in order to avoid this phenomenon, execute the pairsplit command, and then reset the saved (sheltered) original signature if it was previously altered.
The following section describes the procedure for resetting the signature.

E.2.1 Changing the Signature

To change the signature:

1. Use the Disk Management on Windows to set the signature and the partition information for S-VOL.
2. Execute the inqraid –gvinf command to save (shelter) the signature and the volume layout information of S-VOL to the system device.
3. Execute the pairsplit command, then execute the inqraid –svinf command to re-set the saved (sheltered) signature and the volume layout information to the S-VOL. The signature is set as it was previously altered.

   The device object (\Device\HarddiskVolume#) and the Volume {guid} of S-VOL will not be generated when the pair is created using the noread option, and at the time when the Windows is booted. However, the device object (\Device\HarddiskVolume#) and the Volume {guid} will be generated by executing the inqraid –svinf command after the pairsplit command.