

HP StorageWorks RAID Manager XP User's Guide

XP48
XP128
XP512
XP1024
XP12000

eighth edition (March 2005)

part number: T1610-96003

This guide describes HP StorageWorks RAID Manager XP (RM) and provides installation and configuration procedures, RM command usage, and troubleshooting instructions.



© Copyright 2003-2005, Hewlett-Packard Development Company, L.P. All rights reserved.

Hewlett-Packard Company makes no warranty of any kind with regard to this material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

This document contains proprietary information, which is protected by copyright. No part of this document may be photocopied, reproduced, or translated into another language without the prior written consent of Hewlett-Packard. The information contained in this document is subject to change without notice.

Product names mentioned herein may be trademarks of their respective companies.

Hewlett-Packard Company shall not be liable for technical or editorial errors or omissions contained herein. The information is provided "as is" without warranty of any kind and is subject to change without notice. The warranties for Hewlett-Packard Company products are set forth in the express limited warranty statements accompanying such products. Nothing herein should be construed as constituting an additional warranty.

Printed in the U.S.A.

HP StorageWorks Disk Array XP RAID Manager: User's Guide

eighth edition (March 2005)
part number: T1610-96003

Contents

About this guide	9
Intended audience	9
Disk arrays	9
Related documentation	9
Conventions	10
Getting help	10
HP technical support	11
HP storage website	11
HP authorized reseller	12
Revision history	13
Warranty statement	14
HP StorageWorks LUN Security XP Extension disclaimer	16
1 RAID Manager description	17
RAID Manager features and environment	18
Continuous Access (CA)	19
Business Copy (BC)	20
Pairs and pair management	20
RAID Manager instances	22
RAID Manager command device	24
Manually switching command devices	25
2 Installation	27
System and disk array requirements	28
Using RAID Manager with Continuous Access	28
Using RAID Manager with Business Copy	28
Getting ready to install RAID Manager	29
Installing RAID Manager on UNIX systems	30

Installing RAID Manager on Windows systems	32
Installing RAID Manager on MPE/iX systems	33
Installing RAID Manager on OpenVMS systems	38
Configuring services and hosts files	40
Services file configuration	41
Hosts file configuration	41
Paired volume configuration	42
Setting Up the RM instance configuration file	43
Configuration file parameters	43
HORCM_MON section	46
HORCM_CMD section	49
HORCM_DEV Section	52
HORCM_INST section	55
Starting the instances	56
Environment variables for BC	56
Issuing CA commands	57

3 Using RAID Manager 59

Using RAID Manager commands	60
Using RAID Manager commands in scripts	61
Paired CA volume status definitions	62
Paired BC volume status definitions	64
File types and structure	66
Log files	70
Log file format	70
Log directories	72
Output to host log file	72
User-created files	73
User-settable environment variables	74
RM environment variables	74
RM command environment variables	75
RM instance environment variable	76

RM protection	77
Protection facility specification	78
Permission command	79
Protection facility support	79
Command device configuration	80
Commands controlled by RM protection	83
Permitting operations on protected volumes	84
Environment variables	88
Identifying a command device using protection mode	89
Using RAID Manager on a Windows 2000/2003 system with “user” system privileges	90
Windows System Administrator	90
RAID Manager user	95
Sample BAT file	97
LUN Security Extension	99
Guarding options	99
Commands affected	99
Notes and Restrictions	100
Identifying Open LDEV Guard volumes	101

4 RAID Manager command reference 103

General commands	104
Windows NT/2000/2003 commands	105
Data integrity check commands	106
horcctl	107
horcmshutdown	110
horcmstart	111
horctakeover	112
inqraid	116
mkconf	126
paircreate	129
paircurchk	136
pairedisplay	139
pairevwait	149
pairmon	155
pairresync	158
pairsplit	166
pairsyncwait	173
pairvolchk	180
raidar	189

- raidqry 192
- raidscan 195
- Command Options for Windows NT/2000/2003 207
- drivescan 208
- env 210
- findcmddev 211
- mount 212
- portscan 215
- setenv 217
- sleep 218
- sync 219
- umount 222
- usetenv 224
- Data Integrity Check Commands 225
- raidvchkset 226
- raidvchkdsp 233
- raidvchkscan 239

5 Troubleshooting RAID Manager 249

- Error reporting 250
- Operational notes 251
- Error codes 254
- Command return values 256
- Command errors 259

A Configuration file examples 265

- Configuration definition for cascading volumes 266
 - Correspondence between a configuration file and mirror descriptors 266
 - Cascading connection configuration files 267
- CA configuration (remote CA, two hosts) 271
 - CA (remote CA, two host) command examples 274
- CA configuration (local loopback, two hosts) 276
 - CA (local loopback, two hosts) command examples 278
- CA configuration (two RM instances, one host) 280
 - CA (two RM instances, one host) command examples 282
- BC configuration 284
 - BC command examples 287
- Configuration for a BC cascaded connection 293
 - BC cascaded connection command examples 295

	Configuration for a CA/BC cascaded connection	297
	CA/BC cascaded connection command examples	300
	Two-host BC configuration	303
	Two BC mirror configuration	305
	Three-host BC configuration	307
	Device group configuration	309
B	HA Failover and failback	311
	Using RAID Manager in HA environments	312
	HA control script state transitions	312
	Failback after SVOL-SMPL takeover	316
	PVOL-PSUE takeover	321
	S-VOL data consistency function	329
	Takeover-switch function	332
	Swap-takeover function	334
	SVOL-takeover function	336
	PVOL-takeover function	338
	Recovery procedures of HA system configuration	339
	Regression and recovery of CA	342
	CA recovery procedures	343
C	Fibre Channel addressing	345
	Fibre Channel address conversions	346
D	STDIN file formats	351
E	Porting notice for MPE/iX	353
	Porting notice for MPE/iX	354
	Introduction	354
	Restrictions in the current release	354
	Known issues and concerns	357
	MPE/iX startup procedures	359
F	Porting notice for OpenVMS	363
	Porting notice for OpenVMS	364
	Introduction	364
	Requirements and restrictions	364
	Known issues and concerns	371
	DCL command examples	377
	Startup procedures on bash	382

Glossary 387

Index 395

About this guide

This guide describes HP StorageWorks RAID Manager XP (RM) and provides installation and configuration procedures, RM command usage, and troubleshooting instructions. It also has configuration file examples and information about High Availability failover and failback, Fibre Channel addressing, and standard input (STDIN) file formats.

Intended audience

This guide is intended for use by system administrators who have knowledge of:

- Data processing concepts
- Direct access storage device subsystems (disk arrays) and their basic functions
- RAID technology
- Operating system commands and utilities

Disk arrays

Unless otherwise noted, the term *disk array* refers to these disk arrays:

HP Surestore Disk Array XP512
HP Surestore Disk Array XP48
HP StorageWorks Disk Array XP128
HP StorageWorks Disk Array XP1024
HP StorageWorks XP12000 Disk Array

Related documentation

HP provides the following related documentation:

- *HP StorageWorks Continuous Access XP: User's Guide*
- *HP StorageWorks Business Copy XP: User's Guide*
- *HP StorageWorks Command View XP for XP Disk Arrays: User Guide*

Conventions

This guide uses the following text conventions.

page 23	Blue text represents a cross-reference. For the online version of this guide, the reference is linked to the target.
www.hp.com	Underlined blue text represents a website on the Internet. For the online version of this guide, the reference is linked to the target.
literal	Bold text represents literal values that you type exactly as shown, as well as key and field names, menu items, buttons, file names, application names, and dialog box titles.
<i>variable</i>	Italic type indicates that you must supply a value. Italics is also used for manual titles.
input/output	Monospace font denotes user input and system responses, such as output and messages.
<i>Example</i>	The word “example” in italics denotes an example of input or output. The display shown in this guide may not match that of your configuration exactly.
[]	Square brackets indicate an optional parameter.
{ }	Braces indicate that you must specify at least one of the listed options.
	A vertical bar separates alternatives in a list of options.

Getting help

If you have questions after using this guide, contact your HP representative or access our website:

www.hp.com

HP technical support

In North America, call technical support at 1-800-652-6672, available 24 hours a day, 7 days a week.

Outside North America, call technical support at the nearest location. Telephone numbers for worldwide technical support are listed on the HP website under support:

<http://h18006.www1.hp.com/storage/array systems.html>

Be sure to have the following information available before calling:

- Technical support registration number (if applicable)
- Product serial numbers
- Product model names and numbers
- Applicable error messages
- Operating system type and revision level
- Detailed, specific questions

For continuous quality improvement, calls may be recorded or monitored.

HP storage website

Visit the support website for the most current information about HP StorageWorks XP products.

<http://h18006.www1.hp.com/storage/array systems.html>

Consult your HP account representative for information about product availability, configuration, and connectivity.

HP authorized reseller

For the name of your nearest HP authorized reseller, you can obtain information by telephone:

United States 1-800-345-1518

Canada 1-800-263-5868

Or visit: www.hp.com

Revision history

September 1999	OPEN-8 emulation added.
January 2000	Content extensively revised and reorganized.
September 2000	Content extensively revised.
February 2001	Added support of MPE/iX. Content significantly enhanced.
March 2001	Added mkconf command. Content enhanced.
November 2003	Added Oracle Data Validation. Added OpenVMS. Content significantly enhanced.
July 2004	General edit of content, layout, and language. General update to reflect recent changes. Modified installation procedure. Modified document to reflect HP12000 disk arrays. Added “Using RAID Manager with ‘user’ system privileges on a Windows 2000/2003 system.” Added HP StorageWorks LUN Security XP Extension disclaimer.
March 2005	Added -fe option for the <code>raidvchkdsp</code> , <code>raidscan</code> , and <code>pairedisplay</code> commands (display of information about external LUNs). Clarified use of <code>-l</code> option in <code>pairsplit</code> command (applies to Continuous Access only). Added CA-Journal description and operations. Added CA-Journal options to <code>raidscan</code> , <code>pairedisplay</code> , <code>inqraid</code> , <code>raidscan</code> , <code>raidvchkdsp</code> , <code>raidvchkscan</code> , <code>pairedisplay</code> , and <code>pairsyncwait</code> commands.

Warranty statement

HP warrants that for a period of ninety calendar days from the date of purchase, as evidenced by a copy of the invoice, the media on which the Software is furnished (if any) will be free of defects in materials and workmanship under normal use.

DISCLAIMER. EXCEPT FOR THE FOREGOING AND TO THE EXTENT ALLOWED BY LOCAL LAW, THIS SOFTWARE IS PROVIDED TO YOU “AS IS” WITHOUT WARRANTIES OF ANY KIND, WHETHER ORAL OR WRITTEN, EXPRESS OR IMPLIED. HP SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OR CONDITIONS OF MERCHANTABILITY, SATISFACTORY QUALITY, NON-INFRINGEMENT, TITLE, ACCURACY OF INFORMATIONAL CONTENT, AND FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow exclusions of implied warranties or conditions, so the above exclusion may not apply to you to the extent prohibited by such local laws. You may have other rights that vary from country to country, state to state, or province to province.

WARNING! YOU EXPRESSLY ACKNOWLEDGE AND AGREE THAT USE OF THE SOFTWARE IS AT YOUR SOLE RISK. HP DOES NOT WARRANT THAT THE FUNCTIONS CONTAINED IN THE SOFTWARE WILL MEET YOUR REQUIREMENTS, OR THAT THE OPERATION OF THE SOFTWARE WILL BE UNINTERRUPTED, VIRUS-FREE OR ERROR-FREE, OR THAT DEFECTS IN THE SOFTWARE WILL BE CORRECTED. THE ENTIRE RISK AS TO THE RESULTS AND PERFORMANCE OF THE SOFTWARE IS ASSUMED BY YOU. HP DOES NOT WARRANT OR MAKE ANY REPRESENTATIONS REGARDING THE USE OR THE RESULTS OF THE USE OF THE SOFTWARE OR RELATED DOCUMENTATION IN TERMS OF THEIR CORRECTNESS, ACCURACY, RELIABILITY, CURRENTNESS, OR OTHERWISE. NO ORAL OR WRITTEN INFORMATION OR ADVICE GIVEN BY HP OR HP’S AUTHORIZED REPRESENTATIVES SHALL CREATE A WARRANTY.

LIMITATION OF LIABILITY. EXCEPT TO THE EXTENT PROHIBITED BY LOCAL LAW, IN NO EVENT INCLUDING NEGLIGENCE WILL HP OR ITS SUBSIDIARIES, AFFILIATES, DIRECTORS, OFFICERS, EMPLOYEES, AGENTS OR SUPPLIERS BE LIABLE FOR DIRECT, INDIRECT, SPECIAL, INCIDENTAL, CONSEQUENTIAL, PUNITIVE OR OTHER DAMAGES (INCLUDING LOST PROFIT, LOST DATA, OR DOWNTIME COSTS), ARISING OUT OF THE USE, INABILITY TO USE, OR THE RESULTS OF USE OF THE SOFTWARE, WHETHER BASED IN WARRANTY, CONTRACT, TORT OR OTHER LEGAL THEORY, AND WHETHER OR NOT ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Your use of the Software is entirely at your own risk. Should the Software prove defective, you assume the entire cost of all service, repair or correction. Some jurisdictions do not allow the exclusion or limitation of liability for incidental or consequential damages, so the above limitation may not apply to you to the extent prohibited by such local laws.

NOTE. EXCEPT TO THE EXTENT ALLOWED BY LOCAL LAW, THESE WARRANTY TERMS DO NOT EXCLUDE, RESTRICT OR MODIFY, AND ARE IN ADDITION TO, THE MANDATORY STATUTORY RIGHTS APPLICABLE TO THE LICENSE OF THE SOFTWARE TO YOU; PROVIDED, HOWEVER, THAT THE CONVENTION ON CONTRACTS FOR THE INTERNATIONAL SALE OF GOODS IS SPECIFICALLY DISCLAIMED AND SHALL NOT GOVERN OR APPLY TO THE SOFTWARE PROVIDED IN CONNECTION WITH THIS WARRANTY STATEMENT.

HP StorageWorks LUN Security XP Extension disclaimer

HP StorageWorks LUN Security XP Extension provides the ability to place logical volumes into secure states. In these secure states, data on the volumes can not be modified until the retention time specified when the volume is placed in the secured state has elapsed. TO THE FULLEST EXTENT PERMITTED BY LAW, UNLESS EXPRESSLY PROVIDED OTHERWISE UNDER WRITTEN AGREEMENT BETWEEN HP AND CUSTOMER, HP WILL NOT BE LIABLE FOR ANY DIRECT, SPECIAL OR CONSEQUENTIAL DAMAGES (INCLUDING BUT NOT LIMITED TO LOSS OF PROFITS AND LOSS OF DATA) REGARDLESS OF WHETHER SUCH DAMAGES ARE BASED ON TORT, WARRANTY, CONTRACT OR ANY OTHER LEGAL THEORY, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES, ARISING OUT OF ANY PERFORMANCE OF OR IN FURTHERANCE OF UNLOCKING OR OTHERWISE UNSECURING AT THE REQUEST OF THE CUSTOMER DATA THAT HAS BEEN SECURED USING HP StorageWorks LUN Security XP Extension BEFORE THE RETENTION TIME HAS EXPIRED.

RAID Manager description

HP StorageWorks RAID Manager XP (RM) enables you to perform operations with HP StorageWorks Continuous Access XP (CA) and HP StorageWorks Business Copy XP (BC) by issuing commands from a host to the disk array. The RM software interfaces with the host system software and host high availability (HA) software, as well as with the BC and CA software on the disk array.

This book provides instructions for installing and using the RM software on various versions of UNIX, Microsoft Windows, and MPE/iX. It has an extensive command reference and additional materials about configuration, failover, and failback.

RAID Manager features and environment

RAID Manager lets you issue Business Copy (BC) and Continuous Access (CA) commands from a host. These commands can be issued from the command line or built into a script (for example, a ksh, perl script, or an MS-DOS batch file).

You can execute a large number of BC and CA commands in a short period of time by using scripts containing RM commands. In MPE/iX, you can create POSIX command scripts. For more information about scripting, see [“Using RAID Manager commands in scripts” \(page 61\)](#).

RM software consists of the following:

- RM instances (daemons)
- configuration files
- BC/CA commands and shell scripts

RM uses these entities:

- a special volume called a command device
- BC/CA volumes

RM runs in these (and other) environments:

UNIX	RM runs on a UNIX host as the HORCM daemon.
Windows NT/2000/2003	RM runs on a Windows NT/2000/2003 host as a service.
MPE/iX	RM runs in MPE/iX as a job stream. See Appendix E, “Porting notice for MPE/iX” .
OpenVMS	RM runs on OpenVMS as a detached process. See Appendix F, “Porting notice for OpenVMS” .

Continuous Access (CA)

CA copies data from a local HP XP disk array to one or more remote HP XP disk arrays. You can use CA for data duplication, migration, and offsite backup.

RM displays CA volume or group information and allows you to perform CA operations through either the command line, a script (UNIX), or a batch file (Windows).

CA has a number of features that ensure reliable transfers in asynchronous mode, including journaling and protection against link failure.

For effective and complete disaster recovery solutions, CA (and therefore RM) is integrated with many cluster solutions, such as Cluster Extension (CLX) for Windows, Linux, Solaris and AIX, as well as MetroCluster and ContinentalCluster for HP-UX.

CA modes

CA can operate in 3 different modes:

CA-Synchronous (CA-Sync): With CA-Sync all write operations on the primary (source) volume have to be replicated to the secondary (copy) volume before the write can be acknowledged to the host. This mode ensures the highest level of data concurrency possible. Host I/O performance is directly impacted by the distance between the primary and secondary volumes and therefore CA-Sync is recommended for metropolitan distances.

CA-Asynchronous (CA-Async): With CA-Async all write operations on the primary volume are time stamped and stored in a the array system cache, also known as the side file, before the write is acknowledged to the host. The data is then asynchronously replicated to the secondary array and re-applied in sequence to the secondary devices. With CA-Async data is not always current, but due to the unique timestamp implementation data will always be consistent. The side file functions to protect host I/O performance from any temporary degradations of the communication link between the sites. It also acts as a buffer for temporary high write bursts from the host. CA-Async is ideal for long distance replication.

CA-Journal: CA-Journal was introduced with the XP12000 and is only available for XP12000-to-XP12000 replication. CA-Journal works in principal the same as CA-Async, but instead of buffering write I/Os in the more expensive and limited XP array cache (the side file), CA-Journal writes data on special XP LUNS called journal pools. Journal pools can consists of up to 16 physical LDEVs of any size, and can therefore buffer much larger amounts of data. CA-Journal also implements a unique read operation from the remote array, instead of the normal write (push) operation from the local (primary) array, and is therefore much more tolerant of short communication link outages.

Business Copy (BC)

BC software allows you to create and maintain up to nine copies of data on the local disk array. You can use these copies for backup, data duplication, or testing.

BC duplicate volumes are created within the same disk array at hardware speeds.

RM displays BC volume or group information and allows you to perform BC operations through either the command line, a script (UNIX), or a batch file (Windows).

When you use CA to make a duplicate copy of a volume on a remote disk array, and then make up to 9 internal BC copies on the remote disk array from that volume, you can effectively create up to 10 copies of a logical volume on the remote disk array.

Pairs and pair management

Both BC and CA continuously copy data from a primary source volume known as a P-VOL to a secondary volume known as an S-VOL.

The relationship between a P-VOL and an S-VOL is called a pair.

You can use RM's **paircreate** command to establish pairs. Once a pair is established, updates to the P-VOL are automatically and continuously copied to the S-VOL.

There are other commands to manage pairs. You can temporarily suspend copy operations, resync the pair, and delete the pair relationship.

In addition, RM has many commands to display status, manage failover, manage failback, and set the conditions under which BC and CA execute.

RAID Manager instances

Each copy of RM is known as an RM instance. Instances are local or remote and can run on the same host or different hosts. Two RM instances are typically required to manage a BC or CA pair/group.

Local instance The RM instance currently being used, or the instance to which commands are issued. Local instances link to remote instances by using UDP socket services.

Remote instance The RM instance that the local instance communicates with, as configured in the **HORCM_INST** section of an RM instance configuration file. The **HORCM_INST** section defines how device groups link to remote RM instances.

RM also provides failover and operation commands that can support mutual hot standby when used with industry-standard failover software.

There are four possible RM topologies:

- **One host connected to one disk array**

If you are using one host, both RM instances are located on the same host.

- **One host connected to two or more disk arrays**

By connecting the host to two disk arrays you can implement disaster recovery by maintaining duplicate data on two different disk arrays. When you choose this option, the host Logical Volume Manager must not be able to see both sides of the same BC or CA pair or it will become confused.

- **Two or more hosts connected to one disk array**

If you are using two hosts, one RM instance is located on each host. Connecting two hosts to one disk array allows you to maintain copies of the data that are controlled by independent hosts, allowing primary volumes (P-VOLs) to be used by one host while secondary volumes (S-VOLs) are used by the other host (for example, for backup, testing, or data duplication).

- **Two or more hosts connected to two or more disk arrays**

By connecting two hosts to two or more disk arrays, you can implement disaster recovery by maintaining duplicate data at a remote location. Two hosts connected to two disk arrays allows the most flexible CA disaster recovery plan. The remote and local sets of data are administered by different hosts, guarding against host failure. This is the configuration used by HA software (such as HP MetroCluster) in conjunction with RAID Manager's **horctakeover** command, which allows for both failover and failback.

RAID Manager command device

You must designate a special volume on the disk array as the RAID Manager command device. The command device accepts BC or CA control operations. These are seen as in-band SCSI read and write commands, and are executed by the disk array. The volume designated as the command device is used only by RM and is blocked from other user access.

The command device can be any OPEN-x device that the host can access. An RM command device uses a minimum of 16 MB of space. The remaining volume space is reserved for RM and its utilities. You cannot use a Logical Unit Size Expansion (LUSE) volume as a command device; however, you can use the Volume Size Configuration (VSC) feature of Command View XP or LUN Configuration Manager XP to make custom volumes as small as 35 MB.

Caution *Be certain that there is no data on a volume you select as a command device. Data on the volume you select becomes inaccessible.*

Caution *MPE/iX systems will need a dummy volume set. Create this through the VOLUTIL utility program and scratch the volume set before converting to a command device.*

Caution *OpenVMS systems need a LUN 0 device of 35 MB. Note that storage assigned to the LUN 0 device is not accessible from OpenVMS.*

RM issues SCSI read/write commands to the command device. If the command device fails for any reason, all BC and CA commands terminate abnormally and the host cannot issue commands to the disk array.

To avoid data loss and system downtime, you can designate an alternate command device. Then, should RM receive an error notification in reply to a request, RM automatically switches to the alternate command device.

Manually switching command devices

To avoid having commands terminate abnormally during a failure, RM has a command device alternating function, which allows you to manually switch devices.

- **How to define alternate command devices**

You can define two or more command devices in the **HORCM_CMD** section of the configuration definition file. If you specify two or more devices on the same line, they are recognized as alternating control devices for the same array. See “[HORCM_CMD section](#)” (page 49).

- **When the command device switches**

When RM receives an error notification from the operating system, the RM switches to the alternate device.

You can also alternate command devices manually by issuing the RM **horcctl** command. See “[horcctl](#)” (page 107).

- **When to issue the horcctl (alternate command device) command**

Issue the **horcctl** command to alternate command devices before the command device is about to be blocked due to online maintenance.

After completing online maintenance, reissue the command. The first command device is then re-activated.

- **How RM handles multiple command devices at startup**

If only one (of several) command devices from the configuration definition file is available at RM startup, the startup log will contain a warning.

2

Installation

This chapter describes how to install and configure RAID Manager in several environments.

System and disk array requirements

RM requires that BC or CA be installed on the disk array. To determine the requirements for each product, see the following manuals:

HP StorageWorks Business Copy XP: User's Guide

HP StorageWorks Continuous Access XP: User's Guide

HP StorageWorks Command View XP for XP Disk Arrays: User Guide

Using RAID Manager with Continuous Access

Ask your HP representative to configure the disk arrays for CA functions. Make sure of the following:

- The disk arrays have CA license keys installed.
- The sender ports (Initiator for Fibre Channel, RCP for ESCON) and receiver ports (RCU-Target for Fibre Channel and LCP for ESCON) are configured on the local and remote disk arrays.
- The path between the master and remote control units (CUs) is established by using Continuous Access.
- Bidirectional swap is enabled between the primary and secondary volumes. Verify that at least two physical links exist in each direction.
- One or more RM command devices are set, using Command View XP or LUN Configuration Manager XP. If neither Command View XP nor LUN Configuration Manager XP is available, the HP representative can configure devices by using the disk arrays' internal service processors (SVPs).

Using RAID Manager with Business Copy

Ask your HP representative to configure the disk array for BC functions. Make sure of the following:

- The disk array has the BC license key installed.
- One or more RM command devices are set, using Command View XP or LUN Configuration Manager XP. If neither Command View XP nor LUN Configuration Manager XP is available, the HP representative can configure devices by using the disk array's SVPs.

Getting ready to install RAID Manager

You install and configure RM with assistance from the HP representative, as necessary.

Before you install and configure RM, you and your HP representative must complete these preliminary steps:

1. Plan the mapping of the BC and CA disk volume pairs. Determine which volumes to access.
2. Map the disk array interfaces to be used for each host.

After completing the preliminary steps above, installing and configuring RM consists of the following tasks. Task details appear in the subsequent sections.

- **Installing RAID Manager**

Install the RM software on the hosts.

- **Configuring the services and hosts files**

Add a service name/number to the host services file (for example, `/etc/services`) for each instance of RM. Configure the hosts file.

- **Setting up the RM instance configuration file**

Configure paths to one or more RM command devices for each host. All hosts and RM instances can use the same command device for a given disk array. However, it is recommended that each host have a its own command device.

- **Starting the instances**

This procedure includes setting environment variables needed to issue commands to the desired instance of RM.

Installing RAID Manager on UNIX systems

Follow the steps specific to your UNIX system to install RM.

Note: If RM is installed on a host and an active RM instance is running, shut it down before performing the installation.

1. Place the CD-ROM in the CD-ROM drive.
2. Identify the CD-ROM device file to be substituted in the **mount** commands below (for example, **/dev/dsk/c1t1d0**).
3. Log in as root.
su root
4. Create a CD-ROM mount directory and make it accessible to all users.
mkdir -p /cdrom
chmod 777 /cdrom
5. Mount the CD-ROM.

HP-UX

For HP-UX, use the **mount** command with the **-f** option:

```
mount -f cdfs -o ro /dev/dsk/c1t1d0 /cdrom
```

Sun Solaris

For Sun Solaris, use the **mount** command with the **-f** option:

```
mount -f hsfs -o ro /dev/dsk/c0t6d0s2 /cdrom/cdrom0
```

In most cases, Sun Solaris automatically mounts the CD-ROM. If not, use this **mount** command:

```
mount -f hsfs -o ro /vol/dev/dsk/c0t6d0/cdrom0 /cdrom/cdrom0
```

IBM AIX

For IBM AIX, use the **mount** command with the **-rv** option:

```
mount -rv cdrfs /dev/cd0 /cdrom
```

6. Choose a file system for the RM software. You need about 5 MB of disk space. The standard and recommended file system to load the software to is **/opt**.

7. From the **/opt** directory, use **cpio** to unpack the archives. Create the **HORCM** directory if it does not already exist.

cd /opt

mkdir HORCM (*choose the next command according to your OS*)

cat /cdrom/LINUX/rmxc* | cpio -idum (or)

cat /cdrom/AIX/rmxc* | cpio -idum (or)

cat /cdrom/DIGITAL/rmxc* | cpio -idum (or)

cat /cdrom/HP_UX/rmxc* | cpio -idum (or)

cat /cdrom/SOLARIS/rmxc* | cpio -idum

8. Change the directory to **/opt/HORCM** and verify the contents.

cd /opt/HORCM

ls

Example

```
etc horcmuninstall.sh log0 usr
horcminstall.sh log log1
```

9. Create a link from the root directory to the **/opt/HORCM** directory.

ln -s /opt/HORCM /HORCM

10. Run the RM Installer.

/HORCM/horcminstall.sh

This script creates symbolic links in the **/usr/bin** directory for RM commands.

When you complete the installation, create all required configuration files.

Installing RAID Manager on Windows systems

1. Boot the Windows server and log in with administrator access.
2. Insert the RAID Manager CD in the CD-ROM drive.
3. Under the **Start** menu, select **Run**.
4. When the Run window opens, enter **D:\WIN_NT\setup.exe** (where **D** is the letter of your CD-ROM drive) in the Open dialog box and click **OK**.
5. The installation wizard opens. Follow the on-screen instructions to install the RM software.

Installing RAID Manager on MPE/iX systems

Note: If RM is installed on a host and an active RM instance is running, shut it down before performing the installation.

1. Update your system with MPE/iX 6.5 or greater, along with that OS version's latest Power Patch.
2. Install the MPE/iX RAID Manager Patch ID **XPMMX65**.
3. Verify that at least one logical volume on the disk array is configured to function as a command device.

Caution *MPE/iX systems require that the command device be recognized as a dummy volume set. Create this through the VOLUTIL utility program and then scratch the volume before converting it to a command device.*

4. Run the POSIX shell from CI and change your working directory to the temporary directory **/tmp/raidmgr**.

```
: Sh
```

```
Shell/iX> cd /tmp/raidmgr
```

5. Execute the install script

```
Shell/iX> ./RMinstsh
```

This install script requests that you specify a POSIX directory where the RAID Manager executables and log files will be placed. The standard and recommended POSIX directory is **/opt**.

This script creates the necessary POSIX directories. All relevant files are placed under the directory **/opt/HORCM**. The RAID Manager executables are placed under **/opt/HORCM/usr/bin**. A symbolic link (**/HORCM**) that points to **/opt/HORCM** is created under the root directory.

6. Once the above installation completes successfully, create the device files:

```
Shell/iX> mknod /dev/ldev99 c 31 99 ← LDEV devices  
Shell/iX> mknod /dev/ldev100 c 31 100  
Shell/iX> mknod /dev/cmddev c 31 102 ← Command device
```

The 31 in the above example is called the major number. The 99, 100, 102 are called minor numbers. For RAID Manager, always specify **31** as the major number. The minor number should correspond to the LDEV numbers as configured in sysgen. Create device files for all the LDEVs configured through sysgen and for the command device. The device link file for the command device should be called **/dev/cmddev**.

7. Add a service entry for each RM instance in the **SERVICES.NET.SYS** file.

```
horcm0    6100g        #RaidManager instance 0  
horcm1    6100g        #RaidManager instance 1
```

8. Each host running an instance should be listed in the **HOSTS.NET.SYS** file. This allows you to refer to any remote hosts by name as well as by IP address. If a domain name server (DNS) manages host name resolution, then **HOSTS.NET.SYS** is not necessary.
9. Create RM instance configuration files for each instance. This file defines the relationships between RM instances and the link between a volume and an RM instance. When you create an RM configuration file, use this naming convention:

horcm#.conf

Example

```
horcm1.conf
```

where the instance # is the instance number (in this case 1).

The configuration file has four sections:

```
HORCM_MON (page 46)  
HORCM_CMD (page 49)  
HORCM_DEV (page 52)  
HORCM_INST (page 55)
```

An example **horcm.conf** file can be found in the **/HORCM/etc** directory. You can make copies of this file to create a different configuration file for each RM instance. Place the conf files in the **/etc** directory.

You will have to start RAID Manager without a description for **HORCM_DEV** and **HORCM_INST** because the target ID and LUN are not yet known. After RAID Manager is up and running, you can find the target ID and LUN by using the **raidscan -find** command.

10. Create a RAID Manager job for each RM instance in the group **PUB.SYS**. A sample job called **jhorcm** is in the **/etc** directory.

Example:

```
!job jraidmr1, manager.sys;pri=cs
!setvar TZ "PST8PDT"
!xeq sh.hpbin.sys '/HORCM/usr/bin/horcmstart.sh 1'
!eoj
```

The above example will start RAID Manager instance 1. In the example, the time zone variable (TZ) has to be set differently for each time zone.

11. Start the RAID Manager daemon by streaming the job.

```
: stream jraidmr1.pub.sys
```

Verify that the job is running by executing the **SHOWJOB** command:

```
: SHOWJOB
JOBNUM  STATE  IPRI  JIN  JLIST      INTRODUCED  JOB NAME
#S2     EXEC           20  20      THU  5:29P  MANAGER.SYS
#J15    EXEC          10S  LP      FRI  5:08P  JRAIDMR1,MANAGER.SYS
#J16    EXEC          10S  LP      FRI  5:08P  JRAIDMR2,MANAGER.SYS
```

- Get the physical mapping of the available LDEVs to fill in the HORCM_DEV and HORCM_INST sections of the **horcm1.conf** file. Invoke the shell and change your working directory to **/HORCM/usr/bin**. Execute:

```
:sh
Shell/iX> cd /HORCM/usr/bin
Shell/iX> export HORCMINST=1
Shell/iX> ls /dev/* | ./raidscan -find
```

DEVICE_FILE	UID	S/F	PORT	TARG	LUN	SERIAL	LDEV	PROD_ID
/dev/cmddev	0	S	CL1-D	1	0	35393	22	OPEN-3-CM
/dev/ldev407	0	S	CL1-E	8	0	35393	263	OPEN-3
/dev/ldev408	0	S	CL1-E	9	0	35393	264	OPEN-3
/dev/ldev409	0	S	CL1-E	10	0	35393	265	OPEN-3
/dev/ldev410	0	S	CL1-E	11	0	35393	266	OPEN-3
/dev/ldev411	0	S	CL1-E	12	0	35393	267	OPEN-3
/dev/ldev412	0	S	CL1-E	13	0	35393	268	OPEN-3

- Now fill in the HORCM_DEV and HORCM_INST sections in your **/etc/horcm#.conf** files.

*Sample Configuration
for Instance 1:*

```
#
#/******For HORCM_MON*****/
HORCM_MON
#ip_address          service          poll(10ms)          timeout(10ms)
NONE                 horcm0             1000                 3000
#/****** For HORCM_CMD*****/
HORCM_CMD
#dev_name            dev_name            dev_name
/dev/cmddev0
#/****** For HORCM_DEV*****/
HORCM_DEV
#dev_group           dev_name            port#                TargetID             LU#                  MU#
VG01                 oradb1              CL1-E                8                    0                    0
VG02                 oradb2              CL1-E                9                    0                    0
#/****** For HORCM_INST *****/
HORCM_INST
#dev_group           ip_address          service
VG01                 HSTB                horcm1
VG02                 HSTC                horcm1
```

14. Shut down the RAID Manager daemon within the shell and the current working directory **/HORCM/usr/bin**.

```
Shell/iX> ./horcmshutdown.sh 1
```

Restart the RAID Manager job using the completed RM configuration file:

```
: stream jraidmrl.pub.sys
```

For additional information, see Appendix E, [Porting notice for MPE/iX \(page 353\)](#).

Installing RAID Manager on OpenVMS systems

Privileges for using RAID Manager

- A user account for RAID Manager must have the same privileges as “SYSTEM” (that is, it must be able to use the SCSI class driver and Mailbox driver directly). Some OpenVMS system administrators may not allow RAID Manager to run from the system account. In this case, create another account on the system, such as “RMadmin” that has the same privileges as “SYSTEM.”
- RAID Manager uses the Mailbox driver for communication between RAID Manager components. So, the RAID Manager command processor and RM daemon (called HORCM) must have the same privileges.

If the RAID Manager command processor and HORCM execute with different privileges, then the RAID Manager command processor will hang or be unable to attach to the daemon.

- RAID Manager also requires that the logical name `sys$posix_root` exist on the system. Therefore, you must define `sys$posix_root` before installing RAID Manager.

It is recommended that the you define the following in LOGIN.COM before RM installation:

```
$ DEFINE/TRANSLATION= (CONCEALED, TERMINAL)
SYS$POSIX_ROOT " Device:[directory] "
$ DEFINE DCL$PATH SYS$POSIX_ROOT:[horcm.usr.bin],
SYS$POSIX_ROOT:[horcm.etc]
$ DEFINE/TABLE=LNM$PROCESS_DIRECTORY
LNM$TEMPORARY_MAILBOX LNM$GROUP
$ DEFINE DECC$ARGV_PARSE_STYLE ENABLE
$ SET PROCESS/PARSE_STYLE=EXTENDED
```

In the above defines, `Device:[directory]` is the `SYSS$POSIX_ROOT`.

To install RAID Manager on an OpenVMS system:

Install RAID Manager by using the file

HP-AXPVMS-RMXP-V0115-4-1.PCSI

1. Insert and mount the installation media.
2. Execute the following command.

```
$ PRODUCT INSTALL RMXP  
/source=Device:[PROGRAM.RM.OVMS]/LOG -  
_$_ /destination=SYS$POSIX_ROOT:[000000]
```

where Device:[PROGRAM.RMOVMS] is where file
HP-AXPVMS-RMXP-V0115-4-1.PCSI exists.

3. Confirm the installation:

```
$ raidqry -h  
Model   : Raid-Manager-XP/OpenVMS  
Ver&Rev: 01.15.02  
:  
:
```

For additional information, see Appendix F, [Porting notice for OpenVMS \(page 363\)](#).

Configuring services and hosts files

Part of configuring RM consists of editing the services and hosts files on the hosts that run RM instances.

UNIX

The services and hosts files are contained in this directory:

`/etc`

Windows NT/2000/2003

The services and hosts files are contained in this directory:

`%systemroot%\system32\drivers\etc`

OpenVMS

The services file is contained in this directory:

`SYSSYSROOT:[000000.TCPIP$ETC]SERVICES.DAT`

The hosts file is contained in this directory:

`SYSSYSROOT:[SYSEXE]HOST.DAT`

MPE/iX

The services and hosts files are contained in the MPE group directory:

`NET.SYS`

Services file configuration

Each BC and CA pair has a primary volume (P-VOL), the volume that contains the data to be copied, and a secondary volume (S-VOL), the volume that receives the data from the primary volume. Each of these volumes is linked to at least one instance of RM for the purpose of pair creation, suspension, and deletion. Each instance of RM can manage multiple volumes (on up to four arrays) and manage either P-VOLs or S-VOLs. See “[Setting Up the RM instance configuration file](#)” (page 43).

Example Instance 0 of RM is linked to the P-VOL for pairs a, b, and c. Instance 1 is linked to the S-VOL for pairs a, b, and c. To create, suspend, or delete pair a, both RM instances would communicate and cooperate.

Reminder: Instances can be on the same or different host systems. The host that is running the instance must have access to the volumes to which it is linked and have access to a disk array command device for the array.

To configure the services file:

1. Edit the **services** file on each system.
2. Add a **udp** service entry for each RM instance that runs on the host and each RM instance referenced in the configuration file. The service number selected must be unique to the **services** file and in the range 1024 to 65535.

Example

horcm0	11000/udp	#RaidManager instance 0
horcm1	11001/udp	#RaidManager instance 1

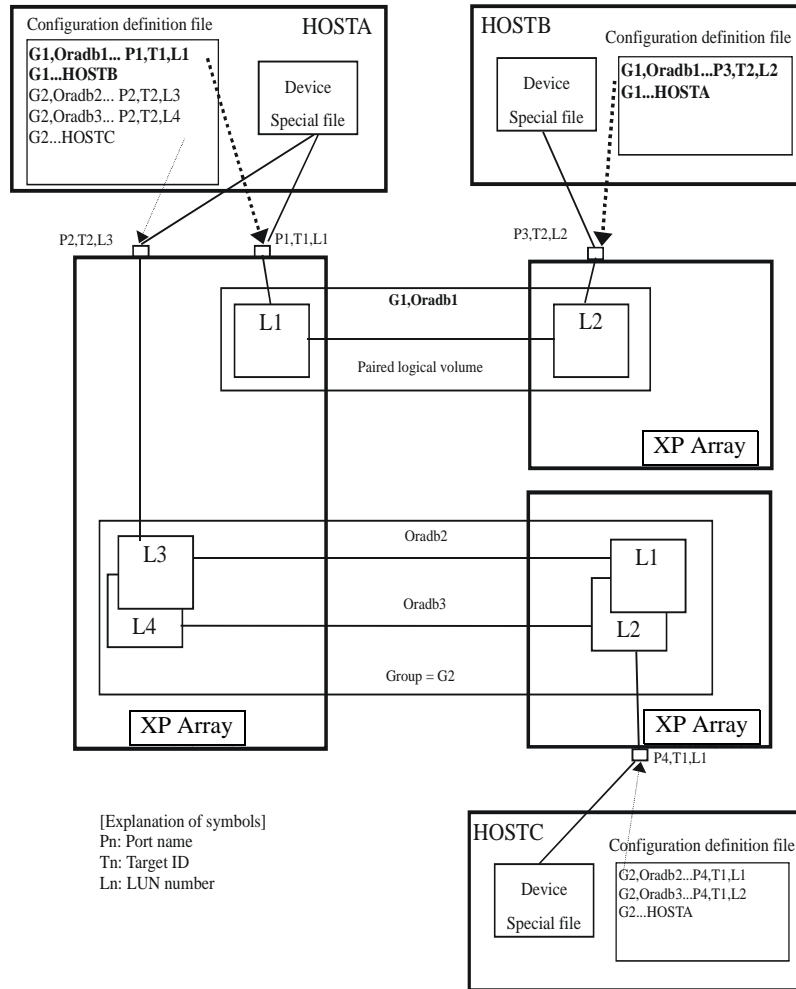
Hosts file configuration

Each host running an RM instance should be configured in the **hosts** file (for example, **/etc/hosts**). This lets you refer to any remote host by either its name or IP address.

If DNS (domain name server) manages host name resolution on your network, **hosts** file editing is not necessary.

Paired volume configuration

Users describe the connection between physical volumes used by the servers and the paired logical (named) volumes (and the names of the remote servers connected to the volumes) in a configuration definition file. See the figure below.



Setting Up the RM instance configuration file

The RM instance configuration file defines the link between a volume and an RM instance. This file also defines the relationships between RM instances and between the physical and logical names for volumes. You must create a configuration file for each RM instance.

When you create an RM configuration file, follow this naming convention, where *instance* is the instance number:

horcm*instance*.conf

Example horcm0.conf

As previously stated, the configuration file has four sections:

HORCM_MON (page 46)

HORCM_CMD (page 49)

HORCM_DEV (page 52)

HORCM_INST (page 55)

You can use the **mkconf** command to create a configuration file. See [mkconf \(page 126\)](#) for usage information.

If the level of detail provided in the following pages is not sufficient, ask your HP representative to consult the HP internal document:

RAID Manager XP Basic Specifications

Configuration file parameters

The configuration file contains all parameters and values for a RM instance. The following table lists the constraints on various configuration file parameters.

Parameter	Default Value	Type	Limit
<i>IP_address</i>	<i>None</i>	Character string	63 characters
<i>host_name</i>	<i>None</i>	Character string	31 characters
<i>service_name</i> or <i>service_number</i>	<i>None</i>	Character string or numeric value	15 characters
<i>poll_value</i> (10 ms increments)	1000	Numeric value	None
<i>timeout_value</i> (10 ms increments)	3000	Numeric value	None
<i>device_name</i> for <i>HORCM_DEV</i>	<i>None</i>	Character string	31 characters
<i>dev_group</i>	<i>None</i>	Character string	31 characters
<i>port</i>	<i>None</i>	Character string	31 characters
<i>target_ID</i>	<i>None</i>	Numeric value	7 characters
<i>LUN</i>	<i>None</i>	Numeric value	7 characters
<i>mirror_unit</i>	0	Numeric value	7 characters
<i>RM_group</i>	<i>None</i>	Character string	31 characters
<i>dev_name</i> for <i>HRCM_CMD</i>	<i>None</i>	Character string	63 characters

For examples of configuration files, see Appendix A, “[Configuration file examples](#)” (page 265).

HP-UX

An example **horcm.conf** file can be found in the **/HORCM/etc** directory.

Windows NT/2000/2003

An example **horcm.conf** file can be found in the **C:\HORCM\etc** directory.

MPE/iX

See Appendix E, [Porting notice for MPE/iX \(page 353\)](#).

Open VMS

See Appendix F, [Porting notice for OpenVMS \(page 363\)](#).

HORCM_MON section

Syntax **HORCM_MON**
{ *host_name* | *IP_address* } { *service_name* | *service_number* } *poll_value*
timeout_value }

host_name Name of the host on which this RM instance runs.

IP_address IP address of the host on which this RM instance runs. Specify *NONE* when two or more network cards are installed in the server, or several networks (subnets) are configured, and you want to use this RM feature to listen on all networks.

service_name Service name that was configured in the host services file.

service_number Service number that was configured in the host services file.

poll_value Specifies a monitoring interval for paired volumes. By making this interval longer, the RM daemon load is reduced, but it may take longer to notice a change in pair status.

If this interval is set to **-1**, paired volumes are not monitored. A value of **-1** should be specified when two or more instances of RM run on the same machine and one is already monitoring the pair.

timeout_value Specifies the remote server communication timeout period.

Description The **HORCM_MON** section describes the host name or IP address and the port number (and the paired volume error monitoring interval) of the local host.

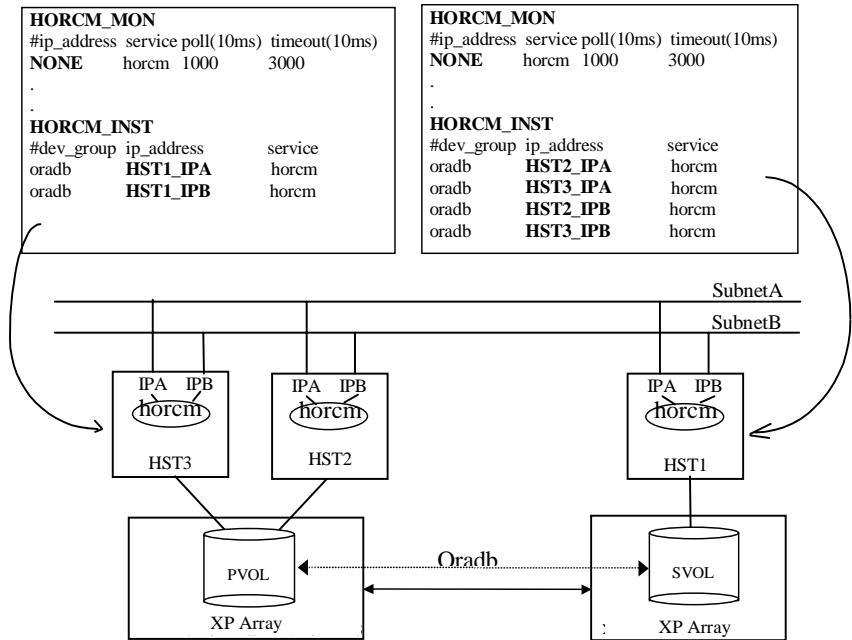
Examples **HORCM_MON**
blue horcm1 1000 3000

The RM instance is running on system blue with service name horcm1, a poll value of 10 seconds, and a timeout value of 30 seconds.

HORCM_MON
NONE horcm1 1000 3000

The RM instance is running on system NONE, indicating two or more network cards are installed in the server, or several networks (subnets) are configured, and the RM listens on all networks. The service name is horcm1, a poll value of 10 seconds, and a timeout value of 30 seconds.

The **raidqry -r group** command executed on each host allows the user to examine the multiple network configurations. The following figure shows that the volume group known as **oradb** is controlled from the right side of the diagram by host HST1 (using either subnet A or B) and from the left side of the diagram by either HST2 or HST3 (using either subnet A or B).



HORCM_CMD section

Syntax **HORCM_CMD**
command_device [command_device] . . .

Description The **HORCM_CMD** section defines the disk devices RM uses to communicate with the disk array. When a RM command is initiated, command data is written to the special disk array command device. The disk array reads this data and carries out the appropriate actions.

Multiple command devices can be defined in this section of the configuration file to provide alternate command devices and paths in the event of failure.

It is recommended that each host has a unique command device. A command device should not be accessed by more than one host. Multiple instances on the same host can use the same command device.

To configure command devices, use LUN Configuration Manager. If LUN Configuration Manager is not available, an HP representative can configure the command devices by using the internal disk array service processor (SVP).

Examples

HP-UX **HORCM_CMD**
/dev/rdisk/c2t3d0 /dev/rdisk/c6t2d4

This example defines two device files as paths to a command device. These devices can be pvlincs to the same volume on the disk array, or may be different command devices. Placing the second command device on the same line implies that it is an alternate within the same array.

```
HORCM_CMD  
#unitID0 (Array 1)  
/dev/rdisk/c1t3d5  
#unitID1 (Array 2)  
/dev/rdisk/c2t3d5
```

This HP-UX example shows multiple disk arrays connected to the host. One RM instance can control multiple disk arrays. To enable this feature, the different command devices have to be specified on different lines. RM uses unit IDs to control multiple disk arrays. A device group can span multiple disk arrays (sync-CA only). The unit ID must be appended for every volume device name in the **HORCM_DEV** section, as shown in the following figure.

```

HORCM_MON
#ip_address service poll (10ms) timeout (10ms)
HST1          horcm    1000         3000

HORCM_CMD
#unitID 0... (seq#30014)
#dev_name      dev_name dev_name
/dev/rdisk/c0t0d0 /dev/rdisk/c0t0d1
#unitID 1... (seq#30015)
#dev_name      dev_name dev_name
/dev/rdisk/c1t0d0

HORCM_DEV
#dev_group  dev_name port#   TargetID  LU#  MU#
oradb       oradb1  CL1-A    3      10
oradb       oradb2  CL1-A    3      11
oralog      oralog1 CL1-A    5      0
oralog      oralog2 CL1-A1   5      0
oralog      oralog3 CL1-A1   5      1

HORCM_INST
#dev_group  ip_address service
oradb       HST2      horcm
oradb       HST3      horcm
oralog      HST3      horcm

```

*Windows NT/2000/
2003* **HORCM_CMD**
 \\.\PHYSICALDRIVE3

This example shows the path to a shared command device in Windows.

\\.\Volume{GUID}

This example shows the use of a Volume GUID for the command device in Windows.

MPE/iX See Appendix E, [Porting notice for MPE/iX \(page 353\)](#).

OpenVMS See Appendix F, [Porting notice for OpenVMS \(page 363\)](#).

HORCM_DEV Section

Syntax HORCM_DEV

device_group device_name port target_ID LUN [mirror_unit]

device_group Each device group contains one or more volumes. This parameter gives you the capability to act on a group of volumes with one RM command. The device group can be any user-defined name up to 31 characters in length.

device_name User-defined and unique to the instances using the device groups. It can be up to 31 characters in length and is a logical name that can be used instead of the physical Port/TID/LUN/MU# designation.

port Disk array I/O port through which the volume is configured to be accessed. Port specification is not case sensitive (CL1-A= cl1-a= CL1-a= cl1-A).

target_ID SCSI/Fibre target ID assigned to the volume.

LUN Decimal logical unit number assigned to the volume.

mirror_unit Used when you are making multiple BC copies from a P-VOL. The mirror unit is a number ranging from 0 to 2 and has to be explicitly supplied for all BC volumes.

If *mirror_unit* is left blank it will be assumed that CA-Sync or CA-Async is being used. The number is not a count of the number of copies to be made but rather a label for a specific P-VOL to S-VOL relationship.

CA-Journal will allow up to four copies from a P-VOL. The mirror unit for a CA-Journal volume is indicated by an “h” and a number ranging from 0 to 3. If *mirror_unit* is omitted, the value of h0 will be assumed. Mirror unit value “h1”, “h2” and “h3” are valid only for CA-Journal operations.

```

Example  HORCM_MON
           #ip_address  service  poll(10ms  timeout(10ms)
           HST1         horcm     1000     3000

           HORCM_CMD
           #dev_name    dev_name  dev_name
           /dev/rsd0e

           HORCM_DEV
           #dev_group  dev_name  port#    TargetID  LU#    MU#
           Group      dev       CL1-A    3         1
           Group1     dev1     CL1-A    3         1     0
           Group2     dev2     CL1-A    3         1     1
           Group3     dev3     CL1-A    3         1     2
           Group4     dev4     CL1-A    3         1     h1

           HORCM_INST
           #dev_group  ip_address  service
           Group      HST2       horcm
           Group1     HST3       horcm

```

The validity of the mirror descriptor is shown for various pair states in the table below.

Feature	SMPL		P-VOL		S-VOL	
	MU#0	MU#1-3	MU#0	MU#1-3	MU#0	MU#1-3
CA	Valid	Invalid	Valid	Invalid	Valid	Invalid
CA-Journal	Valid	Valid	Valid	Valid	Valid	Valid

Description The **HORCM_DEV** section describes the physical volumes corresponding to the paired volume names. Each volume listed in **HORCM_DEV** is defined on a separate line.

Example **HORCM_DEV**
group1 g1-d1 CL1-A 12 1 0

This example shows a volume defined in device group1 known as device g1-d1. It is accessible through disk array unit 0 and I/O port CL1-A. The SCSI target ID is 12, the LUN is 1, and the BC mirror unit number is 0.

You can use RM to control multiple disk arrays with one RM instance by specifying the unit ID appended to the port. This example refers to the example in the [HORCM_CMD section \(page 49\)](#).

HORCM_DEV

```
group1  g1_d1  CL1-A  12  0  
group2  g2-d1  CL1-A1 12  0
```

This example shows that the volume pair with the device name g2-d1 resides on disk array unit 1 while the volume pair with device name g1-d1 resides on disk array unit 0.

Tip In the case of Fibre Channel, if the host reports a different target ID and LU# than **raidscan**, use the **raidscan** value.

Related information

To see configuration file examples, and to see how devices belonging to different unit IDs are configured, see Appendix A, [Configuration file examples \(page 265\)](#).

HORCM_INST section

Syntax HORCM_INST

```
device_group { host_name | IP_address } { service_name |  
service_number }
```

device_group Defined in the HORCM_DEV section. Each group defined in HORCM_DEV must be represented in the HORCM_INST section only once for every remote RM instance.

host_name Host name of the host on which the remote instance runs. The remote instance can run on the same host as the local instance.

IP_address IP address of the host on which the remote instance runs. The remote instance can run on the same host as the local instance.

service_name Service name that was entered into the services file for the remote instance.

service_number Service number that was entered into the services file for the remote instance.

Description The HORCM_INST section defines how RM groups link to remote RM instances.

Example The example below shows that the opposite side of the pairs contained within the group called group1 are serviced from a RM residing on host yellow that listens on a UDP port defined in **/etc/services** to have the name horcm0.

```
HORCM_INST  
group1 yellow horcm0
```

Starting the instances

After the RM instance configuration files have been set up, you can start the instances. See the instructions for your operating system:

HP-UX

Run this shell command on each host that runs an RM instance:

```
/usr/bin/horcstart.sh [ instance_number ] [ instance_number ] . . .
```

If you do not specify an instance number, the command uses the value stored in the **HORCMINST** environment variable. The default value is 0.

Windows NT/2000/2003

From the DOS prompt, under the **\HORCM\etc** directory, type this command:

```
horcstart instance_number [ instance_number ] . . .
```

MPE/iX

See Appendix E, [Porting notice for MPE/iX \(page 353\)](#).

OpenVMS

Run RM instance as a detached process. See Appendix F, [Porting notice for OpenVMS \(page 363\)](#).

Environment variables for BC

By default, all RM operations affect CA volumes. To enable RM commands to control BC operations, set the **HORCC_MRCF** environment variable to **1**.

RM commands are issued to the local instance host. To specify which instance is the local instance, set the **HORCMINST** environment variable,

as in the following examples of environment variables, where *n* is the value of the RM instance for the RM command.

UNIX

For UNIX ksh, use the **export** command:

```
export HORCC_MRCF=1  
export HORCMINST=n
```

For UNIX csh, use the **setenv** command:

```
setenv HORCC_MRCF=1  
setenv HORCMINST=n
```

Windows NT/2000/2003

For Windows NT/2000/2003, use the **set** command:

```
set HORCC_MRCF=1  
set HORCMINST=n
```

MPE/iX

For MPE/iX, use the **setenv** command.

```
setenv HORCC_MRCF 1  
setenv HORCMINST n
```

OpenVMS

For OpenVMS, set the environment variable using symbol.

```
HORCC_MRCF := 1  
HORCMINST := 0
```

Issuing CA commands

To issue CA commands, the **HORCC_MRCF** environment variable must be removed and the **HORCMINST** environment variable must be set.

UNIX

Setting a null value is not sufficient.

For UNIX ksh, use the **unset** command:

```
unset HORCC_MRCF  
set HORCMINST=n
```

For UNIX csh, use the **unsetenv** command:

```
unsetenv HORCC_MRCF  
setenv HORCMINST=n
```

Windows NT/2000/2003

For Windows NT/2000/2003, use the **usetenv** command option:

```
raidscan -x usetenv HORCC_MRCF  
raidscan -x setenv HORCMINST n
```

Related Information For syntax descriptions, see **usetenv** (page 224) and **setenv** (page 217).

MPE/iX

Within the POSIX shell, use the **unset** command:

```
unset HORCC_MRCF  
set HORCMINST=n
```

OpenVMS

For Open VMS, use the following command:

```
$DELETE/SYMBOL HORCC_MRCF
```

3

Using RAID Manager

This chapter describes commands, scripts, functions, and procedures for RAID Manager (RM).

Using RAID Manager commands

To create and manage CA and BC pairs with RM, use the following commands:

paircreate	Establishes a primary to secondary pair relationship between volumes. See “ paircreate ” (page 129).
pairdisplay	Displays the state of volumes. See “ pairdisplay ” (page 139).
pairsplit	Suspends or deletes a paired volume. See “ pairsplit ” (page 166).
pairresync	Restores a volume from a PSUE/PSUS/SSWS state to a PAIR state. See “ pairresync ” (page 158).

Before issuing RM commands, set the **HORCMINST** environment variable to refer to the instance number you want to use for the local RM instance.

Continuous Access

After installing CA, the system administrator must conduct operation tests for recovery and maintenance, to check for possible failures that can occur. The HP representatives can identify the possible causes of a possible failure by using the SVP.

Caution *For CA, make sure to unset the **HORCC_MRCF** variable. Do not set the **HORCC_MRCF** variable to 0.*

Business Copy

For BC, set **HORCC_MRCF=1**.

Using RAID Manager commands in scripts

An RM script is a list of instructions contained in a host file to automate a series of CA and BC operations. The host reads the script file and carries out each command as if it were typed in individually.

Using RM host scripting, you can execute a large number of CA and BC commands in rapid sequence.

Paired CA volume status definitions

Each pair of CA volumes consists of a primary volume (P-VOL) and secondary volume (S-VOL). Each pair has a pairing status.

The P-VOL controls the status for the pair, which is reflected in the status of the S-VOL. The major CA pair statuses are:

- SMPL
- PAIR
- PSUS
- COPY
- PSUE
- PFUS

When you issue a CA command, the status usually changes. A read or write request from the host is allowed or rejected, depending on the status of the paired volume, as shown in the figure on the next page.

Caution

Terminology and functionality differ somewhat between the RM CLI interface and the RC GUI interface. For instance:

- *The terms “suspend” and “split” may have opposite meanings*
- *S-VOL read/write options while suspended may differ*
- *The GUI allows you to choose/force a PSUE state*


For more detail, refer to the following manuals:

HP StorageWorks Business Copy XP: User’s Guide

HP StorageWorks Continuous Access XP: User’s Guide

Note that the data at the async-CA S-VOL is assured to be consistent, but is only current in PSUS state.

If one of the volumes making up an aggregated LUSE volume is PSUE status, the LUSE volume will be reported as PDUB (dubious) status.



Status	Pairing status	Primary	Secondary
SMPL	Unpaired volume	R/W enabled	R/W enabled
PAIR	Paired/duplicated volume, consistency assured	R/W enabled	R* enabled
COPY	In paired state, but (out of order) copying to the secondary volume is not yet completed	R/W enabled	R* enabled
PSUS	In paired state, but updating the secondary volume is suspended. The P-VOLs and S-VOLs track the differences in the changed data. The P-VOL data is typically dominant.	R/W enabled	See Note 1. R/W enabled
PSUE (Error)	PSUS status due to an internal or link error	See Note 2. R/W enabled	R* enabled
PFUS	PSUS status due to a side file full error	See Note 2. R/W enabled	R* enabled

Note 1: Valid when reading and writing are enabled using an option of the pair splitting.
 Note 2: Reading and writing are enabled as long as no errors occur in the primary volume.
 R*: Reading disabled when **-m noread** is specified in the **paircreate** command.

Paired BC volume status definitions

Each paired BC volume consists of a primary volume (P-VOL) and secondary volume (S-VOL). Each volume maintains its own pair status.

The P-VOL controls the pair state that is typically reflected in the status of the S-VOL. The major BC pair statuses are:

- SMPL
- PAIR
- PSUS
- COPY
- RCPY
- PSUG

The status can be changed when a RM command is issued. A read or write request from the host is allowed or rejected according to the status, as shown in the following figure.

Caution

Terminology and functionality differ somewhat between the RM CLI interface and the RC GUI interface. For instance:

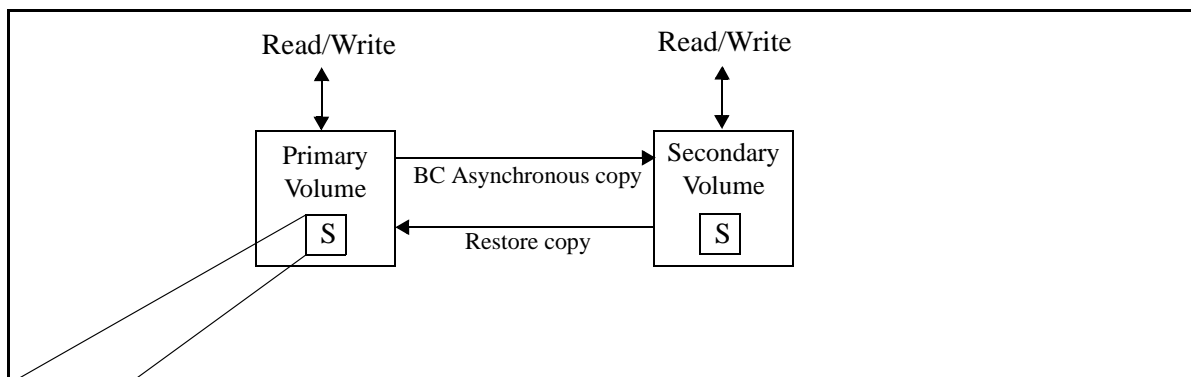
- *The terms “suspend” and “split” may have opposite meanings*
- *S-VOL read/write options while suspended may differ*
- *The GUI allows you to choose/force a PSUE state*

For more detail, refer to the following manuals:

HP StorageWorks Business Copy XP: User’s Guide

HP StorageWorks Continuous Access XP: User’s Guide

Note that reading the BC S-VOL in any state but PSUS is “at your own risk” in that the data is likely to be inconsistent as well as not current.



Status	Pairing status	Primary	Secondary
SMPL	Unpaired volume	R/W enabled	R/W enabled
PAIR	Paired/duplicated volumes. Data in the primary and secondary volumes are not assured to be identical.	R/W enabled	R* enabled
COPY	In paired state, but copying to the secondary volume is not completed. The P-VOL/S-VOL are not assured to be identical.	R/W enabled	R* enabled
RCPY	This reverse copy state infers copying from the secondary to the primary volume by using the -restore option of pairresync .	R* enabled	R enabled
PSUS	In paired state, but updating the secondary volume is suspended. The P-VOL and S-VOL notes any new changes in a bitmap.	R/W enabled	R/W enabled
PSUE (Error)	PSUS status due to an internal failure (or RC GUI command). The P-VOL and S-VOL maintain a delta bitmap.	R/W enabled (See Note 1)	R* enabled

Note 1 : Valid when reading and writing are enabled, as far as no failure occurs in the P-VOL.

R* : Reading disabled when the user specified the **-m noread** option in the paircreate command.

File types and structure

The RM product includes files supplied for the user, log files created internally, and files created by the user. These files are stored in the server's local disk. See the following tables.

Title	File name, Location	Executable for Command
HORCM (RM)	/etc/horcmgr	<i>none</i>
HORCM_CONF	/HORCM/etc/horcm.conf	<i>none</i>
Takeover	/usr/bin/horctakeover	horctakeover
Make configuration file	/usr/bin/mkconf.sh	mkconf
Volume Accessibility check	/usr/bin/paircurchk	paircurchk
Pair generation	/usr/bin/paircreate	paircreate
Pair splitting/suspending	/usr/bin/pairsplit	pairsplit
Pair resynchronization	/usr/bin/pairresync	pairresync
Event waiting	/usr/bin/pairevtwait	pairevtwait
Error notification	/usr/bin/pairmon	pairmon
Volume checking	/usr/bin/pairvolchk	pairvolchk
Pair configuration confirmation	/usr/bin/pairdisplay	pairdisplay
RAID scan	/usr/bin/raidscan	raidscan
RAID activity report	/usr/bin/raidar	raidar
Connection confirmation	/usr/bin/raidqry	raidqry
<i>(continued)</i>		

Title	File name, Location	Executable for Command
Trace control	/usr/bin/horcctl	horcctl
Synchronization waiting command	/usr/bin/pairsyncwait	pairsyncwait
HORCM (RM) activation script	/usr/bin/horcmstart.sh	horcmstart.sh
HORCM shutdown script	/usr/bin/horcmshutdown.sh	horcmshutdown.sh
Connection confirmation	/HORCM/usr/bin/inqraid*	inqraid
Oracle validation setting	/usr/bin/raidvchkset	raidvchkset
Oracle validation confirmation	/usr/bin/raidvchkdsp	raidvchkdsp
Oracle validation confirmation	usr/bin/raidvchkscan	raidvchkscan
*The inqraid command is provided only for Linux, HP-UX, Solaris, MPE/iX, and OpenVMS.		

Title	File name, Location	Command file
HORCM (RM)	\HORCM\etc\horcmgr.exe	<i>none</i>
HORCM_CONF	\HORCM\etc\horcm.conf	<i>none</i>
Takeover	\HORCM\etc\horctakeover.exe	horctakeover
Make configuration file	\HORCM\etc\mkconf.exe	mkconf
Accessibility check	\HORCM\etc\paircurchk.exe	paircurchk
Pair generation	\HORCM\etc\paircreate.exe	paircreate
<i>(continued)</i>		

Title	File name, Location	Command file
Pair split/suspend	\HORCM\etc\pairsplit.exe	pairsplit
Pair resynchronization	\HORCM\etc\pairresync.exe	pairresync
Event waiting	\HORCM\etc\pairevtwait.exe	pairevtwait
Error notification	\HORCM\etc\pairmon.exe	pairmon
Volume checking	\HORCM\etc\pairvolchk.exe	pairvolchk
Pair configuration confirmation	\HORCM\etc\pairdisplay.exe	pairdisplay
RAID scanning	\HORCM\etc\raidscan.exe	raidscan
RAID activity reporting	\HORCM\etc\raidar.exe	raidar
Connection confirmation	\HORCM\etc\raidqry.exe	raidqry
Trace control	\HORCM\etc\horcctl.exe	horcctl
HORCM activation script	\HORCM\etc\horcmstart.exe	horcmstart
HORCM shutdown script	\HORCM\etc\horcmshutdown.exe	horcmshutdown
Synchronous waiting	\HORCM\etc\pairsyncwait.exe	pairsyncwait
Connection confirmation	\HORCM\etc\inqraid.exe	inqraid
Takeover	\HORCM\usr\bin\horctakeover.exe	horctakeover
Accessibility check	\HORCM\usr\bin\paircurchk.exe	paircurchk
Pair generation	\HORCM\usr\bin\paircreate.exe	paircreate
Pair split/suspend	\HORCM\usr\bin\pairsplit.exe	pairsplit
Pair resynchronization	\HORCM\usr\bin\pairresync.exe	pairresync
Event waiting	\HORCM\usr\bin\pairevtwait.exe	pairevtwait
Volume check	\HORCM\usr\bin\pairvolchk.exe	pairvolchk
Synchronization waiting command	\HORCM\usr\bin\pairsyncwait.exe	pairsyncwait
Pair configuration confirmation	\HORCM\usr\bin\pairdisplay.exe	pairdisplay
RAID scan	\HORCM\usr\bin\raidscan.exe	raidscan
<i>(continued)</i>		

Title	File name, Location	Command file
Connection confirmation	\HORCM\usr\bin\raidqry.exe	raidqry
Oracle validation setting	\HORCM\usr\bin\raidvchkset	raidvchkset
Oracle validation confirmation	\HORCM\usr\bin\raidvchkdsp	raidvchkdsp
Oracle validation confirmation	\HORCM\usr\bin\raidvchkscan	raidvchkscan
Tool	\HORCM\Tool\chgacl.exe	chgacl

Windows NT/2000/2003 command notes:

- \HORCM\etc\ commands are used when issuing commands interactively from the console.
- If the user issues these commands without any arguments, RM will enter interactive mode.
- \HORCM\usr\bin commands are not interactive.
- \HORCM\usr\bin commands are used for programatic execution from a user application.

Log files

RM and RM commands write internal logs and trace information to help the user identify the causes of RM failures and to keeps records of the transition history of pairs.

Log file format

Log files provided are the startup log file, error log file, trace file, and core file, which are located as shown below. HOST denotes the host name, and PID denotes the process ID within that host.

UNIX Systems

startup log files	HORCM startup log <code>\$HORCM_LOG/horcm_HOST.log</code>
	Command log <code>\$HORCC_LOG/horcc_HOST.log</code>
error log file	HORCM error log <code>\$HORCM_LOG/horcmllog_HOST/horcm.log</code>
trace files	HORCM trace <code>\$HORCM_LOG/horcmllog_HOST/horcml_PID.trc</code>
	Command trace <code>\$HORCM_LOG/horcmllog_HOST/horcc_PID.trc</code>
core files	HORCM core <code>\$HORCM_LOG/core_HOST_PID/core</code>
	Command core <code>\$HORCM_LOG/core_HOST_PID/core</code>

Windows NT/2000/2003 Systems

startup log files	HORCM startup log \$HORCM_LOG\horcm_HOST_log.txt
	Command log \$HORCC_LOG\horcc_HOST_log.txt
error log file	HORCM error log \$HORCM_LOG\horcmlog_HOST\horcm_log.txt
trace files	HORCM trace \$HORCM_LOG\horcmlog_HOST\horcm_PID_trc.txt
	Command trace \$HORCM_LOG\horcmlog_HOST\horcc_PID_trc.txt
core files	HORCM core \$HORCM_LOG\core_HOST_PID\core
	Command core \$HORCM_LOG\core_HOST_PID\core

MPE/iX Systems

startup log files	HORCM startup log \$HORCM_LOG/horcem_HOST.log
	Command log \$HORCC_LOG/horcc_HOST.log
error log file	HORCM error log \$HORCM_LOG/horcmllog_HOST/horcem.log
trace file	HORCM trace \$HORCM_LOG/horcmllog_HOST/horcem_PID.trc

OpenVMS Systems

startup log file	sys\$posix_root :[horcm.log]
------------------	-------------------------------------

Log directories

The log directories for the RM instance specify the command log files using the environment variables:

\$HORCM_LOG A trace log file directory specified using the environment variable **HORCM_LOG**. The HORCM (RM) log file, trace file and core file (as well as the command trace file and core file) are stored in this directory. If you do not specify an environment variable, **/HORCM/log/curlog** becomes the default.

\$HORCC_LOG A command log file directory specified using the environment variable **HORCC_LOG**. If you do not specify an environment variable, the directory **/HORCM/log n** (n is the instance number) becomes the default.

While CA is running, log files are stored in the **\$HORCM_LOG** directory. When RM starts up, the log files created are saved automatically in the **\$HORCM_LOGS** directory shown below.

RM “in operation” log file directory

\$HORCM_LOG = /HORCM/log n /curlog
 n is the instance number.

RM “automatic archives” log file directory

\$HORCM_LOGS = /HORCM/log n /tmplog
 n is the instance number.

Output to host log file

If you cannot create RM log files or an error occurs before the log files are created, error logs are output in the system log file.

If an RM instance activation fails, check the system log file, identify the error cause, and take any necessary action.

User-created files

When constructing the RM environment, the system administrator should make a copy of the **HORCM_CONF** file, edit the file for the system environment, and save the file:

UNIX

/etc/horcm.conf
or
/etc/horcm n .conf

where n is the instance number.

Windows NT/2000/2003

\WINNT\horcm.conf
or
\WINNT\horcm n .conf

where n is the instance number.

MPE/iX

/etc/horcm.conf
or
/etc/horcm n .conf

where n is the instance number.

OpenVMS

sys\$posix_root : [etc]horcm n .conf

where n is the instance number.

User-settable environment variables

When activating RM or initiating a command, you can specify any of the following environment variables:

- RM Environment Variables
- RM command Environment Variables
- RM instance Environment Variables
- environment variable for BC commands

RM environment variables

\$HORCM_CONF

Specifies the name of the RM configuration file.
Default: **/etc/horcm.conf**

\$HORCM_LOG Specifies the name of the RM log directory.
Default: **/HORCM/log/curlog**

\$HORCM_TRCSZ

Specifies the size of the RM trace file in kilobytes. The file size of 1 MB is used by default if no size is specified.

You cannot change the trace file size of RM in real time using the **horcctl** command.

\$HORCM_TRCLVL

Specifies the RM trace level (between 0 and 15). If you specify a negative value, the trace mode is canceled. If you do not specify a level, tracing becomes level 4 by default.

You can change the trace level of RM in real time by using the **horcctl -c -l** command.

\$HORCM_TRCBUF

Specifies the RM trace mode. If you specify this

environment variable, data is written to the trace file in nonbuffered mode. If you do not specify it, data is written in buffered mode.

The trace mode of RM can be changed in real time by using the **horcctl -c -b** command.

\$HORCM_TRCUENV

This variable specifies whether to use the trace control parameters (TRCLVL and TRCBUF trace types) as they are when a command is issued. When you specify this environment variable, the latest set trace control parameters are used. If you do not specify it, the default trace control parameters for RM commands are used, and tracing becomes level 4, and trace mode is set to buffer mode.

\$HORCM_FCTBL

This variable changes the fibre address conversion table number when the target ID, which is indicated by the **raidscan** command, is different from the target ID used by the host.

RM command environment variables

When issued, CA commands use these environment variables:

\$HORCC_LOG Specifies a command log directory name. If this variable is not specified, the following directory is used:

/HORCM/log*n*

(*n* is the instance number)

\$HORCC_TRCSZ

Specifies the size of the command trace file in kilobytes. If you do not specify a size, the default trace size for CA commands is used. This default trace size is the trace size used by CA.

The default trace size for CA commands can be changed in real time by using the **horcctl -d -s** command.

\$HORCC_TRCLVL

Specifies the command trace level (between 0 and 15). If you specify a negative value, the trace mode is canceled. If you do not specify a level, the default trace level for CA commands is used. This tracing is level 4 by default (or the CA level). You can change the default trace level for CA commands in real time using the **horcctl -d -l** command.

\$HORCC_TRCBUF

This variable specifies the command trace mode. If you specify this environment variable, data is written to the trace file in nonbuffer mode. If you do not specify it, the default trace mode for CA commands is used. This default tracing mode is buffered mode (or the CA trace mode). You can change the default trace mode for CA commands in real time using the **horcctl -d -b** command.

RM instance environment variable

The **\$HORCMINST** variable specifies the RM instance number when operating two or more RM instances on a single server.

You must specify an instance number in the command execution environment and the RM activation environment.

RM protection

The RAID Manager protection facility restricts RM volume control operations to volumes that:


- the host is allowed to see, with or without host-based LUN security (Secure LUN XP)
- are listed in the RM configuration file.

To avoid inconsistency, RM security cannot be controlled within RM itself.

RM security is determined by command device definition within the SVP, Remote Console, or via SNMP. Upon definition, the protection facility for each command device can be enabled by setting an attribute. RM refers to this attribute when it first recognizes the command device.

Command devices with protection ON permit access to volumes that are not only on their list of allowed volumes, but are also host viewable.

The following figure shows the definition of a protected (access refused) volume:



Protection facility specification

Only permitted volumes and volumes visible to the host can be listed in the **horcm.conf** file. A volume must fulfill two requirements to be considered “permitted” by the RM protection facility:

- It is host viewable (for example, with the HP supplied Inquiry tool).
- It is a volume listed in the **horcm.conf** file.

RM manages volume mirror descriptors (MU# for CA, BC0/BC1/BC2) as a unit. The permitted volumes in the following table lie at the intersection (the dark shaded areas) of volume device files, representing volumes that the host is allowed to see (row), and volumes that are listed in the **horcm.conf** file (column).

Volumes on horcm.conf	Mirror descriptor defined in horcm.conf								
	CA		BC0		BC1		BC2		
	L	not	L	not	L	not	L	not	
Unknown	N	N	N	N	N	N	N	N	N
/dev/rdisk/c0t0d0	P	N	P	N	P	N	P	N	N

Guide to table terms:

- L** The volume mirror descriptor is listed in the **horcm.conf** file.
- Not** The volume mirror descriptor is not listed in the **horcm.conf** file.
- Unknown** Volumes that the host cannot see, even though the volumes were listed in the **horcm.conf** file.
- P** Permitted
- N** Not permitted

Permission command

To allow initial access to a protected volume, the Permission command must be executed. This command is the **-find inst** option of **raidscan**; see “**raidscan**” (page 195). It is executed by **/etc/horcmgr** automatically upon RM startup. With security enabled, RM permits operations on a volume only after the Permission command is executed. Operations target volumes listed in the **horcm.conf** file.

The command compares volumes in the **horcm.conf** file to all host viewable volumes. Results are noted within RM in an internal table of protected and permitted volumes based on the **horcm.conf** file and the results of the Inquiry command. The Inquiry result is based on the LUN security for that host; you must configure LUN security before beginning RM operation.

Attempts to control protected volumes are rejected with the error code **EX_ENPERM**.

Protection facility support

XP12000	Protection facility supported.
XP1024/128	Protection facility supported.
XP512/48	Protection facility supported.
XP256	Not supported. XP256 can protect volumes only by using the protection mode of RM.

Sequent (Dynix/ptx)

Not supported. If a command device is set to enable protection mode, it is ignored by RM.

Digital UNIX

Not supported. If a command device is set to enable protection mode, it is ignored by RM.

MPE/iX Not supported (only SCSI connections). MPE/iX can protect volumes by using only the protection mode of RM.

OpenVMS Not supported. If a command device is set to enable protection mode, it is ignored by RM.

Command device configuration

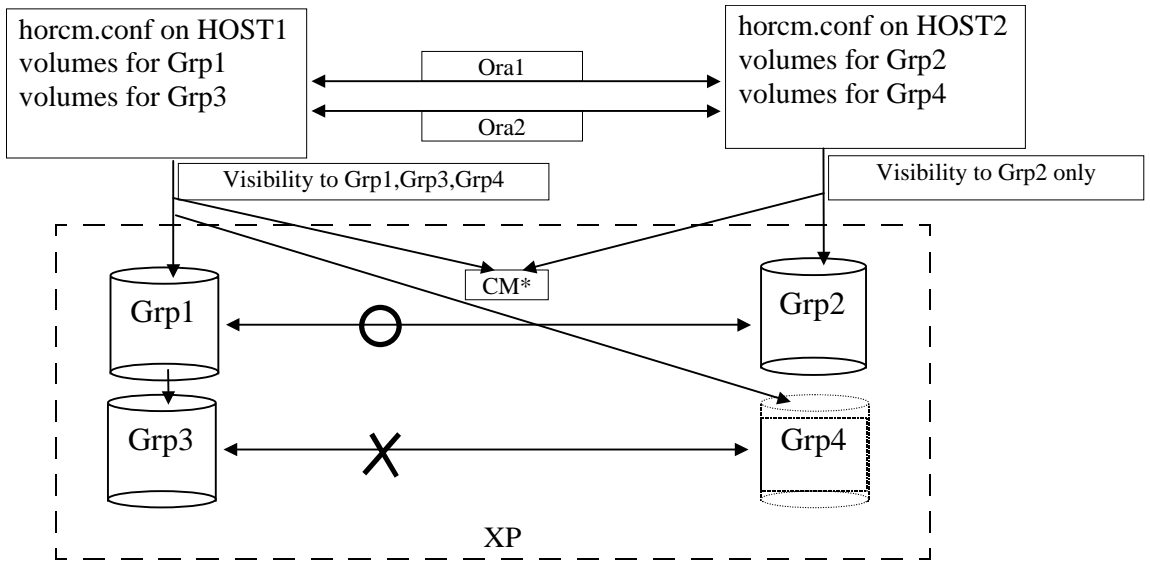
You can use both protected and unprotected modes in a single array by enabling or disabling the protection facility of each command device. As a minimum configuration, it is possible to have two command devices, one with protection enabled and the other disabled.

Protection mode is enabled for the host that uses a LUN security enabled command device.

LUN visibility from two host configuration

The following figure shows a two host protection mode configuration sharing one array. Operations directed at Ora2 are rejected because of no visibility for Grp4 from HOST2.

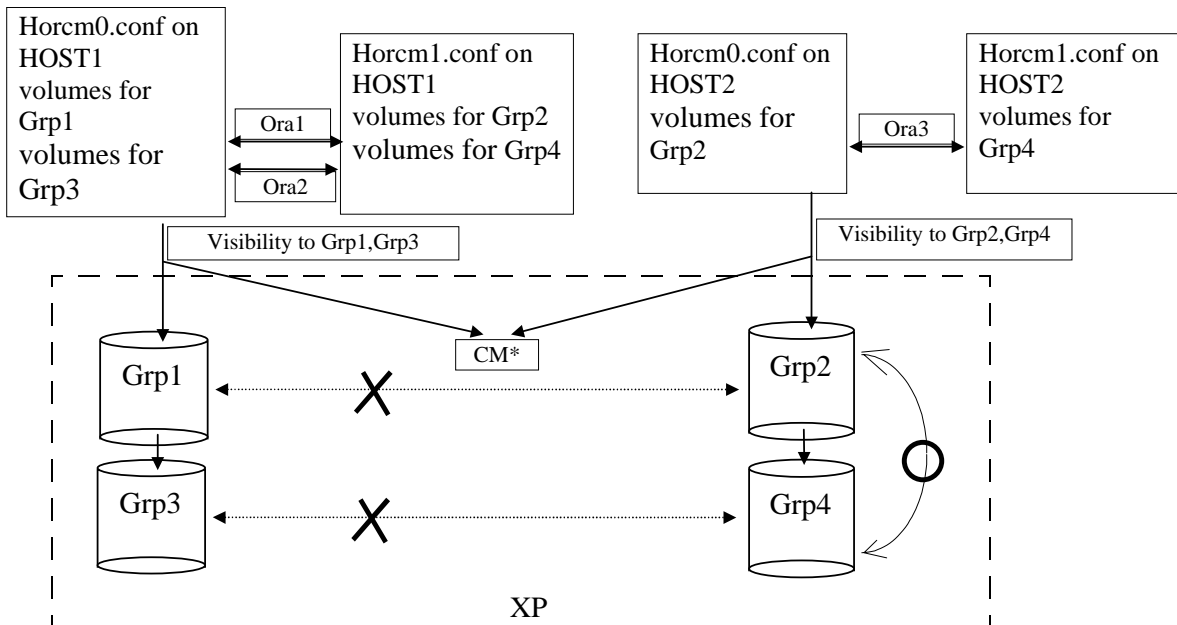
CM* represents a command device with protection set to ON at creation time.



LUN visibility from one host configuration

The following figure shows a one host protection mode configuration sharing one array. Ora1 and Ora2 control operations are rejected because of no visibility to Grp2 and Grp4 from HOST1. If HOST1 uses a command device with protection set to OFF at creation time, then Ora1 and Ora2 volume pairs can be controlled.

CM* represents a command device with protection ON.



Protection is implemented completely within RM. If RM uses a command device with protection enabled, RM will check permissions on all paired volumes with which it deals.

Commands controlled by RM protection

The following commands are controlled by RM protection:

- **horctakeover**, **paircurchk**, **paircreate**, **pairsplit**, **pairresync**, **pairvolchk**, **pairevtwait**, **pairsynwait**

When these commands are issued to non-permitted volumes, RM rejects the request with an error code of **EX_ENPERM**.

- **pairdisplay**

The **pairdisplay** command has no RM protection restrictions. Using this command, you can confirm whether volumes are permitted or not. Non-permitted volumes are shown without any LDEV# information. LDEV# information is shown as ****.

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-D , 3, 0-0) 35013 ****.-----,----- -
oradb  oradev1(R)  (CL1-D , 3, 1-0) 35013 ****.-----,----- -
```

- **raidscan**

The **raidscan** command shows all volumes without restriction because it does not use the HORCM_DEV and HORCM_INST fields in the **horcm.conf** file.

To identify permitted volumes with **raidscan**, use the **-find** option (supported with version 01.05.00). This option shows the device file name and array serial number information. You can use **raidscan -find** to create the **horcm.conf** file, because only permitted volumes (from the host's perspective) will be displayed.

Example (HP-UX)

```
# ioscan -fun | grep rdsd | raidscan -find
DEVICE_FILE      UID  S/F  PORT   TARG  LUN   SERIAL  LDEV  PRODUCT_ID
/dev/rdsd/c0t3d0  0    F    CL1-D   3     0     35013   17   OPEN-3
/dev/rdsd/c0t3d1  0    F    CL1-D   3     1     35013   18   OPEN-3
```

- **raidscan -find inst**

RM recognizes permitted volumes as a result of executing **raidscan -find inst** (the Permission command). This command issues a SCSI inquiry to the specified device file to get the array Ser# and volume LDEV# from the XP array. Then, it cross checks volumes in the **horcm.conf** file against host viewable volumes and stores the result within the RM instance.

The following example shows the relationship between device files and the **horcm.conf** file upon manual execution of the Permission command in an HP-UX environment. Operations to all volumes displayed by **ioscan** will be allowed.

*Example
(HP-UX)*

```
# ioscan -fun | grep rdsk | raidscan -find inst
DEVICE_FILE      Group   PairVol  PORT   TARG  LUN  M   SERIAL  LDEV
/dev/rdsk/c0t3d0  oradb  oradev1  CL1-D  3     0   -   35013   17
/dev/rdsk/c0t3d0  oradb  oradev1  CL1-D  3     0   0   35013   17
```

Permitting operations on protected volumes

Protection Mode requires a recognition step to check the host accessibility of volumes against the **horcm.conf** file at RM startup, using a command device with protection ON. This registration process (① or ②) is executed automatically by **/etc/horcmgr** at RM startup.

- ① The following is executed to register permitted volumes in a file (**\$HORCMPerm**). If the **\$HORCMPerm** file already exists, then it will use the existing file without doing a new **ioscan** (②).

If you want to permit even fewer volumes, edit the device file list in the **\$HORCMPerm** file. If you try to add device files that **ioscan** does not see (due to nonexistence or a LUN security product), an error will be returned at access time. This file is simply the text output (device files only) of a prior **ioscan** with the non XP device files removed.

Naming the \$HORCMPerm file

UNIX systems

The \$HORCMPerm variable is set by default to either `/etc/horcmperm.conf` or `/etc/horcmperm*.conf` (where * is the RM instance number).

```
Example (HP-UX) 'cat $HORCMPerm | /HORCM/usr/bin/raidscan -find inst'
# The following is an example to show permitted LVM
# Volume groups.
# For MU# 0
vg00 /dev/rdisk/c0t3d0 /dev/rdisk/c0t3d1
vg00 /dev/rdisk/c0t3d2 /dev/rdisk/c0t3d3

# For MU# 1
vg01 /dev/rdisk/c0t3d0 /dev/rdisk/c0t3d1
vg01 /dev/rdisk/c0t3d2 /dev/rdisk/c0t3d3
```

Verifying a group for vg01:

The following example verifies whether a LVM volume group is mapped to a group (MU#1 for BC) in the `horcm.conf` file correctly.

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

OR

```
# vdisplay -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/rdisk/c0t3d0 oradb1  oradev1   c0t3d0       1   35013   17
/dev/rdisk/c0t3d1 oradb1  oradev2   c0t3d1       1   35013   18
/dev/rdisk/c0t3d2 oradb   oradev3   c0t3d2       1   35013   19
/dev/rdisk/c0t3d3 -        -         -            1   35013   20
```

Windows NT/2000/2003 systems

The \$HORCMPerm variable is set by default to either `\WINNT\horcmperm.conf` or `\WINNT\horcmperm*.conf` (where * is the instance number).

```
'type $HORCMPerm | x:\HORCM\etc\raidscan.exe -find inst'
# The following is an example to permit DB Volumes.
# Note: a numerical value is interpreted as Harddisk#.
# DBO For MU# 0
```

```
Hd0-10
harddisk12 harddisk13 harddisk17

# DB1 For MU# 1
hd20-23
```

Verifying a group for DB1:

The following example shows how to verify whether a DB volume group is correctly mapped to a horcm.conf group (MU#1 for BC).

```
D:\HORCM\etc> set HORCC_MRCF=1
D:\HORCM\etc> echo hd20-23 | raidscan -find verify 1 -fd
DEVICE_FILE      Group   PairVol  Device_File  M   SERIAL  LDEV
Harddisk20       oradb1  oradev1  Harddisk20   1   35013   17
Harddisk21       oradb1  oradev2  Harddisk21   1   35013   18
Harddisk22       oradb   oradev3  Harddisk22   1   35013   19
Harddisk23       -       -        -            1   35013   20
```

- ② If no \$HORCMPerm file exists, then the following commands can be manually executed to permit the use of all volumes the host is currently allowed to see (LUN security products may or may not be in place).

HP-UX

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

Linux

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

Solaris

```
'ls /dev/rdisk/* | /HORCM/usr/bin/raidscan -find inst'
```

AIX

```
'lsdev -C -c disk | grep hdisk | /HORCM/usr/bin/raidscan -find inst'
```

If the **lsdev** command does not show the TID and LUN (for example, 2F-00-00-2,0) in the column output for the devices as shown below, then the **-d[g] raw device** option (on all commands) and **raidscan -find** will be unable to find target devices.

```
# lsdev -C -c disk
hdisk1  Defined    04-02-01      Other FC SCSI Disk Drive
```

This happens when a Fibre Channel adapter is used with a different device driver (for example, an Emulex adapter with an AIX driver).

MPE/iX

```
'callci dstat | /HORCM/usr/bin/raidscan -find inst'
```

Windows NT/2000/2003

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

The MAX volume to be scanned is **1000** by default.

Important This registration process is not without price because it is executed automatically upon **/etc/horcmgr** startup without checking for protection mode in order to validate the **-fd** option. Permitted volume registration brings a performance degradation in **horcmstart.sh** (RM startup), but the RM daemon runs as usual, depending on how many devices a host has.

If you want RM to start up faster in non-protection mode, then you can set **\$HORCMPerm** to **SIZE 0 byte** as a dummy file or set **HORCMPerm=MGRNOINST**. At that time, the **-fd** option will show the Device_File name as Unknown. Afterwards, you can use **raidscan -find inst** in order to validate the **-fd** option.

Environment variables

\$HORCMPROMOD

This environment variable sets protection mode ON by force. If your command device was created with protection mode OFF, this parameter forces protection mode ON, as shown in the table below.

Original Command Device Setting	HORCMPROMOD	Resulting Mode
Protection mode ON	No Effect (already on) →	Protection mode ON
Protection mode OFF	Variable not set →	Protection mode OFF
	Variable set →	Protection mode ON

\$HORCMPerm

This variable is used to specify the RM permission file. If no file name is specified, the default is **/etc/horcperm.conf**, or **/etc/horcperm*.conf** (where * is the instance number).

- If a RM permission file exists, then **/etc/horcmgr** executes the following command to permit the volumes listed in the file.

(HP-UX) `'cat $HORCMPerm | /HORCM/usr/bin/raidscan -find inst'`

(Windows
NT/2000/2003)

```
'type $HORCMPerm | x:\HORCM\etc\raidscan.exe -find inst'
```

- If no RM permission file exists, then **/etc/horcmgr** executes this built-in command to permit all volumes owned by the host.

(HP-UX)

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

(Linux)

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

(Solaris)

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

(AIX)

```
'lsdev -C -c disk | grep hdisk | /HORCM/usr/bin/raidscan -find inst'
```

(Tru64 UNIX)

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

(Digital UNIX)

```
'ls /dev/rrz* | /HORCM/usr/bin/raidscan -find inst'
```

(DYNIX/ptx)

```
'/etc/dumpconf -d | grep sd | /HORCM/usr/bin/raidscan -find inst'
```

(MPE/iX)

```
'callci dstat | /HORCM/usr/bin/raidscan -find inst'
```

(Windows
NT/2000/2003)

```
'x:\HORCM\etc\raidscan.exe -pi $PhysicalDrive -find inst'
```

- If \$HORCMPerm is set to MGRNOINST, **/etc/horcmgr** does not execute the built-in command.

This is used to execute a system command to permit the volumes specified from a user's shell script.

OpenVMS

```
$ inqraid dka145-146
```

Identifying a command device using protection mode

The SCSI inquiry output can not be changed to identify a command device in protection mode. Therefore, RM needs to provide a way to find the protection mode command device. The **horcctl -D -C** command designates a protection mode command device by appending '*' to the device file name as follows:

(HP-UX)

```
# horcctl -D  
Current control device = /dev/rdsk/c0t0d0*
```

Using RAID Manager on a Windows 2000/2003 system with “user” system privileges

By default, RAID Manager requires Windows system administrator privileges to execute RM commands. This is because RAID Manager needs to open the command device directly as a physical drive.

This section describes how to use the **chgacl.exe** to use RAID Manager commands without Administrator system privileges.

For a user to use “user” privileges to execute RM commands, the system administrator and the user need to perform some the following procedures.

Note: The parameters for the commands shown below are case sensitive.

Windows System Administrator

Allowing a user to use HORCM_CMD by adding a user name to the physical drive

As System Administration, add the user name to the Device Object of the command device in the configuration definition file. This will allow HORCM_CMD to work.

By default, **chgacl.exe** grants read, write and execute permissions. To restrict the permissions, see the section “Allowing different levels of access to a Device Object” ([page 93](#)).

To add a user name to all physical drives:

1. Enter **chgacl /A:<user_name> Phys**

Example

```
C:\HORCM\Tool>chgacl /A:RAdmin Phys
PhysicalDrive0 -> \Device\Harddisk0\DR0
\\.\PhysicalDrive0 : changed to allow 'RAdmin'
PhysicalDrive1 -> \Device\Harddisk1\DR1
\\.\PhysicalDrive1 : changed to allow 'RAdmin'
PhysicalDrive2 -> \Device\Harddisk2\DR2
\\.\PhysicalDrive2 : changed to allow 'RAdmin'
PhysicalDrive3 -> \Device\Harddisk3\DR3
\\.\PhysicalDrive3 : changed to allow 'RAdmin'
```

To add a user name to one or more physical drives:

1. Enter: **chgacl /A:<user_name> <object_name>...**

Example 1 **chgacl /A:RAdmin \\.\PHYSICALDRIVE10**

Example 2 **chgacl /A:RAdmin \\.\PHYSICALDRIVE10 \\.\PHYSICALDRIVE9**

Allowing a user to use the “-x mount/umount” option

If the user needs to use the “-x mount/umount” option of RM commands (for example, **raidscan -x mount Z: \vol2**), add the user name to the volume access control list.

By default, **chgacl.exe** grants read, write and execute permissions. To restrict the permissions, see the section “Allowing different levels of access to a Device Object” ([page 93](#)).

To add a user name to all volumes:

1. Enter **chgacl /A:<user_name> Volume**

Example

```
C:\HORCM\Tool>chgacl /A:RMadmin Volume
Volume{5d0f64b9-3327-11d7-80b8-0002e307aa91} -> \Device\HarddiskVolume9
\\.Volume{5d0f64b9-3327-11d7-80b8-0002e307aa91} : changed to allow 'RMadmin'
Volume{5d0f64b4-3327-11d7-80b8-0002e307aa91} -> \Device\HarddiskVolume6
\\.Volume{5d0f64b4-3327-11d7-80b8-0002e307aa91} : changed to allow 'RMadmin'
Volume{5d0f64b8-3327-11d7-80b8-0002e307aa91} -> \Device\HarddiskVolume8
\\.Volume{5d0f64b8-3327-11d7-80b8-0002e307aa91} : changed to allow 'RMadmin'
Volume{5d0f64b0-3327-11d7-80b8-0002e307aa91} -> \Device\HarddiskVolume2
\\.Volume{5d0f64b0-3327-11d7-80b8-0002e307aa91} : changed to allow 'RMadmin'
Volume{5f7d3408-ae0b-11d8-a414-0002e307aa91} -> \Device\HarddiskVolume11
\\.Volume{5f7d3408-ae0b-11d8-a414-0002e307aa91} : changed to allow 'RMadmin'
Volume{5d0f64ba-3327-11d7-80b8-0002e307aa91} -> \Device\HarddiskVolume10
\\.Volume{5d0f64ba-3327-11d7-80b8-0002e307aa91} : changed to allow 'RMadmin'
Volume{60c51fc2-2fb0-11d7-80b2-806d6172696f} -> \Device\Floppy0
\\.Volume{60c51fc2-2fb0-11d7-80b2-806d6172696f} : changed to allow 'RMadmin'
Volume{60c51fc3-2fb0-11d7-80b2-806d6172696f} -> \Device\CdRom0
```

To add a user name to one or more volumes:

1. Enter: `chgacl /A:<user_name> <Volume{GUID}> ...`

Example

```
chgacl /A:RMadmin
\\.Volume{7dd3ba6b-2f98-11d7-a48a-806d6172696f}
```

You can also use the `\\?\Volume{GUID}` format used by Windows commands such as `mountvol`.

Allowing a user to use the “-x portscan” option

If the user needs to use the “-x portscan” option of RM commands (for example, `raidscan -x mount portscan port0,20`), add the user name to the SCSI port access list.

To add a user name to all SCSI ports:

1. Enter: `chgacl /A:<user_name> Scsi`

Example

```
C:\HORCM\test>chgacl /A:RMadmin Scsi
Scsi3: -> \Device\Scsi\sym_hi2
\\. \Scsi3: : changed to allow 'RMadmin'
Scsi4: -> \Device\Scsi\ql22001
\\. \Scsi4: : changed to allow 'RMadmin'
Scsi0: -> \Device\Ide\IdePort0
\\. \Scsi0: : changed to allow 'RMadmin'
Scsi1: -> \Device\Scsi\aic78xx1
\\. \Scsi1: : changed to allow 'RMadmin'
Scsi2: -> \Device\Scsi\sym_hi1
\\. \Scsi2: : changed to allow 'RMadmin'
```

To add a user name to one or more SCSI ports:

1. Enter: **chgacl /A:<user_name> <SCSIX> ...**

Example 1 **chgacl /A:RMadmin Scsi0**

Example 2 **chgacl /A:RMadmin Scsi0 Scsi1 Scsi2**

Allowing different levels of access to a Device Object

chgacl.exe allows you to set a combination of read, write, execute or “all” access rights to a Device Object. If no permission parameter is given, **chgacl** grants “all” access to the Device Object.

An RM user needs read, write and execute rights to the command device in order to start a HORCM instance.

Enter: **chgacl /A:<user_name> [/P:A-R-W-E] <object_name> ...**

Example (Grant read and write access for the user horcm to all physical drives.)

chgacl /A:horcm /P:R-W Phys

Deleting a user name from the access control list of the Device Object

Caution: The first two commands below may delete the user's privileges to access the system drive (C:\).

To delete a user name from all physical drives:

1. Enter: `chgacl /D:<user_name> Phys`

To delete a user name from all volumes:

1. Enter: `chgacl /D:<user_name> Volume`

To delete a user name from one or more Device Objects:

1. Enter: `chgacl /D: <user_name> <object_name> ...`

Restrictions

The ACL (Access Control List) for the Device Object is set every time Windows boots, so access must be reset every time the system starts up.

Use the Windows Scheduled Tasks application to run a batch file that adds the RM user name to the access list when system reboots.

To add a scheduled task (Windows 2000/Windows 2003):

1. Click **Start**. Click **Control Panel**.
2. Double-click **Scheduled Tasks**. Double-click **Add Scheduled Task**. The Scheduled Task Wizard appears.
3. Click **Next**.
4. Browse to the batch file. Select the batch file.
A sample batch file is shown in this section.
5. Enter a name for the task. Select "When my computer starts"
6. Enter the system administrator password. Enter the password again.
7. Click **Next**. Click **Finish**.

You can redirect the output of the batch file by adding redirection in the batch file. Alternately, you can specify redirection in the Scheduled Task item's Run field in advanced properties (for example, C:\HORCM\add_RM_user.bat > C:\HORCM\logs\add_RM_user.log).

Note: If you change the Windows system administrator's password, this scheduled task will not execute. You will need to modify the task by entering the new password.

When new Device Objects (physical drives) are created, you must update user access for these devices.

RAID Manager user

Establishing the HORCM (/etc/horcmgr) startup environment

By default, the RM configuration definition file is stored in the "%SystemDrive%\windows" or "%SystemDrive%\WINNT" directory. A user with no system administrator privilege is denied writing to these directories. Therefore, the RM user needs to have his or her configuration file in some other directory and set the HORCM_CONF variable to that location.

Example 1 C:\HORCM\etc\>set HORCM_CONF=C:\Documents and Settings\RMadmin\horcm10.conf

C:\HORCM\etc\>set HORCMINST=10

C:\HORCM\etc\>horcmstart

horcmstart must be executed without arguments.

Example 2 Starting two instances:

```
C:\HORCM\etc>set HORCM_CONF=C:\Documents and Settings\RMadmin\horcm101.conf
C:\HORCM\etc>set HORCMINST=101
C:\HORCM\etc>horcmstart
starting HORCM inst 101
HORCM inst 101 starts successfully.
C:\HORCM\etc>set HORCM_CONF=C:\Documents and Settings\RMadmin\horcm5.conf
C:\HORCM\etc>set HORCMINST=5
C:\HORCM\etc>horcmstart
starting HORCM inst 5
HORCM inst 5 starts successfully.
```

Restrictions

Restriction 1. A user without system administrator privilege is not allowed to use the Windows **mountvol** command (although some current Windows 2000 revisions allow a user to mountvol a directory to a volume). Therefore, a user cannot execute the “directory mount” option of RM commands using the **mountvol** command.

For example, “raidscan -x mount C:\test \vol5” will generate an error even though the system administrator has added the user name to the access list of the volume.

Restriction 2. The **inqraid** “gvinf” option uses the “%SystemDrive%:\windows\” or “%SystemDrive%:\WINNT\” directory. Therefore, the user executing this command will get some errors unless the system administrator grants the user write access to the directory.

RM version 01.15.02 and later allows the user to set the HORCM_USE_TEMP variable to prevent the errors.

Example C:\HORCM\etc\>set HORCM_USE_TEMP=1
C:\HORCM\etc\>inqraid \$Phys -gvinf

Restriction 3. The user using the RAID Manager commands and the user starting the HORCM instance must have the same system privileges. The following scenario is an example:

An administrator stated a HORCM instance 5. User A with “user” privileges will not be able to use any RAID Manager commands with HORCM instance 5. This is because even if user A has been added to the access control list for the devices, user A’s RM commands will not be able to communicate with the HORCM instance that was started by another user with different privileges.

RM version 01.15.02 and later allow the user to connect to HORCM by setting the “HORCM_EVERYCLI” environment variable.

Example: C:\HORCM\etc\>set HORCM_CONF=C:\Documents and
Settings\RMadmin\horcm10.conf
C:\HORCM\etc\>set HORCMINST=10
C:\HORCM\etc\>set HORCM_EVERYCLI=1
C:\HORCM\etc\>horcmstart

horcmstart must be executed without arguments.

Sample BAT file

A batch file can be useful in reestablishing user access after a system reboot.

```
@echo off
echo Run at
Date /T
Time /T
rem (1) Allow a user to use HORCM_CMD by adding a
user_name to the physical drive
rem (1a) Add an user name to all physical drives
rem usage: chgac1 /A:<User_name> Phys
c:\horcm\tool\chgac1 /A:RMadmin Phys
rem (1b) Add a user name to one or more physical drives
rem usage: chgac1 /A:<User_name> <Object_name>...
c:\horcm\tool\chgac1 /A:RMadmin \\.\PHYSICALDRIVE0
\\.\PHYSICALDRIVE1
rem (2) Allowing an user to use -x mount/umount option
of the RM commands
rem (2a) Add the user name to ALL volumes
rem usage: chgac1 /A:<user_name> Volume
c:\horcm\tool\chgac1 /A:RMadmin Volume
rem (2b) Add the user name to one or more volumes
rem usage: chgac1 /A:<user_name> <Volume{GUID}>...
```

```
c:\horcm\tool\chgacl /A:RAdmin
\\.\Volume{7dd3ba6b-2f98-11d7-a48a-806d6172696f}
rem (3) Allow a user to use the "-x portscan" option of
RM commands
rem (3a) Add a user name0 to access list of ALL SCSI
ports
rem usage: chgacl /A:<user_name> Scsi
c:\horcm\tool\chgacl /A:RAdmin Scsi
rem (3b) Add the user name to Access List of one or more
SCSI ports
rem usage: chgacl /A:<user_name> <SCSIX>...
c:\horcm\tool\chgacl /A:RAdmin Scsi0 Scsi1
```

LUN Security Extension

HP StorageWorks LUN Security XP Extension is an optional feature that prevents hosts from writing to protected volumes. This is similar to the ORACLE Data Validation command, setting a protection attribute for a specified LU.

Guarding options

RAID Manager supports the following guarding options:

Hiding from inquiry commands. RM conceals the target volumes from SCSI Inquiry commands by responding “unpopulated volume” (0x7F) to the device type.

“SIZE 0” volume. RM replies to SCSI Read Capacity commands with “SIZE 0” for the target volume.

Read protection. RM protects volumes from reading by responding with the check condition of “Illegal function” (SenseKey=0x05, SenseCode=0x2200).

Write protection. RM protects volumes from writing by replying with “Write Protect” in the Mode sense header and by responding with the check condition of “Write Protect” (SenseKey=0x07, SenseCode=0x2700).

S-VOL disabling. RM protects volumes from becoming and S-VOL during pair creation.

Commands affected

RAID Manager has options in the following three commands for setting and verifying guarding:

raidvchkset. This command sets the parameters for guarding specified volumes. ([page 226](#))

raidvchkdsp. This command shows the guarding parameters for specified volumes, based on RM configuration file. ([page 233](#))

raidvchkscan. This command shows the guarding parameter for specified volumes, based on the **raidscan** command. ([page 239](#))

Notes and Restrictions

LUN Security Extension has the following restrictions.

File systems using LUN Security Extension

- In the case of UNIX file system volumes, the volumes must be mounted with the Read Only option by setting Open LDEV Guard after the volumes are unmounted.
- In the case of Windows2003 file systems, you have to use the “-x mount” and “-x amount” options of the above-cited RAID Manager commands.
- In the case of Windows NT or Windows 2000, Open LDEV Guard volumes set to Write Protect (read-only) mode cannot be used as NTFS or FAT file systems.

LVM(VxVM) and LUN Security Extension

When LVM volumes are to be used as Open LDEV Guard volumes, disable LUN Security Extension. Then LVM write commands may be issued to the volumes. Then re-enable Open LDEV Guard.

High Availability cluster server configurations

You should not use LUN Security Extension in HA environments, if HA cluster software will be writing metadata at regular intervals in order to confirm whether its disks are available or not.

Dynamic disk on Windows

LUN Security Extension volumes can not be used with the dynamic disk feature, because dynamic disk does not handle the volumes set to Write

Protect (read-only) mode. LUN Security Extension volumes must use Basic disk only.

License

The LUN Security Extension license key must be installed on the disk array.

Identifying Open LDEV Guard volumes

The inquiry page identifies LUN Security Extension volumes so the user does not use them as normal volumes.

Use **inqraid -fl** with the **-CLI** option. RM appends "*" to the device file name to identify a volume as a LUN Security Extension-protected volume:

Example # `ls /dev/rdsk/c57t4* | ./inqraid -CLI -fl`

DEVICE_FILE	PORT	SERIAL	LDEV	CTG	C/B/12	SSID	R:Group	PRODUCT_ID
c57t4d0*	CL1-D	62496	32	-	s/P/ss	0004	5:01-03	OPEN-3
c57t4d1*	CL1-D	62496	33	-	s/P/ss	0004	5:01-01	OPEN-3
c57t4d2*	CL1-D	62496	34	-	s/P/ss	0004	5:01-02	OPEN-3
c57t4d3*	CL1-D	62496	35	-	s/P/ss	0004	5:01-03	OPEN-3
c57t4d4	CL1-D	62496	36	-	s/P/ss	0004	5:01-01	OPEN-3
c57t4d5	CL1-D	62496	37	-	s/P/ss	0004	5:01-02	OPEN-3
c57t4d6	CL1-D	62496	38	-	s/P/ss	0004	5:01-03	OPEN-3

RAID Manager command reference

This chapter describes the function and syntax for all RM commands.

General commands

RM Command	Description	UNIX	Windows NT/2000/ 2003	MPE/iX	Page
horcctl	Changes and displays the RM internal trace and control parameters.	•	•	•	107
horcmshutdown	Stops RM.	•	•	•	110
horcmstart	A shell script that starts RAID Manager.	•	•	•	111
horctakeover	<i>(CA sync/async only)</i> The host executing horctakeover takes ownership of a pair.	•	•	•	112
inqraid	Displays array information.	•		•	116
mkconf	Makes a configuration file.	•	•	•	126
paircreate	Creates a pair.	•	•	•	129
paircurchk	<i>(CA sync/async only)</i> Checks the consistency of the data on the secondary volume.	•	•	•	136
pairdisplay	Confirms the configuration of a specified pair.	•	•	•	139
pairevtwait	Event waiting command.	•	•	•	149
pairmon	Monitors a pair and reports changes in the pair status.	•	•	•	155
pairresync	Resynchronizes a pair.	•	•	•	158
pairsplit	Splits or suspends a pair.	•	•	•	166
pairsyncwait	Synchronization waiting command.	•	•	•	173
pairvolchk	Checks the attributes of a volume connected to the local or remote hosts.	•	•	•	180
raidar	Reports the I/O activity of a specified LDEV.	•	•	•	189
raidqry	Confirms the connection of the disk array and the open system host.	•	•	•	192
raidscan	Lists the Fibre Channel port, target ID, LUN, and LDEV status.	•	•	•	195

Windows NT/2000/2003 commands

Windows NT/2000/2003 Command	Description	Page
drivescan	Displays the relationship between the hard disk number and physical drive.	208
env	Displays an environment variable.	210
findcmddev	Searches for the command device.	211
mount	Mounts a specified device.	212
portscan	Displays the physical device on a designated port.	215
setenv	Sets an environment variable.	217
sleep	Suspends execution.	218
sync	Flushes remaining unwritten data to the physical drive.	219
umount	Unmounts a specified device.	222
usetenv	Deletes an environment variable.	224

Data integrity check commands

Data Integrity Check Command	Description	Page
raidvchkset	Sets the parameters for validation checking on the specified volumes.	226
raidvchkdsp	Displays the parameters for validation checking on the specified volumes, based on the RM configuration file.	233
raidvchkscan	Displays the parameters for validation checking on the specified volumes, based on the raidscan command.	239

horcctl

Change and display RM internal trace and control parameters

Syntax horcctl -h

horcctl { -d | -c | -l *level* | -b *m* | -q | -s *size(KB)* | -t *type* | -z | -zx }

horcctl -S

horcctl { -D | -C | -u *unitid* }

horcctl { -ND | -NC | -g *group* }

Arguments	-h	Displays Help/Usage and version information.
	-d	Interprets the trace control arguments (-l <i>level</i> , -b <i>m</i> , -s <i>size(KB)</i> , -t <i>type</i>) following this argument as parameters for RM.
	-c	Interprets the trace control arguments (-l <i>level</i> , -b <i>m</i> , -t <i>type</i>) following this argument as parameters for the CA manager.
	-l <i>level</i>	Sets a trace level to the one specified in <i>level</i> . The effective level is between 0 and 15. Specifying a negative value cancels the trace mode. A negative value n is specified as --n , where n is any digit between 1 and 9. For example: horcctl -l --4 Level 4 is the default setting and must not be changed unless directed by an HP service representative. Setting a trace level to other than 4 can impact problem resolution if a program failure occurs. Levels 0 to 3 are for troubleshooting.

When a change option to the trace control parameter is specified, a warning message is displayed, and the command enters interactive mode.

- b** *m* Sets a trace level.

y specifies buffered mode.

n specifies synchronous mode.
- q** Terminates interactive mode and exits this command.
- s** *size(KB)* Changes the default trace buffer size, which is 1 MB, in units of 1,024 bytes.
- t** *type* Sets a trace type to the one specified in *type*. When a *type* is specified, only traces of the specified type are output. A value between 0 and 511 can be specified. Two or more values can be specified.
- z** (*Not for use with MPE/iX*) This option makes this command enter interactive mode.
- zx** (*Not for use with MPE/iX or OpenVMS*) This option prevents using RM in interactive mode.
- S** Shuts down RM.
- D** Displays the RM command device name currently used by RM.

If the command device is blocked due to the online maintenance (microprogram replacement) of the disk array, check the RM command device name before using this argument.
- C** Changes and displays the RM command device being used by the RM.

If the command device is blocked due to the online maintenance (microprogram replacement) of the disk

array, check the RM command device name before using this argument.

By using this argument again after completion of the online maintenance (microprogram replacement), the previous command device is reinstated.

-u *unitid* This argument is in effect when the **-D** or **-C** argument is specified. It specifies the unit ID of a command device as the target.

If this argument is omitted, unit IDs are **0**.

-ND **-g** *group* Displays the network address and port name being used by RM.

-NC **-g** *group* Changes the network address and port name being used by RM and displays the changed network address name.

-g *group* Specifies the group name written in the configuration definition file.

Description The **horcctl** command is used for maintenance (except for the **-S**, **-D**, **-C**, **-ND**, **-NC**, and **-g** arguments) and troubleshooting. When it is issued, the internal trace control parameters of the RM manager and RM commands are changed and displayed.

If the arguments **-l** *level*, **-b** *m*, **-s** *size(KB)*, or **-t** *type* are not specified, the current trace control parameters are displayed.

Example Entering **horcctl -D -C** identifies a protection mode command device by adding "*" to the device name as follows:

```
HP-UX # horcctl -D
Current control device = /dev/rdisk/c0t0d0*
```

horcmshutdown

Stop RM instances

Syntax horcmshutdown.sh [*inst. . .*]

horcmshutdown.exe [*inst. . .*]

Argument *inst* Indicates an instance number corresponding to the RM instance to be shut down.

When omitted, the command uses the value stored in the **HORCMINST** environment variable.

Description The **horcmshutdown** command is an executable for stopping RM instances.

horcmstart

Start RAID Manager instance

Syntax *HP-UX:*
horcmstart.sh [*instance . . .*]

Windows NT/2000/2003
horcmstart.exe [*instance . . .*]

MPE/iX
MPE/iX POSIX cannot launch a daemon process from a POSIX shell. Therefore, you must execute RM as a job in the background by using the **STREAM** command.

Argument *instance* Specifies the RM instance number. If omitted, the command uses the value stored in the **HORCMINST** environment variable. If **HORCMINST** is not set, a null value for instance is used to set the environment variables (**HORCM_COMF**, **HORCM_LOG**, **HORCM_LOGS**).

Description The **horcmstart** command is a executable which starts RM. If RM instance numbers are specified, this executable sets environment variables (**HORCM_CONF**, **HORCM_LOG**, **HORCM_LOGS**) and it starts RM instances.

Returned Values The **horcmstart** command sets either of the following returned values in **exit()**, which allows you to check the execution results.

The command returns **0** upon normal termination.

A nonzero return indicates abnormal termination. For the cause of the error and details, see the execution logs.

Files **/HORCM/loginstance/curlog/horcm_hostname.log**
/HORCM/loginstance/horcm_hostname.log

horctakeover

Take ownership of a pair

CA only

Syntax horctakeover { **-nomsg** | **-g** *group* | **-d** *pair_vol* | **-d[g]** *raw_device* [*MU#*] | **-d[g]** *seq# LDEV#* [*MU#*] | **-q** | **-S** | **-I** | **-t** *timeout* | **-z** | **-zx** }

Arguments **-d** *pair_vol* Specifies a paired volume name written in the configuration definition file. The command executes only for the specified paired volume.

-d[g] *raw_device* [*MU#*]

(*HP-UX, Linux, Solaris, MPE/iX, and Windows NT/2000/2003 only*) Searches the RM instance configuration file (local instance) for a volume that matches the specified *raw_device*. If a volume is found, the command is executed on the paired volume (**-d**) or group (**-dg**).

This option is effective without specification of the **-g** *group* option.

If the volume is contained in two groups, the command is executed on the first volume encountered. If *MU#* is not specified, it defaults to 0.

-d[g] *seq# LDEV#* [*MU#*]

Searches the RM instance configuration file (local instance) for a volume that matches the specified sequence number (*seq#*) and LDEV. If a volume is found, a command is executed on the paired logical volume (**-d**) or group (**-dg**). If the volume is contained in two groups, the command is executed on the first volume encountered. If *MU#* is not specified, it defaults to 0.

seq# is the array serial number. *seq# LDEV#* can be specified in hexadecimal (by addition of 0x) or decimal.

- g** *group* Specifies a group name in the RM instance configuration file.
- The command executes for the entire group unless the **-d** *pair_vol* argument is specified.
- h** Displays Help/Usage and version information.
- I** Executes a PVOL-takeover, which enables the P-VOL for reading and writing by a local host without a remote host. This argument is used when the primary volume is in status or data fence, is not allowing writes, and is in PSUE or PDUB state. If the primary volume is in any other state, then a nop-takeover is executed.
- The target volume of a local host must be a P-VOL.
- nomsg** Suppresses messages to be displayed. It is used to execute this command from a user program.
- This argument must be specified at the beginning of the command arguments.
- q** Terminates interactive mode and exits this command.
- S** Selects and executes an SVOL-takeover. The target volume of a local host must be an S-VOL. If this argument is specified, the **-I** argument is invalid.
- t** *timeout* (*Asynchronous paired volumes only*) Specifies the maximum time in seconds to wait for a resynchronization of P-VOL to S-VOL delta data. If the timeout occurs, **EX_EWSTOT** is returned. This option is required for an asynchronous paired volume; it has no effect for a synchronous paired volumes.
- z** Makes this command enter interactive mode.
- zx** (*Not for use with MPE/iX or OpenVMS*) Prevents using RM in interactive mode.

Description The **horctakeover** meta command (contains many sub-commands) is used in conjunction with HA software, such as MC/Service Guard and CA. It selects and executes one of four actions, depending on the state of the paired volumes: nop-takeover, swap-takeover, SVOL-takeover, or PVOL-takeover.

See [“Takeover-switch function” on page 332](#) for actions taken by **horctakeover**.

The table under the heading [“HA control script state transitions” on page 312](#) lists state transitions resulting from the execution of **horctakeover** in HA control scripts.

Returned Values The **horctakeover** command returns one of the following values in **exit()**.

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
(This returned value depends on microcode level.)

Abnormal termination

Other than the above. For the error cause and details, see the execution logs.

Execution Log File After the SVOL-takeover is executed, if inconsistent volumes exist, they are displayed in the execution log file.

Error Codes The table below lists specific error codes for the **horctakeover** command.

Category	Error Code	Error Message	Value
Volume status unrecoverable	EX_ENQVOL	Unmatched volume status within the group	236
	EX_INCSTG	Inconsistent status in group	229
	EX_EVOLCE	Pair volume combination error	235
	EX_VOLCUR	S-VOL currency error	225
	EX_VOLCUE	Local volume currency error	224
	EX_VOLCRE	Local and remote volume currency error	223
Timer recoverable	EX_EWSTOT	Timeout waiting for specified status	233

To recover from EX_EWSTOT:

If **horctakeover** fails with the EX_EWSTOT error, follow these steps:

1. Wait until the S-VOL state becomes SVOL_PSUS by using the return code of the **pairvolchk -g group -ss** command. Then, attempt the startup again from the HA control script.
2. Attempt to resynchronize the original P-VOL, based on the S-VOL, by using the **pairresync -g group -swaps -c size** command for a Fast Failback operation.

The operation in step 2 may fail with EX_CMDRJE or EX_CMDIOE. This will cause an ESCON link and/or site failure.

If this operation fails, the HA control script will report the following message:

After a recovery from failure, please try the **pairresync -g group -swaps -c size** command.

To avoid the above recovery steps, the timeout value should be just less than (for example, 30 seconds) the startup timeout value for the HA control script.

inqraid

Display array information

HP-UX, Linux, Solaris, AIX, and MPE/iX only

Syntax `inqraid { -h | quit | -inqdump | -f[x][p][l][g][w] | -find[c] | special_file | -CLI [W|WP|WN] | -sort | -CM | -inst | -gvinf | -svinf | -fv | -fp | fg | -gplba }`

Arguments	-h	Displays Help/Usage.
	-quit	Terminates interactive mode and exits the command.
	-inqdump	Displays SCSI information in hexadecimal dump format.
	-fx	Displays the LDEV number in hexadecimal format.
	-fp	Shows an Oracle validation volume with the -CLI option by appending "*" to the device file name.
	-fl	Shows an LDEV Guard volume with the -CLI option by appending "*" to the device file name.
	-fv	(Windows NT/2000/2003 only) Displays the Volume{GUID} via \$Volume in wide format.
	-fg	Displays a LUN in the host view by finding a host group for XP 128 and XP 1024 arrays.
	-fw	Displays the cascading volume status on STD Inquiry Page. If this option will not be specified, then the display will show four cascading mirrors the same as at present in order to maintain compatibility with the current -CLI option.

```
# ls /dev/rdisk/* | inqraid -CLI -fw
DEVICE_FILE    PORT      SERIAL  LDEV  CTG  C../B/..          SSID R:Group  PRODUCT_ID
c1t2d10s2      CL2-D     62500   266   -    Psss/P/PP-----  0005 1:01-02  OPEN-3-SUN
c1t2d11s2      CL2-D     62500   267   -    ssss/s/ss-----  0005 1:01-02  OPEN-3-SUN
```

-find[c]

Using device special file names provided via STDIN, this option displays information about the corresponding configuration file volume groups through the use of the **inquiry** and **pairedisplay** commands.

This option requires that the **HORCMINST** variable be defined in the command execution environment.

The **-find** option employs the following options of the **pairedisplay** command:

(BC)

```
pairedisplay -d <Seq#> <LDEV#> 0 1 2 -1 [-fx] [-CLI] 2>/dev/null
```

(CA)

```
pairedisplay -d <Seq#> <LDEV#> -1 [-fx] [-CLI] 2>/dev/null
```

The **-findc** option employs the following options of the **pairedisplay** command, and presents the output in an easily parsed CLI format.

(BC)

```
pairedisplay -d <Seq#> <LDEV#> <MU#> -fd -CLI 2>/dev/null
```

(CA)

```
pairedisplay -d <Seq#> <LDEV#> -fd -CLI 2>/dev/null
```

The Seq# and LDEV# are provided via the SCSI Inquiry command.

This option requires the HORCMINST variable to be defined.

special_file

Specifies a device special file name as an argument to the command. If no argument is specified, the command waits for input from STDIN. For STDIN file specification information, see Appendix D, “[STDIN file formats](#)”.

- CLI** Specifies structured output for Command Line Interface parsing. The column data is aligned in each row. The delimiters between columns are either a space or -.
- CLI[W|WP|WN]** (*Not for use with Tru64 or Digital UNIX*) Displays the WWN (World Wide Name) and LUN in CLI format.
- sort[-CM]** Sorts the target devices in Serial#, LDEV# order.
- The **-sort-CM** option displays the command devices listed in the horcm.conf file.
- A unitID is displayed with the Serial#.
- When two or more command devices exist, this option will show multiple device files linked to a command device (an LDEV).
- inst** (*MPE/iX only*) Used to make a device special file such as /dev/ldev* by getting an LDEV number (ldev in MPE/iX terminology) from the “LDEV-” field in the **dstat** command.
- gvinf** (*Windows NT/2000/2003 only*) Retrieves the LUN signature and volume layout information by way of a raw device file provided via STDIN or arguments. This information is saved to a system disk file with the format \WindowsDirectory\VOLssss_IIII.ini where **ssss** is the array serial number and **IIII** is the LDEV#.
- Normally, this option is used to save the LUN signature and volume layout information after it has been written created (and before a **paircreate**).
- svinf[=PTN]** (*Windows NT/2000/2003 only*) Uses SCSI Inquiry to get the Serial# and LDEV# created by **-gvinf** of the RAID for the target device, and sets the signature and volume layout information in file VOLssss_IIII.ini to the target device.

This option will complete correctly even if the `Harddisk#` is changed by the operating system. The signature and volume layout information is managed by the `Serial#` and `LDEV#` of RAID.

The `-svinf=PTN` option specifies a string pattern to select only the pertinent output lines being provided from STDIN.

This option returns `0` upon normal termination. A nonzero return indicates abnormal termination.

-gplba (*Windows NT/2000/2003 only*) Displays usable LBA on physical drives.

Description The `inqraid` command displays the relationship between a host device special file and an actual physical drive in the disk array.

Fields	<code>CLX-Y</code>	Displays the port number of the disk array.
	<code>Ser</code>	Displays the production (serial#) number of the disk array.
	<code>LDEV</code>	Displays an <code>LDEV#</code> within the disk array.
	<code>CA</code>	Displays the attribute (PVOL/SVOL/SMPL) of a volume as CA.
	<code>BC</code>	Displays the attribute (PVOL/SVOL/SMPL) of a volume as BC.
	<code>Group</code>	Displays the physical position of a LDEV in relation to a RAID group in the disk array.
	<code>SSID</code>	Displays the Sub System ID of an LDEV in the disk array.
	<code>CTGID</code>	Displays the CT group ID when the LDEV has been specified as an <code>async-CA P-VOL</code> or <code>S-VOL</code> .

CHNO	(Linux only) Displays the Linux channel number of the device adapter.
TID	(Linux only) Displays the target ID of the hard disk connected to the device adapter port. See Appendix , “Fibre Channel addressing” .
LUN	(Linux only) Displays the logical unit number of the hard disk that connects on the device adapter port.
DEVICE_FILE	Displays the device file name.
M	Displays the MU# of local and remote volumes.
PairVol	Displays the paired volume name (dev_name) within the group defined in the configuration file.
P/S	Displays the volume attribute (P-VOL, S-VOL, or Simplex).
Stat	Displays the status of the paired volume.
R_DEVICE	Displays the device file name of the remote half of the pair.
LK	Indicates the results of a check on the paired volume connection path.
PORT	Displays the disk array port number.
C/B/12	Corresponds to CA volume/BC volume/BC MU#1,2. Displays attributes where: P = P-VOL S = S-VOL s = SMPL
R:Group	Displays the physical position of an LDEV in relation to a RAID group in the disk array.
PRODUCT_ID	Displays the product ID field in the STD inquiry page.

PWWN	Displays the port WWN.
NWWN	Displays the Node WWN.
AL	(Not supported for the Tru64 or Digital UNIX -CLIWP and -CLIWN options) Displays the AL_PA of the port.

Examples

Examples using the **-find** option:

Linux

```
ls /dev/sd* | inqraid -find
Group PairVol(L/R) (Port#,TID,LU),Seq#,LDEV#.P/S,Status,Fence, Seq#,P-LDEV# M
oradb oradev2(L) (CL2-N , 3, 2) 8071 22..SMPL -----,----- - -
->/dev/sdc
```

HP-UX

```
# echo /dev/rdisk/c23t0d0 /dev/rdisk/c23t2d3 | ./inqraid -find
Group PairVol(L/R) (Port#,TID,LU-M),Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M
horc1 dev00(L) (CL2-J , 0, 0-0)61456 192..S-VOL SSUS,----- 193 -
->/dev/rdisk/c23t0d0
Group PairVol(L/R) (Port#,TID,LU-M),Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M
horc1 dev10(L) (CL2-J , 2, 3-0)61456 209..S-VOL SSUS,----- 206 -
->/dev/rdisk/c23t2d3
```

Examples using the **-findc** option:

HP-UX

```
# echo /dev/rdisk/c23t0d0 /dev/rdisk/c23t2d3 | ./inqraid -findc
DEVICE_FILE M Group PairVol P/S Stat R_DEVICE M P/S Stat LK
c23t0d0 0 horc1 dev00 S-VOL SSUS c23t0d1 0 P-VOL PSUS OK
/dev/rdisk/c23t0d0[1] -> No such on the group
/dev/rdisk/c23t0d0[2] -> No such on the group
DEVICE_FILE M Group PairVol P/S Stat R_DEVICE M P/S Stat LK
c23t2d3 0 horc1 dev10 S-VOL SSUS c23t2d2 0 P-VOL PSUS OK
/dev/rdisk/c23t2d3[1] -> No such on the group
/dev/rdisk/c23t2d3[2] -> No such on the group

# echo /dev/rdisk/c23t0d0 /dev/rdisk/c23t2d3 | ./inqraid -findc -CLI
DEVICE_FILE M Group PairVol P/S Stat R_DEVICE M P/S Stat LK
c23t0d0 0 horc1 dev00 S-VOL SSUS c23t0d1 0 P-VOL PSUS OK
c23t2d3 0 horc1 dev10 S-VOL SSUS c23t2d2 0 P-VOL PSUS OK
```

An example using the **-CLI** option:

Linux

```
# ls /dev/sd* | ./inraid -CLI
DEVICE_FILE  PORT    SERIAL  LDEV CTG   C/B/12  SSID R:Group  PRODUCT_ID
sdh          CL2-B   30053   23   2    S/P/ss  0004 5:02-01  OPEN-3
sdi          CL2-B   30053   14   -    -        -      -        OPEN-3-CM
sdj          -       -       -    -    -        -      -        -
```

An example using the **-CLIW** option:

```
# echo /dev/rdisk/c23t0d0 /dev/rdisk/c23t0d1 | ./inraid -CLIW
DEVICE_FILE  WWN                AL PORT  LUN   SERIAL  LDEV  PRODUCT_ID
c23t0d0      500060e802f01018  C3 CL2-J  0     61456   192  OPEN-3
c23t0d1      500060e802f01018  C3 CL2-J  1     61456   193  OPEN-3

DEVICE_FILE  WWN                AL PORT  LUN   SERIAL  LDEV  PRODUCT_ID
c0t2d3       5000E000E0005000  3E CL1-A  20    30015   2054 OPEN3-CVS
```

An example using the **-sort -CM** option:

HP-UX

```
#ioscan -fun | grep rdsk | ./inraid -sort -CM -CLI
HORCM_CMD
#dev_name          dev_name          dev_name
#UnitID 0 (Serial# 30012)
/dev/rdsk/c0t3d0   /dev/rdsk/c1t2d1
#UnitID 1 (Serial# 30013)
/dev/rdsk/c2t3d0
```

An example using the **-inst** option:

MPE/iX

```
shell/iX> callci dstat | ./inraid -inst -CLI
DEVICE_FILE  PORT    SERIAL  LDEV CTG   C/B/12  SSID R:Group  PRODUCT_ID
ldev100      CL1-L   35013   19   -    s/s/ss  0004 5:01-01  OPEN-3
ldev101      CL1-L   35013   35   -    -        -      -        OPEN-3-CM
```

An example using the **-gvinf** option follows. This example saves the volume information for all physical drives.

```
D:\HORCM\etc>inraid $Phys -gvinf -CLI
\\.\PhysicalDrive0:
# Harddisk0 -> [VOL61459_448_DA7C0D91] [OPEN-3 ]
\\.\PhysicalDrive1:
# Harddisk1 -> [VOL61459_449_DA7C0D92] [OPEN-3 ]
\\.\PhysicalDrive2:
# Harddisk2 -> [VOL61459_450_DA7C0D93] [OPEN-3 ]
```

An example using the **-svinf=PTN** follows. This example writes signature/volume information to LUNs identified by “Harddisk” in the output of the **pairdisplay** command.

```
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 61459 451..S-VOL SSUS,----- 448 -
URA URA_001 (L) Harddisk4 0 61459 452..S-VOL SSUS,----- 449 -
URA URA_002 (L) Harddisk5 0 61459 453..S-VOL SSUS,----- 450 -
```

```
D:\HORCM\etc>pairdisplay -l -fd -g URA | inqraid -svinf=Harddisk
[VOL61459_451_5296A763] -> Harddisk3 [OPEN-3 ]
[VOL61459_452_5296A760] -> Harddisk4 [OPEN-3 ]
[VOL61459_453_5296A761] -> Harddisk5 [OPEN-3 ]
```

Additional Information If you create an S-VOL with the “Noread” option and reboot the Windows 2000/2003 system, the system will be unable to create a Device object (\Device\HarddiskVolume#) and Volume {GUID} for that SVOL. A Device object (\Device\HarddiskVolume#) and Volume{GUID} can be created, using the **-svinf** option to the **inqraid** command (on a suspended SVOL).

\Device\HarddiskVolume#(number) is assigned in sequential order by the **-svinf** option. This number will be valid as long as the system configuration does not change.

Use the **-svinf -sort** option to cause signature writes to occur in LDEV# order as follows.

```
D:\HORCM\etc>echo hd5 hd4 hd3 | inqraid -svinf -sort
[VOL61459_451_5296A763] -> Harddisk3 [OPEN-3 ]
[VOL61459_452_5296A760] -> Harddisk4 [OPEN-3 ]
[VOL61459_453_5296A761] -> Harddisk5 [OPEN-3 ]
```

An example using the **-fv** option.

```
C:\HORCM\etc>inqraid -CLI $Vol -fv
DEVICE_FILE PORT SERIAL LDEV
CTG H/M/12 SSID R:Group PRODUCT_ID
Volume{cec25efe-d3b8-11d4-aead-00c00d003b1e}\Vol13\Dsk0 CL2-D 62496 256
- - - - OPEN-3-CVS-CM
```

An example using the **-fp** option:

```
# ls /dev/rdsk/c57t4* | ./inqraid -CLI -fp
DEVICE_FILE PORT SERIAL LDEV CTG C/B/12 SSID R:Group PRODUCT_ID
c57t4d0* CL1-D 62496 32 - s/P/ss 0004 5:01-03 OPEN-3
```

```

c57t4d3*    CL1-D    62496    35    -    s/P/ss    0004    5:01-03    OPEN-3
c57t4d4    CL1-D    62496    36    -    s/P/ss    0004    5:01-01    OPEN-3
c57t4d5    CL1-D    62496    37    -    s/P/ss    0004    5:01-02    OPEN-3

```

The following examples display the relationship between a special file and the actual physical drive in the disk array, by using the **inqraid** and system commands.

```

HP-UX # ioscan -fun | grep rdsk | ./inqraid
/dev/rdsk/c0t2d0 ->[HP] CL2-D Ser = 30053 LDEV = 8 [HP ] [OPEN-3 ]
          CA = SMPL BC[MU#0 = SMPL MU#1 = SMPL MU#2 = SMPL]
/dev/rdsk/c0t2d1 ->[HP] CL2-D Ser = 30053 LDEV = 9 [HP ] [OPEN-3 ]
          CA = SMPL BC[MU#0 = SMPL MU#1 = SMPL MU#2 = SMPL]
/dev/rdsk/c0t4d0 ->[HP] CL2-D Ser = 30053 LDEV = 14 [HP] [OPEN-3 CM]

```

```

Linux # ls /dev/sd* | ./inqraid
/dev/sdg ->CHNO = 0 TID = 1 LUN = 6
          [HP] CL2-B Ser = 30053 LDEV = 22 [HP ] [OPEN-3 ]
          CA = SMPL BC[MU#0 = SMPL MU#1 = SMPL MU#2 = SMPL]
/dev/sdh ->CHNO = 0 TID = 1 LUN = 7
          [HP] CL2-B Ser = 30053 LDEV = 23 [HP ] [OPEN-3 ]
          CA = SMPL BC[MU#0 = SMPL MU#1 = SMPL MU#2 = SMPL]
/dev/sdi ->CHNO = 0 TID = 4 LUN = 0
          [HP] CL2-B Ser = 30053 LDEV = 14 [HP ] [OPEN-3-CM ]

```

```

Solaris # ls /dev/rdsk/* | ./inqraid
/dev/rdsk/c0t2d1 -> [HP] CL2-D Ser = 30053 LDEV = 9 [HP ] [OPEN-3 ]
          CA = P-VOL BC[MU#0 = SMPL MU#1 = SMPL MU#2 = SMPL]
          RAID5[Group 2- 1] SSID = 0x0008 CTGID = 3
/dev/rdsk/c0t4d0 -> [HP] CL2-D Ser = 30053 LDEV = 14 [HP ] [OPEN-3-CM ]

```

```

MPE/iX shell/iX>ls /dev/* | ./inqraid 2>/dev/null
/dev/ldev009 -> [HP] CL2-D Ser = 30053 LDEV = 9 [HP ] [OPEN-3 ]
          CA = P-VOL BC[MU#0 = SMPL MU#1 = SMPL MU#2 = SMPL]
          RAID5[Group 2- 1] SSID = 0x0008 CTGID = 3
/dev/cmddev -> [HP] CL2-D Ser = 30053 LDEV = 14 [HP ] [OPEN-3-CM ]

```

```

AIX # lsdev -C -c disk | grep hdisk | ./inqraid
hdisk1 -> [HP] CL2-D Ser = 30053 LDEV = 9 [HP ] [OPEN-3 ]
          CA = P-VOL BC[MU#0 = SMPL MU#1 = SMPL MU#2 = SMPL]
          RAID5[Group 2- 1] SSID = 0x0008 CTGID = 3
hdisk2 -> [HP] CL2-D Ser = 30053 LDEV = 14 [HP ] [OPEN-3-CM ]

```

Additional Information

If the **lsdev** command does not show the TID and LUN (for example, 2F-00-00-2,0) on the column output for the devices as shown below, then the **inqraid** command and **-d[g] raw_device** option for all commands will be unable to find a target device.

```

# lsdev -C -c disk
hdisk1 Defined 04-02-01 Other FC SCSI Disk Drive

```

This occurs when a Fibre Channel adapter and device driver are different (for example, an Emulex adapter with an AIX driver).

*Windows
NT/2000/2003*

```
C:\HORCM\etc> echo hd1-2 | inqraid ( or inqraid hd1-2 )
Harddisk 1 -> [HP] CL2-D Ser = 30053 LDEV = 9 [HP ] [OPEN-3 ]
                CA = P-VOL BC[MU#0 = SMPL MU#1 = SMPL MU#2 = SMPL]
                RAID5[Group 2- 1] SSID = 0x0008 CTGID = 3
Harddisk 2 -> [HP] CL2-D Ser = 30053 LDEV = 14 [HP ] [OPEN-3-CM ]
```

Tru64

```
# ls /dev/rdisk/dsk* | ./inqraid
/dev/rdisk/dsk10c -> [HP] CL2-D Ser = 30053 LDEV = 9 [HP] [OPEN-3 ]
                CA = P-VOL BC[MU#0 = SMPL MU#1 = SMPL MU#2 = SMPL]
                RAID5[Group 2- 1] SSID = 0x0008 CTGID = 3
/dev/rdisk/dsk11c -> [HP] CL2-D Ser = 30053 LDEV = 14 [HP] [OPEN-3-CM]
```

DYNIX/ptx

```
# dumpconf -d | grep sd | ./inqraid
sd1 -> [HP] CL2-D Ser = 30053 LDEV = 9 [HP ] [OPEN-3 ]
                CA = P-VOL BC[MU#0 = SMPL MU#1 = SMPL MU#2 = SMPL]
                RAID5[Group 2- 1] SSID = 0x0008 CTGID = 3
Sd2 -> [HP] CL2-D Ser = 30053 LDEV = 14 [HP ] [OPEN-3-CM ]
```

mkconf

Make a configuration file

Syntax `mkconf.sh { -g group | -m MU# | -i inst# | -s service | -a }`

Windows NT/2000/2003 or OpenVMS only:

`mkconf.exe { -g group | -m MU# | -i inst#] | -s service | -a | -c drive# }`

Arguments	(none)	Using the mkconf command without any arguments displays help/usage information.
	<code>-g group</code>	Specifies the dev_group name to be used within the newly created configuration file. If not specified, VG will be used.
	<code>-m MU#</code>	Specifies the mirror descriptor for BC volumes. CA volumes do not specify a mirror descriptor.
	<code>-i inst#</code>	Specifies the instance number for RM.
	<code>-s service</code>	Specifies the service name to be used in the newly created configuration file. If not specified, 52323 will be used as a default.
	<code>-a</code>	Used to add a new volume group within the newly created configuration file.
	<code>-c drive#</code>	<i>(Windows NT/2000/2003 only)</i> Specifies the range of drive numbers to be searched for existing command devices. If not specified, hd0-99 will be used as the default.
	<code>-c <DKA#-#</code>	<i>(OpenVMS only)</i> Specifies the range of drive numbers to be searched for existing command devices. If not specified, DKA0-10000 DGA0-10000 will be used as the default.

Description The **mkconf** command is used to make a configuration file from a special file (raw device file) provided via STDIN. It executes the following steps:

1. Make a configuration file containing only the HORCM_CMD section by executing **inraid -sort -CM -CLI**.
2. Start a RM instance without a HORCM_DEV and HORCM_INST section, which is just enough to execute the **raidscan** command for the next step.
3. Make a configuration file including the HORCM_DEV and HORCM_INST sections by executing **raidscan -find conf** using a special file (raw device file) provided via STDIN. For STDIN file specification information, see Appendix D, “[STDIN file formats](#)” .
4. Start a RM instance using the newly created configuration file.
5. Execute **raidscan -find verify** to verify the correspondence between host device files and the newly created configuration file.

The configuration file is created with the name **horcm*.conf** within the current directory. A RM log directory is created with the name **log*** within the current directory.

You may have to modify the *ip_address* and *service* parameters within the newly created configuration file as the need arises.

Example This example demonstrates the usage of the **mkconf** command and the resulting configuration file.

HP-UX

```
# cd /tmp/test
# cat /etc/horcmperm.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/rdisk/c23t0d0	ORA	ORA_000	CL2-J	0	0	0	61456	192
/dev/rdisk/c23t0d1	ORA	ORA_001	CL2-J	0	1	0	61456	193
/dev/rdisk/c23t0d2	ORA	ORA_002	CL2-J	0	2	0	61456	194
/dev/rdisk/c23t0d3	ORA	ORA_003	CL2-J	0	3	0	61456	195
/dev/rdisk/c23t0d4	ORA	ORA_004	CL2-J	0	4	0	61456	256
/dev/rdisk/c23t0d5	ORA	ORA_005	CL2-J	0	5	0	61456	257
/dev/rdisk/c23t0d6	ORA	ORA_006	CL2-J	0	6	0	61456	258
/dev/rdisk/c23t0d7	-	-	-	-	-	0	61456	259

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 *.conf
```

Configuration file:

```
# 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            3000

HORCM_CMD
#dev_name        dev_name      dev_name
#UnitID 0 (Serial# 61456)
/dev/rdsk/c23t3d0

HORCM_DEV
#dev_group      dev_name      port#      TargetID      LU#      MU#
# /dev/rdsk/c23t0d0  SER =      61456 LDEV = 192 [ FIBRE FCTBL = 4 ]
ORA             ORA_000      CL2-J      0             0             0
# /dev/rdsk/c23t0d1  SER =      61456 LDEV = 193 [ FIBRE FCTBL = 4 ]
ORA             ORA_001      CL2-J      0             1             0
# /dev/rdsk/c23t0d2  SER =      61456 LDEV = 194 [ FIBRE FCTBL = 4 ]
ORA             ORA_002      CL2-J      0             2             0
# /dev/rdsk/c23t0d3  SER =      61456 LDEV = 195 [ FIBRE FCTBL = 4 ]
ORA             ORA_003      CL2-J      0             3             0
# /dev/rdsk/c23t0d4  SER =      61456 LDEV = 256 [ FIBRE FCTBL = 4 ]
ORA             ORA_004      CL2-J      0             4             0
# /dev/rdsk/c23t0d5  SER =      61456 LDEV = 257 [ FIBRE FCTBL = 4 ]
ORA             ORA_005      CL2-J      0             5             0
# /dev/rdsk/c23t0d6  SER =      61456 LDEV = 258 [ FIBRE FCTBL = 4 ]
ORA             ORA_006      CL2-J      0             6             0
# ERROR [CMDDEV] /dev/rdsk/c23t0d7  SER =      61456 LDEV = 259 [ OPEN-3-CM ]

HORCM_INST
#dev_group      ip_address    service
ORA             localhost     52323
```

paircreate

Create a pair relationship

Syntax paircreate { **-nomsg** | **-g** *group* | **-h** | **-d** *pair_vol* | **-d[g]** *raw_device* [*MU#*] | **-d[g]** *seq# LDEV* [*MU#*] | **-q** | **-vr** | **-vl** | **-f** *fence* [CTGID] | **-m** *mode* | **-c** *size* | **-nocopy** | **-split** | **-z** | **-zx** | **-jp** *ID* | **-js** *ID* }

Arguments **-c** *size* Specifies the number of tracks that are concurrently copied. The number can range from 1 to 15. The default value of 3 is used if this argument is not specified.

-d *pair_vol* Specifies a paired volume name written in the configuration definition file. The command executes only for the specified paired volume.

-d[g] *raw_device* [*MU#*]

(*HP-UX, Linux, Solaris, MPE/iX, AIX, and Windows NT/2000/2003 only*) Searches the RM configuration file (local instance) for a volume that matches the specified raw device. If a volume is found, the command is executed on the paired volume (**-d**) or group (**-dg**).

This option is effective without specification of the **-g** *group* option.

If the specified *raw_device* is listed in multiple device groups, this applies to the first one encountered.

-d[g] *seq# LDEV#* [*MU#*]

Searches the RM instance configuration file (local instance) for a volume that matches the specified sequence # and LDEV. If a volume is found, the command is executed on the paired logical volume (**-d**) or group (**-dg**).

This option is effective without specification of the **-g** *group* option.

If the specified LDEV is listed in multiple device groups, this applies to the first one encountered.

seq # LDEV # can be specified in a hexadecimal (by addition of 0x) or decimal.

-f fence [CTGID] (*CA only*) Specifies a data-consistency level.

Valid values:

async CTGID (*CA-Async and CA-Journal only*)

data

status

never

CTGID (CT group ID) is assigned automatically, but the **async** option terminates with **EX_ENOCTG** when beyond the maximum number of CT groups.

Maximum number:

XP256	16	(0-15)
XP512	64	(0-63)
XP1024	128	(0-127)
XP12000	256	(0-255)

The *CTGID* option forces creation of paired volumes for a given CTGID group.

-g group Specifies the group to be paired; the group name is specified in the **HORCM_DEV** section of the RM instance configuration file.

The command executes for the entire group unless the **-d pair_vol** argument is specified.

-h Displays Help/Usage, and version information.

-m mode The following modes may be specified:

noread (*BC only*) Specifies that the S-VOL is unreadable while the paired volumes are in the PAIR state. This mode is useful for hiding S-VOLs. By

default, the S-VOL is readable even when in the PAIR state.

cyl (*XP1024/XP12000 CA only*) Specifies that a bitmap table is managed with each cylinder on CA volumes.

trk (*XP1024/XP12000 CA only*) Specifies that a bitmap table is managed with each track on CA volumes.

If **cyl** or **trk** are not specified, the default bitmap table identified below is used.

RAID	CA Volumes	Default Bitmap table
XP1024	OPEN-3/9	Track
	OPEN-E/L/M	Cylinder
Others	Don't care	Cylinder

If there is not enough shared memory to maintain track level information, error EX_CMDRJE is returned.

dif (*BC only*) Use at paircreate to cause the S-VOL bitmap table (used to create a differential backup) to designate all tracks changed since paircreate.

inc (*BC only*) Use at paircreate to cause the S-VOL bitmap table (used for incremental backup) to designate all tracks changed since the last re-synchronization

grp [CTGID] (*XP128/XP1024/XP12000 BC only*) Use at paircreate to group specified pairs into a consistency group, allowing a consistent split of multiple devices at exactly the same point in time. This applies when doing a split using the pairsplit -g <group> command (except **-S** or **-E** option).

A CTGID (CT Group ID) is assigned automatically if you do not specify the CTGID option in the command. If CTGID is not specified and the maximum number of CT groups already exist, an EX_ENOCTG error will

occur. Therefore, the CTGID option can forcibly assign a volume group to an existing CTGID.

- nocopy** (CA only) Creates paired volumes without copying data. The data consistency of SMPL volumes is assured by the user.
- nomsg** Suppresses messages that are displayed when this command is executed. If used, this argument must be specified at the beginning of a command argument.
- q** Terminates interactive mode and exits this command.
- split** (BC only) Splits the paired volume after completing the pairing process.

-split works differently based on the microcode version:

XP256 microcode 52-46-xx or over
XP512 microcode 01-10-00/xx or over
XP1024 and XP12000:

This option will return immediately with the PVOL_PSUS and SVOL_COPY state changes. The SVOL state will be changed to SVOL_SSUS after all data is copied.

XP256 microcode 52-46-yy or under
XP512 microcode 01-10-00/xx or under:

After the command is executed, the volume status will be PVOL_COPY and SVOL_COPY. The PVOL and SVOL states will be changed to PVOL_PSUS and SVOL_SSUS after all data is copied.

- vl or -vr** Required. Specifies the direction of the P-VOL to S-VOL relationship. Specifies which set of volumes, **r** (remote) or **l** (local), is the primary (P-VOL) set. Local disks are determined by how the **HORCMINST** environment variable is set.

-vl specifies the volumes defined by the local RM instance as the primary volumes.

-vr specifies the volumes defined by the remote RM instance as the primary volumes while the local RM instance controls the secondary volume.

-z Makes this command enter interactive mode.

-zx *(Not for use with MPE/iX or OpenVMS)* Prevents using RM in interactive mode.

-jp ID *(CA-Journal only)* Specify a journal group ID for a P-VOL

-js ID *(CA-Journal only)* Specify a journal group ID for an S-VOL

Description The **paircreate** command establishes a primary to secondary pair relationship between volumes. This command generates a new paired volume from SMPL volumes. The default action pairs a logical group of volumes as defined in the RM instance configuration file.

HP-UX

Caution *Before issuing this command, ensure that the secondary volume is not mounted on any HP-UX system. If the secondary volume is mounted during the **paircreate** command, change the pair status to SMPL, unmount the secondary volume, and reissue the **paircreate** command.*

MPE/iX

Caution *Before issuing this command, ensure that the secondary volume is not mounted on any MPE/iX system. If it is, VSCLOSE that volume set and de-configure the LDEVs using IOCONFIG, the online device configuration utility program.*

Returned Values This command sets either of the following returned values in **exit()**, which allows you to check the execution results.

The command returns **0** upon normal termination.

A nonzero return indicates abnormal termination. For the error cause and details, see the execution logs.

(CA only) If the target volume is under maintenance, this command can't report copy rejection if an error occurs.

Examples Establish a BC pairing between the volumes in group **vg01**. The volumes in the local instance are used as the P-VOLs.

```
paircreate -g vg01 -vl
```

Create a BC volume pair that corresponds to disk device **/dev/rdisk/c5t1d0** as the S-VOL (using the remote instances volume as the P-VOL):

```
paircreate -d /dev/rdisk/c5t1d0 -vr
```

If the volume is part of multi-volume group, only the volume specified by the **-d** argument is set up as a pair.

Create a BC group pair out of the group that contains the seq# 35611 and LDEV 35. Use the volumes defined by the local instance as the P-VOLs:

```
paircreate -d 35611 35 -vl
```

In this example, all volumes that are part of the group that contains this LDEV are put into the PAIR state. Because MU# was not specified, it defaulted to 0.

Error Codes The table lists specific error codes for the **paircreate** command.

Category	Error Code	Error Message	Value
Volume status unrecoverable	EX_ENQVOL	Unmatched volume status within the group	236
	EX_INCSTG	Inconsistent status in group	229
	EX_INVVOL	Invalid volume status	222
	EX_INVSTP	Invalid pair status	228
	EX_ENQSIZ	Unmatched volume size for paring	212
Resource unrecoverable	EX_ENOCTG	Not enough CT groups in the RAID	217
	EX_ENXCTG	No CT groups left for OPEN Vol use	215

paircurchk

Check S-VOL data consistency

CA only

Syntax `paircurchk { -g group | -d pair_vol | -d[g] raw_device [MU#] | -d[g] seq# LDEV# [MU#] | -h | -nomsg | -q | -z | -zx }`

Arguments `-d pair_vol` Specifies a paired volume name written in the configuration definition file. The command executes only for the specified paired volume.

`-d[g] raw_device [MU#]`
(HP-UX, Linux, Solaris, MPE/iX, AIX, and Windows NT/2000/2003 only) Searches the RM configuration file (local instance) for a volume that matches the specified raw device. If a volume is found, the command is executed on the paired volume (`-d`) or group (`-dg`).

This option is effective without specification of the `-g group` option.

If the specified `raw_device` is listed in multiple device groups, this applies to the first one encountered.

`-d[g] seq# LDEV# [MU#]`
Searches the RM instance configuration file (local instance) for a volume that matches the specified sequence # and LDEV. If a volume is found, the command is executed on the paired logical volume (`-d`) or group (`-dg`).

This option is effective without specification of the `-g group` option.

If the specified LDEV is listed in multiple device groups, this applies to the first one encountered.

`seq # LDEV #` can be specified in hexadecimal (by the addition of `0x`) or decimal.

-g <i>group</i>	Specifies a group name in the RM instance configuration file.
	The command executes for the entire group unless the -d <i>pair_vol</i> argument is specified.
-h	Displays Help/Usage and version information.
-nomsg	Suppresses the messages displayed when this command is executed. It is used to execute this command from a user program.
	If used this argument must be specified at the beginning of command arguments.
-q	Terminates interactive mode and exits this command.
-z	Makes this command enter interactive mode.
-zx	<i>(Not for use with MPE/iX or OpenVMS)</i> Prevents using RM in interactive mode.

Description The **paircurchk** command displays pairing status in order to allow the operator to verify the completion of pair generation or pair resynchronization. This command is also used to confirm the paired volume connection path (physical link of paired volume to the host).

The granularity of the reported data is based on the volume or group.

Returned Values This command sets either of the following returned values in **exit()**, which allows you to check the execution results.

Normal termination

0. (OK. Data is consistent.)

Abnormal termination:

Other than 0. (For the error cause and details, refer to the execution logs.)

Output Fields Group The group name (*dev_group*) described in the configuration definition file.

Pair vol	The paired volume name (<i>dev_name</i>) within a group described in the configuration definition file.
Port targ# lun#	The port number, target ID, and LUN described in the configuration definition file.
LDEV#	The LDEV number.
Volstat	The attribute of a volume.
Status	The status of the paired volume.
Fence	The fence level of the paired volume.
To be	The data consistency of the secondary volume.

Example

```
# paircurchk -g oradb
Group Pair vol Port targ# lun# LDEV# Volstatus Status Fence To be...
oradb oradb1 CL1-A 1 5 145 S-VOL PAIR NEVER Analyzed
oradb oradb2 CL1-A 1 6 146 S-VOL PSUS STATUS Suspected
```

Error Codes The table lists specific error codes for the **paircurchk** command.

Category	Error Code	Error Message	Value
Volume status unrecoverable	EX_VOLCUR	S-VOL currency error	225

pairedisplay

Confirm pair configuration

Syntax `pairedisplay { -g group | -d pair_vol | -d[g] raw_device [MU#]
| -d[g] seq# LDEV# [MU#] | -c [[-1]] | -f[x|c|d|m|e] | -CLI | -FCA
[MU#] | -FBC [MU#] | -h | -q | -z | -m mode }`

Arguments	[-1]	Displays the local paired volume status.
	-c	Checks the paired volume connection path (physical link from paired volume to the host) and only illegally paired volumes are displayed. If this option is not specified, the status of the specified volumes is displayed without checking their path to the host.
	-CLI	Specifies structured output for Command Line Interface parsing. The column data is aligned in each row. The delimiters between columns are either a space or “-”. If you specify the -CLI option, pairedisplay will not display the cascading mirror (MU#1-4).
	-d pair_vol	Specifies a paired volume name written in the configuration definition file. The command executes only for the specified paired volume.
	-d[g] raw_device [MU#]	(<i>HP-UX, Linux, Solaris, Windows NT/2000/2003, AIX, and MPE/iX only</i>) Searches the RM instance configuration file (local instance) for a volume that matches the specified <i>raw_device</i> . If a volume is found, the command is executed on the paired volume (-d) or group (-dg). If the volume is contained in two groups, this command executes for the first volume encountered only. If <i>MU#</i> is not specified, it defaults to 0.
	-d[g] seq# LDEV# [MU#]	

Searches the RM instance configuration file (local instance) for a volume that matches the specified sequence number (*seq#*) and LDEV. If a volume is found, a command is executed on the paired logical volume (**-d**) or group (**-dg**). If the volume is contained in two groups, this command executes for the first volume encountered only. If *MU#* is not specified, it defaults to 0.

seq# is the array serial number. *seq# LDEV#* can be specified in hexadecimal (by addition of 0x) or decimal.

-f[x|c|d|m|e]

-fx displays the LDEV number in hexadecimal.

-fc displays the copy operation rate and a completion percentage. Detects and displays the status (PFUL, PFUS) and confirms **SSWS** state as an indication of **SVOL_SSUS-takeover**. This option is also used to display the copy operation progress, the Side File percentage or the BITMAP percentage for asynchronous pair volumes.

-fd displays the relationship between the Device_File and the paired volumes, based on the group (as defined in the local instance configuration definition file). If Device_File column shows **unknown** to either the local or the remote host (instance), then it shows a volume that is not recognized on the current host, and pair operations are rejected (except the local option **-l** in protection mode).

-fm displays the Bitmap mode.

-fe displays the serial number and LDEV number of the external LUNs mapped to the LDEV and additional informations for the pair volume. This option is invalid if **-m all** or **-m cas** are specified.

Example (CA)

```
# pairdisplay -g horc0 -fdxe
```

```

Group ... LDEV#.P/S,Status,Fence,Seq#,P-LDEV# M CTG JID AP EM E-Seq# E-LDEV#
horc0 ... 41.P-VOL PAIR ASYNC ,63528 40 - 0 - 2 - - -
horc0 ... 40.S-VOL PAIR ASYNC ,----- 41 - 0 - - - - -

```

Example (BC)

```

# pairdisplay -g horc0 -fe
Group ... Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M CTG CM EM E-Seq# E-LDEV#
horc0 ... 63528 65.P-VOL COPY,63528 64 - - N - - -
horc0 ... 63528 64.S-VOL COPY,----- 65 - - N - - -

```

CTG. For CA-Async and CA-Journal, displays the CT group ID, and “Fence” is shown as ASYNC. For BC, displays the CT group ID only at the time volumes are split.

JID. The journal group ID for the P-VOL or S-VOL. If the volume is not a CA-Journal volume, “-” will be displayed.

AP. The number of active paths in to the P-VOL. If this is not known, “-” will be displayed.

CM. Copy mode. “N” is for non snapshot. “S” is for snapshot.

EM. External connection mode. “H” represents a mapped E-LUN hidden from the host. “V” represents a mapped E-LUN visible to the host. “-” represents an unmapped E-LUN.

E-Seq#: The production (serial) number of the external LUN, “-” represents an unknown number.

E-LDEV#: The LDEV# of the external LUN. “-” represents an unknown number.

-FCA [MU#]

Displays a cascading CA volume in a combination CA and BC environment. If the **-I** option is specified, this option displays a cascading CA volume on a local host (near site). If no **-I** option is specified, then this option displays a cascading CA volume on a remote host (far site).

The MU# specifies the cascading mirror descriptor for CA-Journal.

-FBC [MU#] Specifies a cascading BC volume in a combination BC and CA environment. If the **-l** option is specified, this option displays a cascading BC volume on a local host (near site). If no **-l** option is specified, then this option displays a cascading BC volume on a remote host (far site).

The **-m mode** option cannot be specified.

-g group Specifies a group name in the RM instance configuration file. Group names are defined in the **HORCM_DEV** section of the RM instance configuration file.

The command executes for the entire group unless the **-d pair_vol** argument is specified.

-h Displays Help/Usage, and version information.

-l Displays the paired volume status of the local host (which issues this command).

-m mode Displays the status of mirror descriptors for specified pair logical volumes and volume pair status. The cascading volume *mode* option can be designated as **cas** or **all**.

The **cas** option displays only MU#0 (plus used MU#s). The **all** option displays all MU#s whether used or not. The mode option displays all cascading mirrors (MU#1-4).

```
# pairdisplay -g URA -CLI -fd -m all
Group   PairVol L/R Device_File  M  Seq# LDEV# P/S Status  Seq# P-LDEV# M
MURA   MURA_001 L  c1t2d7s2    0  62500 263 S-VOL PAIR   -    262  -
-       - L      c1t2d7s2    1  62500 263 SMPL      -    -    -
-       - L      c1t2d7s2    2  62500 263 SMPL      -    -    -
URA    URA_001 L  c1t2d7s2    -  62500 263 SMPL      -    -    -
-       - L      c1t2d7s2    h1 62500 263 SMPL      -    -    -
URA    URA_001 R  c1t2d8s2    0  62500 264 SMPL      -    -    -
```

-	-	R	c1t2d8s2	1	62500	264	SMPL	-	-	-	-
-	-	R	c1t2d8s2	2	62500	264	SMPL	-	-	-	-
URA	URA_001	R	c1t2d8s2	-	62500	264	SMPL	-	-	-	-
-	-	R	c1t2d8s2	h1	62500	264	SMPL	-	-	-	-

- q** Terminates interactive mode and exits this command.
- z** Makes this command enter interactive mode.
- zx** (*Not for use with MPE/iX or OpenVMS*) Prevents using RM in interactive mode.

Description The **pairdisplay** command displays the pairing status of a volume or group of volumes. This command is also used to confirm the configuration of paired volumes.

Volumes are defined in the **HORCM_DEV** section of the RM instance configuration files.

Output Fields	Group	The group name (<i>dev_group</i>) described in the configuration definition file.
	PairVol (L/R)	The paired volume name (<i>dev_name</i>) of the group described in the configuration definition file. L is the local host. R is the remote host.
	P, T#, L#	(<i>CA only</i>) The port number, target ID, and LUN described in the configuration definition file.
	Port# ID LU-M	(<i>BC only</i>) The port number, target ID, LUN, and MU# described in the configuration definition file.
	Seq#	The disk array serial number.
	LDEV#	The LDEV number.
	P/S	The (P-VOL, S-VOL) attribute of a volume.
	Status	The status of the paired volume.

- Fence (CA only) The fence level of the paired volume.
- P-LDEV# Displays the LDEV# of a primary pair partner.
- M = "W" **P-VOL and PSUS state:** indicates that S-VOL is suspending with R/W enabled.
S-VOL and SSUS state: indicates that S-VOL has been altered since entering SSUS state.
- M = "-" **S-VOL and SSUS state:** indicates that S-VOL has NOT been altered since entering SSUS state.
P-VOL and PSUS state: indicates that S-VOL is suspending with Read only.
- M = "N" **COPY/RCPY/PAIR/PSUE state:** indicates that the volume is Read-disabled.
- M = "C" Indicates a bitmap table is managed with each cylinder.
- M = "T" Indicates a bitmap table is managed with each track.
- M = "U" Indicates a bitmap is unknown.
- % The table below shows percentages for Async-CA, Sync-CA, and BC.

State Volume	Async-CA			Sync-CA			BC			
	COPY	PAIR	OTHER	COPY	PAIR	OTHER	COPY	PAIR	PVOL_PSUS SVOL_COPY	OTHER
PVOL	CR	SF	BM	CR	BM	BM	CR	CR	BM	CR
SVOL		SF	BM		BM	BM	CR	CR	CR	CR

CR Shows the copy operation rate.

BM Shows the (identical) percentage of PVOL & SVOL BITMAPs.

SF Shows the Side File percentage for each CT group in relation to a 100% full side file in cache.

The following is an arithmetic expression using the High Water Mark (HWM) as 100% of a side file space:

$$\text{HWM (\%)} = 30 / \text{Side File space (30 to 70)} * 100$$

Examples

(BC Only)

```
# pairdisplay -g oradb
```

```
Group Pair Vol (L/R) (Port#,TID,LU-M), Seq#, LDEV#...P/S, Status, Seq#, P-LDEV# M
oradb oradb1(L) (CL1-A, 1, 1-0) 30053 18 ...P-VOL PAIR 30053 19 -
oradb oradb1(R) (CL1-D, 1, 1-0) 30053 19 ...S-VOL PAIR ---- 18 -
```

(CA only)

```
# pairdisplay -g oradb -fcx
```

```
Group Pair Vol (L/R) (P,T#,L#), Seq#, LDEV# P/S, Status, Fence, %, P-LDEV# M
oradb oradb1(L) (CL1-B, 1,0) 1234 64 P-VOL PAIR Never, 75 C8 -
oradb oradb1(R) (CL1-A, 1,0) 5678 C8 S-VOL PAIR Never, --- 64 -
```

The following shows the output when using **-CLI**. The format aligns the column data in each row, making it easier to parse. The delimiters between columns are either a space or **-**.


```
Group PairVol L/R Port# TID LU-M Seq# LDEV# P/S Status Seq# P-LDEV# M
homrcf1 deva0 L CL1-D 3 5 0 30053 271 P-VOL PAIR 30053 263 -
homrcf1 deva1 L CL1-D 3 5 1 30053 271 SMPL - - - -
homrcf1 deva2 L CL1-D 3 5 2 30053 271 SMPL - - - -
```

The following example uses **-fd**.

```
# pairdisplay -g oradb -fd
```

```
Group PairVol (L/R) Device_File M ,Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M
oradb oradev1(L) c0t3d0 0 35013 17..P-VOL COPY, 35013 18 -
oradb oradev1(R) c0t3d1 0 35013 18..S-VOL COPY, 35013 17 -
```

The following figure shows cascading volumes, using the **-m** option.



The following example uses **-m cas**. This option displays the cascaded volumes at either end of the designated CA pair that are assigned either BC bitmaps (LU0-0) or CA bitmaps (LU0).

```
# pairdisplay -g oradb -m cas
Group  PairVol (L/R) (Port#,TID,LU-M), Seq#, LDEV#.P/S, Status, Seq#, P-LDEV# M
oradb  oradev1(L) (CL1-D , 3, 0-0) 30052 266...SMPL ----, ----- ---- -
oradb  oradev1(L) (CL1-D , 3, 0) 30052 266...P-VOL COPY, 30053 268 -
oradb1 oradev11(R) (CL1-D , 3, 2-0) 30053 268...P-VOL COPY, 30053 270 -
oradb2 oradev21(R) (CL1-D , 3, 2-1) 30053 268...P-VOL PSUS, 30053 272 W
oradb  oradev1(R) (CL1-D , 3, 2) 30053 268...S-VOL COPY, ----- 266 -
```

The following examples use **-m all**. This argument displays all bitmaps, whether in use or not, that can be employed with the volumes involved in the designated CA pair.

```
# 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-D , 3, 0-0) 30052 266...SMPL ----, ----- ---- -
----- (L) (CL1-D , 3, 0-1) 30052 266...SMPL ----, ----- ---- -
----- (L) (CL1-D , 3, 0-2) 30052 266...SMPL ----, ----- ---- -
oradb  oradev1(L) (CL1-D , 3, 0) 30052 266...P-VOL PAIR, 30053 268 -
oradb1 oradev11(R) (CL1-D , 3, 2-0) 30053 268...P-VOL COPY, 30053 270 -
oradb2 oradev21(R) (CL1-D , 3, 2-1) 30053 268...P-VOL PSUS, 30053 272 W
----- (R) (CL1-D , 3, 2-2) 30053 268...SMPL ----, ----- ---- -
oradb  oradev1(R) (CL1-D , 3, 2) 30053 268...S-VOL COPY, ----- 266 -
```

A MU# of **0** (not 0-0) designates a CA volume.

```
# 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-D , 3, 0-0) 30052 266...SMPL ----, ----- ---- -
----- (L) (CL1-D , 3, 0-1) 30052 266...SMPL ----, ----- ---- -
----- (L) (CL1-D , 3, 0-2) 30052 266...SMPL ----, ----- ---- -
oradb  oradev1(L) (CL1-D , 3, 0) 30052 266...P-VOL PAIR, 30053 268 -
```

```
# pairdisplay -g URA -CLI -fd -m all
Group  PairVol L/R Device_File M Seq# LDEV# P/S Status Seq# P-LDEV# M
MURA  MURA_001 L c1t2d7s2 0 62500 263 S-VOL PAIR - 262 -
-      - L c1t2d7s2 1 62500 263 SMPL - - -
-      - L c1t2d7s2 2 62500 263 SMPL - - -
URA   URA_001 L c1t2d7s2 - 62500 263 SMPL - - -
-      - L c1t2d7s2 h1 62500 263 SMPL - - -
URA   URA_001 R c1t2d8s2 0 62500 264 SMPL - - -
```

-	-	R	c1t2d8s2	1	62500	264	SMPL	-	-	-	-
-	-	R	c1t2d8s2	2	62500	264	SMPL	-	-	-	-
URA	URA_001	R	c1t2d8s2	-	62500	264	SMPL	-	-	-	-
-	-	R	c1t2d8s2	h1	62500	264	SMPL	-	-	-	-

pairevwait

Wait for event completion

Syntax `pairevwait -h`

`pairevwait { -g group | -d pair_vol | -d[g] raw_device [MU#] | -d[g] seq# LDEV# [MU#] | -FCA [MU#] | -FBC [MU#] | -h | -s status . . . | -t timeout [interval] | -nowait | -l | -nomsg | -q | -z }`

- Arguments**
- `-d pair_vol` Specifies a paired volume name written in the configuration definition file. The command executes only for the specified paired volume.
 - `-d[g] raw_device [MU#]`
(*HP-UX, Linux, Solaris, Windows NT/2000/2003, AIX, and MPE/iX only*) Searches the RM configuration file (local instance) for a volume that matches the specified raw device. If a volume is found, the command is executed on the paired volume (`-d`) or group (`-dg`).

This option is effective without specification of the `-g group` option.

If the volume is contained in two groups, the command is executed on the first volume encountered. If *MU#* is not specified, it defaults to 0.
 - `-d[g] seq# LDEV# [MU#]`
Searches the RM instance configuration file (local instance) for a volume that matches the specified sequence # and LDEV. If a volume is found, the command is executed on the paired logical volume (`-d`) or group (`-dg`).

This option is effective without specification of the `-g group` option.

If the volume is contained in two groups, the command is executed on the first volume encountered only.

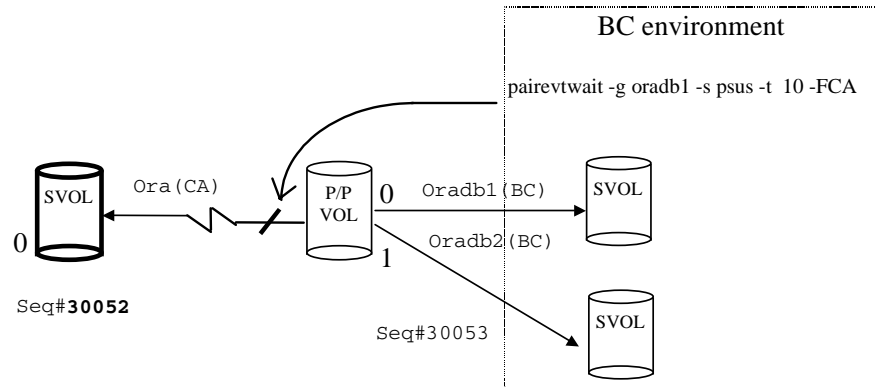
- FCA** [*MU#*] Used to forcibly specify, for event waiting, a CA P-VOL that is also a BC P-VOL. If the **-l** option is specified, then the status of a cascading CA volume on a local host (near site) is tested. If no **-l** option is specified, then this option tests the status of a cascading CA volume on a remote host (far site).
- The target CA volume must be SMPL or P-VOL.
- The *MU#* specifies the cascading mirror descriptor for CA-Journal.
- FBC** [*MU#*] Used to forcibly specify, for event waiting, a CA P-VOL that is also a BC P-VOL. If the **-l** option is specified, then this option tests the status of a cascading BC volume on a local host (near site). If no **-l** option is specified, then this option tests the status of a cascading BC volume on a remote host (far site).
- The target BC volume must be SMPL or P-VOL.
- g** *group* Specifies a group name in the RM instance configuration file.
- The command executes for the entire group unless the **-d** *pair_vol* argument is specified.
- h** Displays Help/Usage and version information.
- l** When this command cannot use a remote host because it is down, this option allows execution of this command by a local host only.
- The target volume of a local host must be SMPL or P-VOL.
- nomsg** Used to suppress messages when this command is executed from a user program.
- This option must be specified at the beginning of the command arguments.

- nowait** Causes the pairing status is reported immediately.
When this option is specified, the **-t** and **-s** options are ignored.
- q** Terminates interactive mode and exits this command.
- s status** Specifies the waiting status (SMPL, COPY [including RCPY], PAIR, PSUS, or PSUE). If two or more statuses are specified following **-s**, waiting occurs according to the logical OR of the specified statuses. This argument is not valid when the **-nowait** argument is specified.
- t timeout [interval]**
Specifies the amount of time, in one-second intervals, to wait for the specified state. If *[interval]* is not specified, the, the default value is used. This argument is not valid when the **-nowait** argument is specified. If the interval is specified as greater than 1999999, a warning message is displayed.
- z** Makes this command enter interactive mode.
- zx** (*Not for use with MPE/iX or OpenVMS*) Prevents using RM in interactive mode.

Description The **pairevtwait** command waits for completion of the **paircreate** and **pairresync** commands. It also checks the status of those commands. It waits (sleeps from the viewpoint of the process) until the paired volume status becomes identical to a specified status. When the desired status has been achieved, or the timeout period has elapsed, the command exits with the appropriate return code.

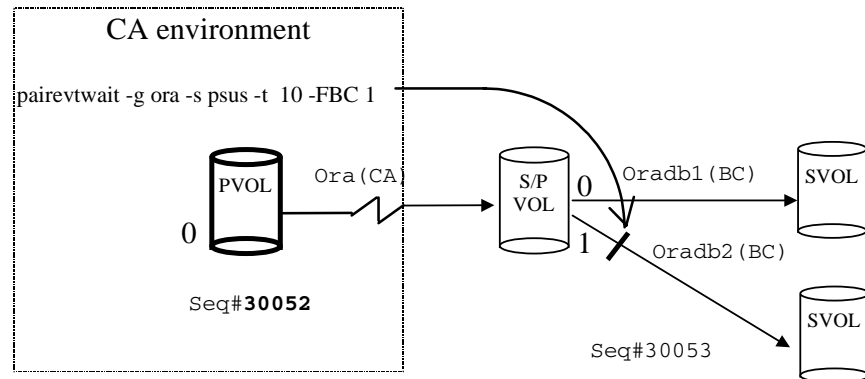
CA Operation

The figure below shows the usage of the **-FCA** option. In the example, the command waits up to 10 seconds for the designated CA pair to reach PSUS state by specifying the BC group name to which it is cascaded.



BC Operation

The figure below shows the usage of the **-FBC** option. In the example, the command tests the status of the intermediate S-VOL/P-VOL (MU#1) through a specified pair group in a CA environment.



Returned Values This command sets one of the following returned values in `exit()`, which allows you to check the execution results.

When the `-nowait` argument 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

Abnormal termination

Other than 6 to 127 (For the error cause and details, see the execution logs.)

When the `-nowait` argument is not specified:

Normal termination

- 0: The status is identical to the specified status.

Abnormal termination

Other than 0 to 127. (For the error cause and details, see the execution logs.)

Error Codes The table lists specific error codes for the **pairevwait** command.

Category	Error Code	Error Message	Value
Volume status unrecoverable	EX_ENQVOL	Unmatched volume status within the group	236
	EX_INCSTG	Inconsistent status in group	229
	EX_INVVOL	Invalid volume status	222
	EX_EVOLCE	Pair volume combination error	235
	EX_EWSUSE	Pair suspended at WAIT state	234
Timer recoverable	EX_EWSTOT	Timeout waiting for specified status	233
	EX_EWSLTO	Timeout waiting for specified status on the local host	232

pairmon

Report pair transition status

Syntax pairmon { **-D** | **-allsnd** | **-q** | **-resevt** | **-nowait** | **-s status . . .** | **-z** }

Arguments	-allsnd	Reports all pairing status transition events.
	-D	Selects the default report mode. One event is reported (and cleared) if there is pairing status transition information to be reported. If there is no information, the command waits. The report modes consists of three flags: -allsnd , -resevt , and -nowait .
	-h	Displays Help/Usage and version information.
	-nowait	When this option is specified, the pairing status is reported immediately.
	-q	Terminates interactive mode and exits this command.
	-resevt	Reports events if there is pairing status transition information and then resets all the events.
	-s status	Specifies the waiting status (SMPL, COPY [including RCPY], PAIR, PSUS, or PSUE). If two or more statuses are specified following -s , waiting occurs according to the logical OR of the specified statuses. This argument is not valid when the -nowait argument is specified.
	-z	Makes this command enter interactive mode.
	-zx	<i>(Not for use with MPE/iX or OpenVMS)</i> Prevents using RM in interactive mode.

Description The **pairmon** command is sent to the RM (daemon) to report the transition of pairing status. When an error or status transition is detected, this command outputs an error message.

Events exist on the pair state transfer queue for RM. Resetting an event correlates to the deletion of one or all events from the pair state transfer queue. If the command does not reset, the pair state transfer queue is maintained.

The table below shows the results of **pairmon** argument combinations.

-D	-nowait	-resevt	-allsnd	Actions
-D				When RM does not have an event, this option waits until an event occurs. If more events exist, it reports one event. This option clears the event that it reports.
Invalid			-allsnd	When RM does not have an event, this option waits until an event occurs. If more than one event exists, it reports all events without clearing them.
Invalid		-resevt		When RM does not have an event, this option waits until an event occurs. If multiple events exist, it reports one event and resets (clears) all events.
Invalid		-resevt	-allsnd	When RM does not have an event, this option waits until an event occurs. If multiple events exist, it reports all events and then clears them.
Invalid	-nowait			When RM does not have an event, this option reports “no event” immediately. If multiple events exist, it reports one event. This option resets (clears) the event which it reported.
Invalid	-nowait		-allsnd	When RM does not have an event, this option reports “no event” immediately. If more events exist, then it reports all events.

(continued)

-D	-nowait	-reseat	-allsnd	Actions
Invalid	-nowait	-reseat		When RM does not have an event, this option reports “no event” immediately. If multiple events exist, then it reports one event and resets all events.
Invalid	-nowait	-reseat	-allsnd	When RM does not have an event, this option reports “no event” immediately. If multiple events exist, then it reports all events and resets them.

Output Fields		
Group		The group name (dev_group) defined in the configuration definition file.
Pair vol		The paired volume name (dev_name) within the group, defined in the configuration definition file.
Port targ# lun#		The port number, TargetID, and LUN defined in the configuration definition file.
LDEV#		The LDEV number.
Oldstat		The “old” pair status when the status of a pair volume has changed.
Newstat		The “new” pair status when the status of a pair volume has changed.
code		The internal status code when the status of a pair volume has changed.

Example

```
# pairmon -allsnd -nowait
Group Pair vol Port targ# lun# LDEV# Oldstat code -> Newstat code
oradb oradb1 CL1-A 1 5 145 SMPL 0x00 -> COPY 0x01
oradb oradb2 CL1-A 1 6 146 PAIR 0x02 -> PSUS 0x04
```

pairresync

Resynchronize a pair

Syntax `pairresync { -nomsg | -g group | -d pair_vol | -d[g] raw_device [MU#] | -d[g] seq# LDEV# [MU#] | -c size | -FCA [MU#] | -FBC | -h | -l | -nomsg | -q | -restore | -swap[s|p] | -z | -zx }`

Arguments `-c size` Used to specify the number of tracks (1 to 15) copied in parallel. If omitted, the default is the value used at time of **paircreate**.

`-d pair_vol` Specifies a paired volume name written in the configuration definition file. The command executes only for the specified paired volume.

`-d[g] raw_device [MU#]`
(*HP-UX, Linux, Solaris, Windows NT/2000/2003, AIX, and MPE/iX only*) Searches the RM configuration file (local instance) for a volume that matches the specified raw device. If a volume is found, the command is executed on the paired volume (`-d`) or group (`-dg`).

This option is effective without specification of the `-g group` option.

If the specified `raw_device` is listed in multiple device groups, this applies to the first one encountered.

`-d[g] seq# LDEV# [MU#]`
Searches the RM instance configuration file (local instance) for a volume that matches the specified sequence # and LDEV. If a volume is found, the command is executed on the paired logical volume (`-d`) or group (`-dg`).

This option is effective without specification of the `-g group` option.

If the specified LDEV is listed in multiple device groups, this applies to the first one encountered.

seq # LDEV # can be specified in hexadecimal (by the addition of 0x) or decimal format.

-FCA [*MU#*] Used to resync a CA P-VOL that is also a BC P-VOL. If the **-l** option is specified, this option resynchronizes a cascading CA volume at the local host (near site). If no **-l** option is specified then this option re-synchronizes a cascading CA volume at the remote host (far site).

The target CA volume must be a P-VOL and the **-swap[s | p]** option cannot be specified.

The *MU#* specifies the cascading mirror descriptor for CA-Journal.

-FBC [*MU#*] Used to forcibly resync a BC pair whose P-VOL is a CA S-VOL. If the **-l** option is specified, this option resynchronizes a cascading BC volume on a local host (near site). If no **-l** option is specified, this option resynchronizes a cascading BC volume on a remote host (far site).

The BC volume designated by the *MU#* must be a P-VOL.

-g *group* Specifies a group name to resynchronize. The group names are defined in the **HORCM_DEV** section of the RM instance configuration file.

The command is executed for the entire group unless the **-d** *pair_vol* argument is specified.

-h Displays Help/Usage and version information.

-l Allows a local host (connected to the P-VOL) to resynchronize P-VOL to S-VOL even though the remote host is down.

- nomsg** Used to suppress messages when this command is executed from a user program.
- This option must be specified at the beginning of the command arguments
- q** Terminates interactive mode and exits this command.
- restore** *(BC only) (Optional)* Copies differential data from the secondary volume to the primary volume. (The S-VOL must not be mounted on any host while this command is executing.)

If the **-restore** option is not specified, the P-VOL is copied to the S-VOL. If the **-restore** option is used, the P-VOL must not be host mounted while the command is executing. If the target volume is currently under maintenance, this command cannot execute copy rejection in case of trouble.

If mode 80 is turned ON at the SVP, this option takes time to complete the S-VOL to P-VOL copy (**pairevtwait** will signal its completion). However, at completion, the P-VOL and S-VOL LUNs will still point to the same LDEVs (physical disks) as before.

If mode 80 is turned OFF at the SVP, this option takes virtually no time (**pairevtwait** still signals completion) because the P-VOL LUN will now be associated with the LDEVs that used to be associated with the S-VOL (and vice versa). This allows virtually immediate P-VOL access while it continues to copy to the S-VOL in the background. To avoid noticing a performance change after using this option, the P-VOL and S-VOL should use the same RAID type and the same speed disks (for example, 10k RPM).

-swap[s|p]

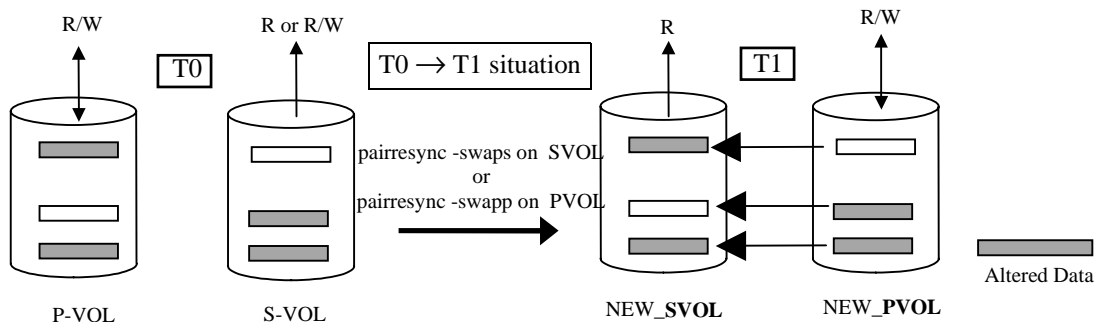
(CA only) The **-swaps** option is executed from the S-VOL when there is no host on the P-VOL side to help. A remote host must be connected to the S-VOL. Typically executed in PSUS (SSWS) state (after a **horctakeover**) to facilitate fast failback without requiring a full copy.

Unlike **-swaps**, **-swapp** requires the cooperation of hosts at both sides. It is the equivalent of **-swaps**, executed from the original P-VOL side.

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 P-VOL/S-VOL designations are swapped.

The application can continue to run at the remote failover site during this operation. At completion, the remote failover site will own the P-VOL. When desired, a very fast **horctakeover** will allow a fast failback of the application from the recovery site to the original site.

The following figure describes the **-swap[s|p]** operation. The left side of the diagram shows T0 (time zero) for both the P-VOL and S-VOL, before command execution. The right side shows T1, after command execution.



- z Makes this command enter interactive mode.
- zx *(Not for use with MPE/iX or OpenVMS)* Prevents using RM in interactive mode.

Description The **pairresync** command resumes the updating the secondary volume based on the primary volume to reestablish pairing. If no data has been written in the secondary volume, differential P-VOL data is copied. If data has been written in the secondary volume, differential data from the P-VOL is copied to the S-VOL. The changes on the SVOL are overwritten. The **-swap** option updates the PVOL based on the SVOL so that the PVOL becomes the SVOL and the SVOL becomes the PVOL. Pair resynchronization can be specified even while the primary volume is being accessed. When the **pairresync** command is issued, any write access to the secondary volume is disabled.

The **pairresync** command puts a paired volume currently in the suspend state [PSUS or SSUS] into a PAIR state.

This command can be applied to each paired logical volume or each group.

Continuous Access

This command is rejected if any P-VOL or S-VOL is undergoing maintenance, such as copy resynchronization. If the P-VOL or S-VOL is in a blocked state, an error code is returned.

UNIX

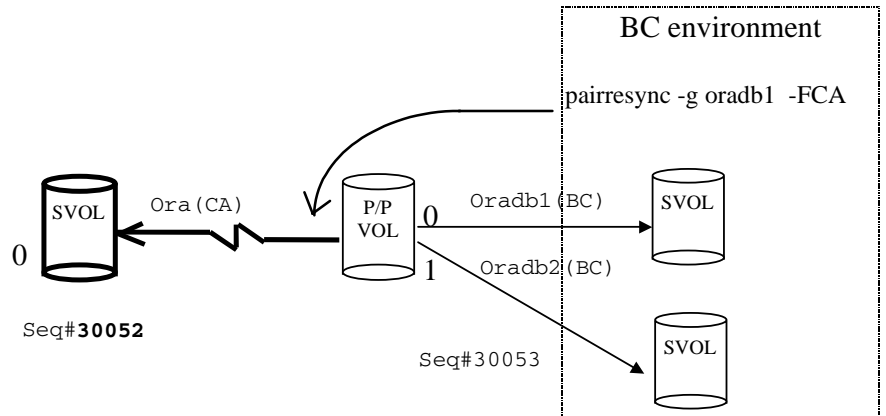
Caution *Because data in the secondary volume is renewed by **pairresync**, the secondary volume must not be in a mounted state on any UNIX system.*

MPE/iX

Caution *Before issuing this command, ensure that the secondary volume is not mounted on any MPE/iX system. If it is, VSCLOSE that volume set and de-configure the LDEVs using IOCONFIG, the online device configuration utility program.*


CA Operation

The following figure shows the usage of the **-FCA** option. In the example, the command resynchronizes a CA pair by specifying the name of a cascaded BC group.



BC Operation

The following figure shows the usage of the **-FBC** option. In the example, the command resynchronizes a BC pair (MU#1) by specifying the MU# and the CA group to which it is cascaded.



Returned Values This command sets either of the following returned values in **exit()**, which allows you to check the execution results.

The command returns **0** upon normal termination.

A nonzero return indicates abnormal termination. For the error cause and details, see the execution logs.

Output Fields	Group	The group name (<i>dev_group</i>) described in the configuration definition file.
	PairVol (L/R)	The paired volume name (<i>dev_name</i>) of the group described in the configuration definition file. L is the local host. R is the remote host.
	P, T#, L#	(<i>CA only</i>) The port number, target ID, and LUN described in the configuration definition file.
	Port# ID LU-M	(<i>BC only</i>) The port number, target ID, LUN, and MU# described in the configuration definition file.
	Seq#	The disk array serial number.
	LDEV#	The LDEV number.
	P/S	The (P-VOL, S-VOL) attribute of a volume.
	Status	The status of the paired volume.
	Fence	(<i>CA only</i>) The fence level of the paired volume.
	Copy%	The copy operation rate (identical for P-VOL and S-VOL).
	P-LDEV#	Displays the LDEV# of a primary pair partner.
	M=W	(<i>Valid for PSUS state only</i>) In the P-VOL case, this designates “suspended” with S-VOL R/W enabled.

In the S-VOL case, this designates that the S-VOL can accept writes.

M=N

(Valid for COPY/RCPY/PAIR/PSUE state) A listed volume means that reading is disabled.

Example This example shows a **pairresync** on group VG01. **pairedisplay** shows two volumes in the COPY state. The `copy%` value indicates how much of the P-VOL is in sync with the S-VOL.

```
# pairresync -g VG01
# pairedisplay -g VG01 -fc -l
Group  PairVol (L/R) (Port#,TID,LU-M),Seq#,LDEV#.P/S,Status,Copy%,P-LDEV# M
VG01   d1(L)      (CL2-P , 0, 0-0)35641 58..P-VOL COPY, 89 61 -
VG01   d2(L)      (CL2-P , 0, 1-0)35641 59..P-VOL COPY, 96 62 -
```

Error Codes The table lists specific error codes for the **pairresync** command.

Category	Error Code	Error Message	Value
Volume status unrecoverable	EX_ENQVOL	Unmatched volume status within the group	236
	EX_INCSTG	Inconsistent status in group	229
	EX_INVVOL	Invalid volume status	222
	EX_INVSTP	Invalid pair status	228

pairsplit

Split a pair

Syntax `pairsplit { | -c size / -nomsg | -g group | -d pair_vol | -d[g] raw_device [MU#] | -d[g] seq# LDEV# [MU#] | -E | -FBC [MU#] | -FCA [MU#] | -I | -r[w] | -P | -R | -S }`

Arguments `-c size` (*BC only*) Copies differential data retained in the primary volume into the secondary volume, then enables reading and writing from and to the secondary volume (after completion of the copying).

For *size*, specify a track size for copying in a range of 1 to 15. If no track size is specified, the value used for **paircreate** is used.

`-d pair_vol` Specifies a paired volume name written in the configuration definition file. The command executes only for the specified paired volume.

`-d[g] raw_device [MU#]`

(*HP-UX, Linux, Solaris, Windows NT/2000/2003, AIX, and MPE/iX only*) Searches the RM configuration file (local instance) for a volume that matches the specified raw device. If a volume is found, the command is executed on the paired volume (`-d`) or group (`-dg`).

This option is effective without specification of the `-g group` option.

If the specified *raw_device* is listed in multiple device groups, this applies to the first one encountered.

-d[g] *seq# LDEV# [MU#]*

Searches the RM instance configuration file (local instance) for a volume that matches the specified sequence # and LDEV. If a volume is found, the command is executed on the paired logical volume (**-d**) or group (**-dg**).

This option is effective without specification of the **-g** *group* option.

If the specified LDEV is listed in multiple device groups, this applies to the first one encountered.

seq # LDEV # can be specified in hexadecimal (by the addition of 0x) or decimal format.

-E

(*BC only*) Issued to forcibly suspend a paired volume (for example, when a failure occurs). It is not normally used.

-FCA [*MU#*]

Used to forcibly specify a cascading CA volume in a combination CA and BC environment. If the **-I** option is specified, this option splits a cascading CA volume on a local host (near site). If no **-I** option is specified, then this option splits a cascading CA volume on a remote host (far site).

The target CA volume must be a P-VOL, and the **-r** option cannot be specified.

The *MU#* specifies the cascading mirror descriptor for CA-Journal.

-FBC [*MU#*]

Used to forcibly specify a cascading BC volume in a combination BC and CA environment. If the **-I** option is specified, this option splits a cascading BC volume on a local host (near site). If no **-I** option is specified, then this option splits a cascading BC volume on a remote host (far site).

- The target BC volume must be a P-VOL, and the **-E** option cannot be specified.
- g** *group* Specifies which group to split. The group names are defined in the **HORCM_DEV** section of the RM instance configuration file.
- The command executes for the entire group unless the **-d** *pair_vol* argument is specified.
- h** Displays Help/Usage and version information.
- l** When the remote host is down and cannot be used, this option enables a **pairsplit** from a local host.
- (*CA only*) Unless the **-R** option is specified, the target volume of a local host must be a P-VOL.
- nomsg** Suppresses messages. If used, this argument must be the first argument specified.
- P** (*CA only*) Used to bring the primary volume forcibly into write disabled mode. It is issued by the secondary host to suppress data updating by the host possessing the primary volume.
- R** (*CA only*) Used to bring the secondary volume forcibly into SMPL mode. It is issued by the secondary host if the host possessing the primary volume goes down because of a failure or the like.
- r[w]** (*CA only*) Used to specify a mode of access to the secondary volume after paired volumes are split.
- The **-r** option allows read-only access of the secondary volume, **-r** is a default option.
- The **-rw** option enables reading and writing from and to the secondary volume.
- S** (*Optional*) Used to bring the primary and secondary volumes into SMPL mode in which pairing is not maintained. Data consistency is only maintained if

devices are in a suspend status (PSUS). If devices are in a pair status (PAIR), data on the secondary volume will not be consistent and not usable.

Description The **pairsplit** command is used to change the status of a paired volume. This command puts the pair into either PSUS or SMPL state.

For status change from PAIR to PSUS or PSUS to SMPL: Before these state changes are made, all changes made to the P-VOL, up to the point when the command was issued, are written to the S-VOL. If possible the host system must flush any of the host resident buffer cache before executing this command.

For status change from PAIR to SMPL: Changes made on the P-VOL, that are not yet copied to S-VOL will be lost and data consistency on S-VOL will not be enforced. First, change the status from PAIR to PSUS and then to SMPL to ensure consistency on S-VOL in order to use data on S-VOL.

After a pair is put into the PSUS state, changes made to the P-VOL are no longer copied to the S-VOL. However, the changes made to both the S-VOL and the P-VOL are noted and, when the volumes are resynchronized, the changed tracks or cylinders (CA) are resynchronized with the P-VOL. See “pairresync” .

When a pair is put into SMPL state, the pair relationship between the volumes is broken. Changes made to either volume are not recorded. To get the volumes back into a pair relationship, the **paircreate** command must be used.

This command stops updating the secondary volume while maintaining pairing status. When this command is issued, read or read/write access to the secondary volume is enabled and the volume is put into a SSUS state.

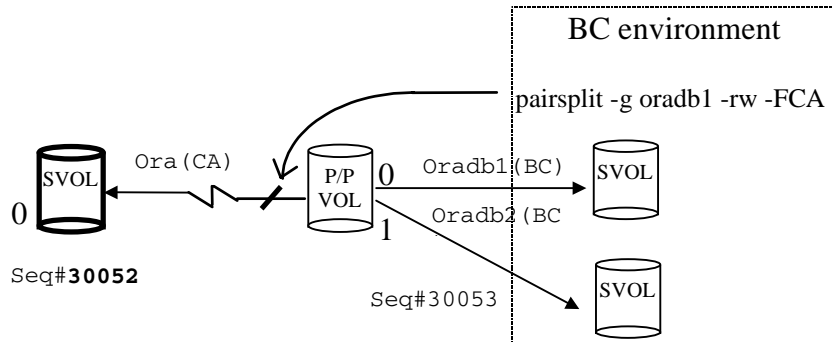
This command can be applied to each paired logical volume or each group. Only one pair splitting argument (**-r**, **-rw**, **-S**, **-R**, or **-P**) can be specified. If several arguments are specified, only the last argument is valid.

MPE/iX

Before you execute this command, the non-written data that remains in the buffer of the host must be given a flush for synchronization. For MPE/iX systems this is VSCOSE of the volume set.

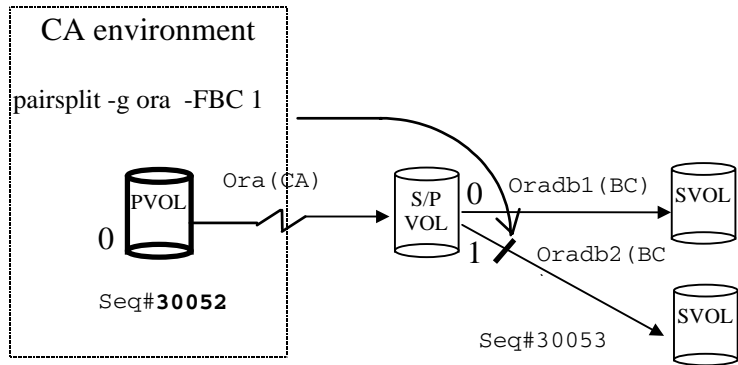
CA Operation

The following figure shows the usage of the **-FCA** option. In the example, the command splits (to PSUS) the CA pair by specifying the name of the BC group to which it is cascaded.



BC Operation

The following figure shows the usage of the **-FBC** option. In the example, the command splits (to PSUS) a BC pair (MU#1) by specifying the MU# and the name of the CA group to which it is cascaded.



Returned Values This command sets either of the following returned values in `exit()`, which allows you to check the execution results.

The command returns **0** upon normal termination.

A nonzero return indicates abnormal termination. For the error cause and details, see the execution logs.

Error Codes The table below lists specific error codes for the `pairsplit` command.

Category	Error Code	Error Message	Value
Volume status unrecoverable	EX_ENQVOL	Unmatched volume status within the group	236
	EX_INCSTG	Inconsistent status in group	229
	EX_INVVOL	Invalid volume status	222
	EX_EVOLCE	Pair volume combination error	235
	EX_INVSTP	Invalid pair status	228
	EX_EWSUSE	Pair suspended at WAIT state	234

pairsyncwait

Synchronization waiting command

Syntax `pairsyncwait { -h | -q | -z | -g group | -d pair_vol | -d[g] raw_device [MU#] | -d[g] seq# LDEV# [MU#] | -m marker | -t timeout | -nowait | -nomsg | -fq }`

Arguments	-h	Displays Help/Usage, and version information.
	-q	Terminates interactive mode and exits this command.
	-z	Makes this command enter interactive mode.
	-zx	<i>(Not for use with MPE/iX and OpenVMS)</i> Prevents you from using RM in interactive mode. This option terminates interactive mode upon RM shutdown.
	-g <i>group</i>	Used to specify a group name that is defined in the configuration definition file. The command is executed for the specified group unless the -d <i>pair_vol</i> option is specified.
	-d <i>pair_vol</i>	Used to specify a logical (named) volume that is 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#*]

(*HP-UX, Linux, Solaris, Windows NT/2000/2003, AIX, and MPE/iX only*) Searches the RM configuration file (local instance) for a volume that matches the specified raw device. If a volume is found, the command is executed on the paired volume (**-d**) or group (**-dg**).

This option is effective without specification of the **-g** *group* option.

If the specified *raw_device* is listed in multiple device groups, this will apply to the first one encountered.

-d[g] <*seq#*> <*LDEV#*> [*MU#*]

Searches the RM instance configuration file (local instance) for a volume that matches the specified sequence # and LDEV. If a volume is found, the command is executed on the paired logical volume (**-d**) or group (**-dg**).

This option is effective without specification of the **-g** *group* option.

If the specified LDEV is listed in multiple device groups, this will apply to the first one encountered.


seq # LDEV # can be specified in hexadecimal (by the addition of *0x*) or decimal format.

-m marker

Used to specify the Q-marker, the async-CA sequence # of the main control unit (MCU) P-VOL. If RM gets the Q-marker from the **-nowait** option, then it can confirm the completion of asynchronous transfer to that point, by using **pairsyncwait** with that Q-marker.

If a Q-marker is not specified, RM uses the latest sequence # at the time **pairsyncwait** is executed. It is also possible to wait for completion from the S-VOL side.

Q-marker format:



-fq

Displays the number of remaining Q-markers in the CT group in the sidefile.

```
# pairsyncwait -g oradb -nowait -fq
UnitID  CTGID      Q-Marker  Status    Q-Num  QM-Cnt
      0      3    01003408ef  NOWAIT      2      120

# pairsyncwait -g oradb -nowait -m 01003408e0 -fq
UnitID  CTGID      Q-Marker  Status    Q-Num  QM-Cnt
      0      3    01003408e0  NOWAIT      2      105

# pairsyncwait -g oradb -t 50 -fq
UnitID  CTGID      Q-Marker  Status    Q-Num  QM-Cnt
      0      3    01003408ef  TIMEOUT     2      5
```

If you specify **-nowait -fq** QM-Cnt shows the number of remaining Q-markers in the CT group.

If you specify **-nowait -m marker -fq** QM-Cnt shows the number of remaining Q-markers from the specified marker in the CT group.

If you do not specify **-nowait** and the display status is “TIMEOUT” QM-Cnt shows the number of remaining Q-markers at timeout.

If the status for the Q-market is invalid (“BROKEN” or “CHANGED”) QM-Cnt will show as “-”.

To determine the remaining data in the CT group:

Remaining data in CT group = Side File capacity * Side File percentage / 100

The side file percentage is the rate shown under the “%” column by the **pairedisplay** command.

The side file capacity is the capacity within 30-70% of the cache setting as the side file.

To determine the average data per Q-marker in the CT group:

Data per Q-Marker = Remaining data in CT group / QM-Cnt

-t timeout Used to specify the *timeout* value to wait for the completion of the remote control unit (RCU) disk fast write (DFW) cache area. It is expressed in units of multiples of 100ms. The MCU gets the latest sequence # from the RCU at regular intervals.

-nowait Used to get the latest sequence # of the MCU P-VOL and CTGID without waiting.

When this option is specified, the latest sequence # of MCU P-VOL is reported immediately, and the **-t timeout** options are ignored.

-nomsg Suppresses messages to be displayed. It is used to execute a command from a user program.

If used, this argument must be specified at the beginning of a command argument.

Description The **pairsyncwait** command is used to confirm that a mandatory write (and all writes before it) has been stored in the DFW (write) cache area of the RCU.

The command gets the latest P-VOL async-CA sequence # of the main control unit (MCU) side file and the sequence # of the most recently received write at the RCU DFW (with the correct CTGID, *group* or *raw_device*) and compares them at regular intervals.

If the RCU sequence # exceeds the value of the designated MCU sequence # within the time specified, this command reports a **0** return code (meaning P-VOL/S-VOL synchronization to the desired point is achieved).

The **-nowait** option shows the latest sequence # (Q-marker) of the designated MCU P-VOL and CTGID. The Q-marker is displayed in 10 hexadecimal characters.

Returned Values This command returns one of the following values in **exit ()**, which allows you to check the execution results.

When the **-nowait option is specified:**

Normal termination

0. The status is **NOWAIT**

Abnormal termination

Other than 0 to 127. (For the error cause and details, see the execution logs.)

When the `-nowait` option is not specified:

Normal termination:

- 0: The status is DONE (Completion of synchronization)
- 1: The status is TIMEOUT (Timeout expired before the desired Q-marker reached the remote array.)
- 2: The status is BROKEN (Q-Marker process is refused and terminated)
- 3: The status is CHANGED (Q-Marker is invalid (old) due to a prior (successfully reported) resynchronization)

Abnormal termination:

Other than 0 to 127 (For the error cause and details, refer to the execution logs)

Tip Specified *group* volume must be P-VOL with status PAIR. Other cases return an error (EX_INVVOL). It is possible to issue `pairsysncwait` from the S-VOL side, but `-m marker` is required.

Output Fields

UnitID	The Unit ID in the case of multiple DKC connections.
CTGID	The CT group ID when the LDEV has been specified as a CA-Async P-VOL or S-VOL.
Q-Marker	The sequence number of MCU P-VOL at the time the command is received.
Status	The status after execution of the command.
Q-Num	Number of processes in the queue waiting for synchronization within the CTGID of the unit.
Q-Cnt	The number of remaining I/Os in the sidefile. CA-Async sends a token called “dummy record set” at regular intervals. Therefore QM-Cnt always shows “2” or “3,” even if the host is doing no writing.

Examples When the `-nowait` option is specified:

```
# pairsyncwait -g oradb -nowait
UnitID  CTGID      Q-Marker  Status      Q-Num
0        3          01003408ef NOWAIT      2
```

When the `-nowait` option is not specified:

```
# pairsyncwait -g oradb -t 100
UnitID  CTGID      Q-Marker  Status      Q-Num
0        3          01003408ef DONE        2
```

```
# pairsyncwait -g oradb -t 1
UnitID  CTGID      Q-Marker  Status      Q-Num
0        3          01003408ef TIMEOUT    3
```

```
# pairsyncwait -g oradb -t 100 -m 01003408ef
UnitID  CTGID      Q-Marker  Status      Q-Num
0        3          01003408ef DONE        0
```

```
# pairsyncwait -g oradb -t 100
UnitID  CTGID      Q-Marker  Status      Q-Num
0        3          01003408ef BROKEN     0
```

```
# pairsyncwait -g oradb -t 100 -m 01003408ef
UnitID  CTGID      Q-Marker  Status      Q-Num
0        3          01003408ef CHANGED   0
```

Error Codes The table below lists specific error codes for the `pairsyncwait` command.

Category	Error Code	Error Message	Value
Volume status unrecoverable	EX_INVVOL	Invalid volume status	222

pairvolchk

Check volume attribute

Syntax pairvolchk { **-h** | **q** | **z** | **-g** *group* | **-d** *pair_vol* | **-d[g]** *raw_device* [*MU#*] | **-FCA** [*MU#*] | **-FBC** [*MU#*] | **-d[g]** *seq# LDEV#* [*MU#*] | **-c** | **-s[s]** | **-nomsg** }

Arguments **-c** Checks the conformability of the paired volumes of the local and remote hosts and reports the volume attribute of the remote host.

If it is not specified, the volume attribute of the local host is not reported.

-d *pair_vol* Specifies a paired volume name written in the configuration definition file. The command executes only for the specified paired volume.

-d[g] *raw_device* [*MU#*]

(*HP-UX, Linux, Solaris, Windows NT/2000/2003, AIX and MPE/iX only*) Searches the RM configuration file (local instance) for a volume that matches the specified raw device. If a volume is found, the command is executed on the paired volume (**-d**) or group (**-dg**).

This option is effective without specification of the **-g** *group* option.

If the specified *raw_device* is listed in multiple device groups, this applies to the first group encountered.

-d[g] *seq# LDEV#* [*MU#*]

This option searches the RM instance configuration file (local instance) for a volume that matches the specified sequence number (*seq#*) and LDEV. If a volume is found, the command is executed on the paired logical volume (**-d**) or group (**-dg**).

This option is effective without specification of the **-g** *group* option.

If the specified LDEV is listed in multiple device groups, this applies to the first group encountered.

seq # LDEV # can be specified in hexadecimal (by the addition of 0x) or decimal.

-FBC [*MU#*] Forcibly specifies a BC pair using the name of a CA group to which it is cascaded. If the **-c** option is not specified, this option acquires the attributes of a cascading BC volume on a local host (near site). If the **-c** option is specified, this option acquires the attributes of a cascading BC volume at the remote host (far site).

-FCA [*MU#*] Forcibly specifies a CA volume by way of its cascaded BC volume name. If the **-c** option is not specified, this option acquires the attributes of a cascading CA volume at the local host (near site). If the **-c** option is specified, this option acquires the attributes of a cascading CA volume at the remote host (far site).

The *MU#* specifies the cascading mirror descriptor for CA-Journal.

-g *group* Specifies a group name in the RM instance configuration file. The command executes for the entire group unless the **-d** *pair_vol* argument is specified.

-h Displays Help/Usage and version information.

-nomsg Suppresses messages to be displayed when this command is executed. It is used to execute a command from a user program.

If used, this argument, must be specified at the beginning of a command argument.

-q Terminates interactive mode and exits this command.

- s[s]** See the status table on [page 187](#). Used to acquire the fine granularity volume state (for example, PVOL_PSUS) of a volume.
- If it is not specified, the generic volume state (for example, P-VOL) is reported.
- z** Makes this command enter interactive mode. Interactive mode is terminated upon RAID Manager shut-down.
- zx** *(Not for use with MPE/iX or OpenVMS)* Prevents using RM in interactive mode.

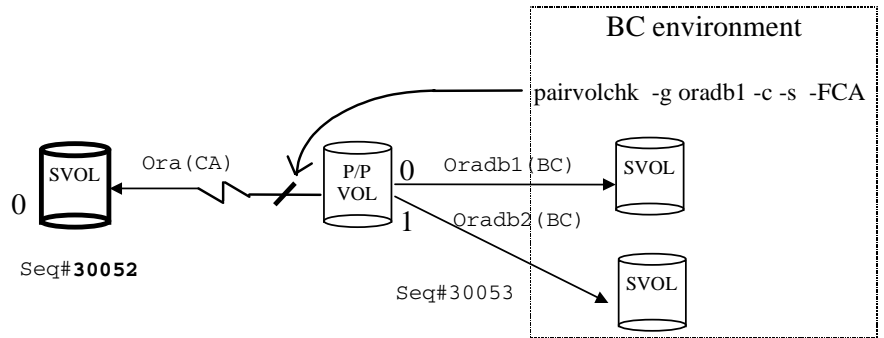
Description The **pairvolchk** command reports the attributes of a volume from the perspective of the local or remote host. This command can be applied to each paired logical volume or each group.

This is the most important command used by high availability (HA) failover software to determine when a failover or failback is appropriate.

The table under the heading “[HA control script state transitions](#)” on [page 312](#) lists state transitions resulting from the execution of **pairvolchk** in HA control scripts.

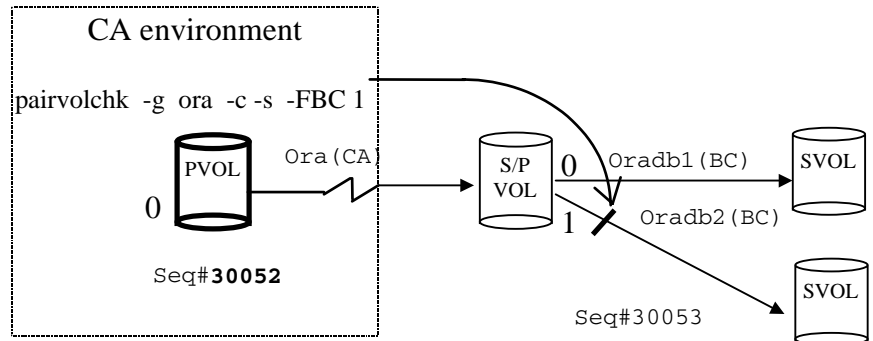
CA Operation

The following figure shows the usage of the **-FCA** option. In the example, the command acquires the fine granularity CA status by specifying the name of a BC group to which it is cascaded.



BC Operation

The following figure shows the usage of the **-FBC** option. In the example, the command acquires the fine granularity status (PVOL_PSUS) of the BC S-VOL/P-VOL (MU#1) by specifying the name of the CA group to which it is cascaded.



Returned Values When the `-s[s]` argument 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:

Other than 0 to 127. (For the error cause and details, see the execution logs.)

236: EX_ENQVOL

237: EX_CMDIOE

235: EX_EVOLCE (Only when the `-c` option is specified)

242: EX_ENORMT (Only when the `-c` option is specified)

216: EX_EXTCTG

214: EX_ENQCTG

The table below shows the error messages associated with the above error codes.

Error Code	Error Message	Return Value
EX_ENORMT	No remote host alive to accept commands or Remote RAID Manager might be blocked (sleeping) while performing I/O.	242
EX_CMDIOE	Control command I/O error	237
	Control command I/O error, or rejected	
EX_ENQVOL	Unmatched volume status within the group	236
EX_EVOLCE	Pair Volume combination error	235
EX_INCSTG	Inconsistent status in group	229
EX_VOLCUR	S-VOL currency error	225
EX_VOLCUE	Local Volume currency error	224
EX_VOLCRE	Local and Remote Volume currency error	223
EX_EXTCTG	Extended CT group between RAIDs	216
EX_ENQCTG	Unmatched CTGID within the group	214

When the `-s[s]` argument 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
- 26: The status is PVOL_PDUB (CA and LUSE volume only)
- 29: The status is PVOL_INCSTG (Inconsistent status in group)
Not returned
- 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
- 36: The status is SVOL_PDUB (CA and LUSE volume only)
- 39: The status is SVOL_INCSTG (Inconsistent status in group)
Not returned

PFUL and PFUS states:

- 42: The status is PVOL_COPY
- 43: The status is PVOL_PAIR
- 44: The status is PVOL_PSUS
- 45: The status is PVOL_PSUE
- 46: The status is PVOL_PDUB (CA and LUSE volumes only)
- 47: The status is PVOL_PFUL
- 48: The status is PVOL_PFUS
- 52: The status is SVOL_COPY or SVOL_RCPY
- 53: The status is SVOL_PAIR
- 54: The status is SVOL_PSUS
- 55: The status is SVOL_PSUE
- 56: The status is SVOL_PDUB (CA and LUSE volumes only)
- 57: The status is SVOL_PFUL
- 58: The status is SVOL_PFUS

Other than 0 to 127 (For the error cause and details, see the execution logs):

- 236:EX_ENQVOL
- 237:EX_CMDIOE
- 235:EX_EVOLCE . . . When the **-c** argument is specified
- 242:EX_ENORMT. . . When the **-c** argument is specified
- 16:EX_EXTCTG
- 214:EX_ENQCTG

When a volume group contains volumes in different states, one state will take precedence and will be reported for the group as shown in the following table.

Option	COPY	PSUE	PDUB	PFUS	PSUS	PFUL	PAIR	Group Status
-s	1	x	x	x	x	x	x	COPY*
	0	1	x	x	x	x	x	PSUE
	0	0	1	x	x	x	x	PDUB
	0	0	0	1	x	x	x	PFUS
	0	0	0	0	1	x	x	PSUS
	0	0	0	0	0	1	x	PFUL
	0	0	0	0	0	0	1	PAIR
-ss	1	x	x	x	x	x	x	COPY*
	0	1	x	x	x	x	x	PSUE
	0	0	1	x	x	x	x	PDUB
	0	0	0	x	x	1	x	PFUL
	0	0	0	x	x	0	1	PAIR
	0	0	0	1	0	0	0	PFUS
	0	0	0	0	1	0	0	PSUS

Explanation of Terms

1 Status is TRUE.

0 Status is FALSE.

x Status is TRUE or FALSE (don't care).

COPY* Status is either COPY or RCPY.

PFUL Since the PFUL state refers to the High Water Mark of the Side File in PAIR state, the PFUL state is displayed as PAIR by all commands except **pairvolchk** and the **-fc** option of the **pairdisplay** command.

PFUS Since the PFUS state is referring to a Suspend state with the Side File Full, the PFUS state is displayed as PSUS by all commands except **pairvolchk** and the **-fc** option of the **pairedisplay** command.

SVOL_PSUS Displayed as SSUS by the **pairedisplay** command.

Error Codes The table lists specific error codes for the **pairvolchk** command.

Category	Error Code	Error Message	Value
Volume status unrecoverable	EX_ENQVOL	Unmatched volume status within the group	236
	EX_EVOLCE	Pair volume combination error	235

Examples

CA Async:

```
# pairvolchk -g oradb
pairvolchk:Volstat is P-VOL.[status=PAIR fence=ASYNC CTGID=2 MINAP=2
```

CA Sync:

```
# pairvolchk -g oradb
pairvolchk : Volstat is P-VOL.[status = PAIR fence = DATA MINAP = 2 ]
```

“MINAP” shows the minimum active paths on specified group on the P-VOL. If the array firmware does not support tracking the number of active paths, then "MINAP" will not be displayed as below.

BC:

```
# pairvolchk -g oradb
pairvolchk : Volstat is P-VOL.[status = PAIR ]
```

BC with CT Group:

```
# pairvolchk -g oradb
pairvolchk : Volstat is P-VOL.[status = PAIR CTGID = 1 ]
```

raidar

Report LDEV activity

Syntax `raidar -h`

`raidar { -h | -p port targ lun [mun] | -pd raw_device | -q
| -s [interval] [count] | -z | -zx }`

Arguments	-h	Displays Help/Usage and version information.
	-p	Specifies a device location of the disk array for a disk array activity. This argument can be used more than once to monitor more than one device. It is only possible to monitor 16 devices at once.
	<i>port</i>	Specifies the name of a port to be reported by selecting it from CL1-A to CL1-R (excluding CL1-I and CL1-O), or CL2-A to CL2-R (excluding CL2-I and CL2-O). For the XP1024, the expanded ports CL3-A up to CL3-R, or CL4-A up to CL4-R can also be selected. For the XP12000, the expanded ports CL3-A up to CL3-R, or CLG-A up to CLG-R can be selected. Port specification is not case sensitive (CL1-A= cl1-a= CL1-a= cl1-A).
	<i>lun</i>	Specifies a LUN of a specified SCSI/Fibre Channel target.
	<i>targ</i>	Specifies a SCSI/Fibre Channel target ID of a specified port.
	<i>mun</i>	(<i>BC only</i>) Specifies the duplicated mirroring descriptor (MU#) for the identical LU under BC in a range of 0 to 2.
	-pd <i>raw_device</i>	(<i>HP-UX, Linux, Solaris, Windows NT/2000/2003, AIX, and MPE/iX only</i>) Allows the designation of an LDEV via the specified <i>raw_device</i> file.

- q** Terminates interactive mode and exits this command.
- s [interval] [count] or -sm [interval] [count]**
Designates the monitoring time interval.
 - s** option interprets the time interval in seconds.
 - sm** option interprets the time interval in minutes.
- interval* Specify from 1 to 60. The default value is 3.
- count* Specifies a number of repetitions. When omitted, the command repeats until cancelled via CNTL-C.
- z** Makes the command enter interactive mode. The interactive mode can be used to enter the **-p** options. This command begins an activity based on the **-s** options. This command returns to non-interactive mode upon receiving CNTL-C. Interactive mode is terminated through the **-q** option or RM shut-down.
- zx** *(Not for use with MPE/iX or OpenVMS)* Prevents using RM in interactive mode.

Description

The **raidar** command reports the I/O activity of a port, target or LUN over a specified time interval. It will report any early termination via CNTL-C. This command can be used regardless of the RM instance configuration definitions.

I/O activity of an S-VOL that is part of an active CA pair (a pair that is in the COPY or PAIR state) shows internal I/O used to maintain the pair as well as user I/O. For BC, only host I/Os are reported on the P-VOL.

For CA, the I/O activity reported for an S-VOL in either COPY or PAIR state reflects the total, not just host based, activity of the volume.

For BC, only the host based I/O activity is reported.

If the volume state changes from S-VOL (COPY or PAIR) to SMPL during the monitoring period, the activity number may be based on some internal and some host I/Os.

Output Fields	IOPS	Displays the IO (reads/writes) per second.
	HIT (%)	Displays the cache hit rate for reads.
	W (%)	Displays ratio for writes.
	IOCNT	Displays the number of reads/writes.

Example

```

# raidar -p c11-a 15 6 -p c11-b 14 5 -p c11-a 12 3 -s 3
TIME[03] PORT T L VOL STATUS IOP HIT(%) W(%) IOCNT
13:45:25 - - - - - - - - -
13:45:28 CL1-A 15 6 SMPL --- 200.0 80.0 40.0 600
          CL1-B 14 5 P-VOL PAIR 133.3 35.0 13.4 400
          CL1-A 12 3 P-VOL PSUS 200.0 35.0 40.6 600

```

sampling interval

Signifies no I/O activity

Related Information Appendix C, “[Fibre Channel addressing](#)” .

raidqry

Confirm disk array connection to host

Syntax raidqry -h

raidqry { -l | -q | -r *group* | -f | -z | -zx }

Arguments	-f	This option is used to display the floatable IP address for the hostname (<i>ip_address</i>) described in a configuration definition file.
	-h	Displays Help/Usage and version information.
	-l	Displays a configuration of the local host connected to the disk array.
	-q	Terminates interactive mode and exits this command.
	-r <i>group</i>	Displays the configuration of the remote host and the disk array connected with the designated group.
	-z	Makes this command enter interactive mode.
	-zx	(<i>Not for use with MPE/iX or OpenVMS</i>) Prevents using RM in interactive mode.

Description The **raidqry** command displays the configuration of the connected host and disk array.

Output Fields	No	Specifies the group names (by number) in the order in which they are defined in the configuration file.
	Group	When using the -r option, this item shows the group name (<i>dev_group</i>) described in the configuration definition file.

Floatable Host	When using the -f option, this item displays the first 30 characters of the host name (<i>ip_address</i>) described in the configuration definition file. The -f option interprets the host name as utilizing a floatable IP for the host.
HORCM_ver	When the -l option is specified, this shows the version of the CA of the local host. When the -r option is specified, this item shows the version of the CA on the remote host for the specified group.
Uid	Shows unit ID of the disk array connected to the local host when the -l option is specified. If the -r option is specified, the information is for the disk array connected to the remote host.
Serial#	Shows the production serial number of the disk array connected to the local host when the -l option is specified. If the -r option is specified, the information is for the disk array connected to the remote host.
Micro_ver	Shows microcode version of the disk array connected to the local host when the -l option is specified. If the -r option is specified, the information is for the disk array connected to the remote host.
Cache (MB)	Shows logical cache capacity (in MB) in the disk array. When the -l option is specified, the cache capacity is for the local disk array. When the -r option is specified, the cache capacity shown is for the remote disk array.

Display Example

```
# raidqry -l
No Group Hostname HORCM_ver Uid Serial# Micro_ver Cache(MB)
1 --- HOSTA 01-00-03/03 0 30053 52-35-02/02 256
1 --- HOSTA 01-00-03/03 1 30054 52-35-02/02 256

# raidqry -r oradb
No Group Hostname HORCM_ver Uid Serial# Micro_ver Cache(MB)
1 oradb HOSTA 01-00-03/03 0 30053 52-35-02/02 256
2 oradb HOSTB 01-00-03/03 0 30053 52-35-02/02 256
1 oradb HOSTA 01-00-03/03 1 30054 52-35-02/02 256
2 oradb HOSTB 01-00-03/03 1 30054 52-35-02/02 256

# raidqry -l -f
No Group Floatable Host HORCM_ver Uid Serial# Micro_ver Cache(MB)
1 --- FH001 01-00-03/03 0 30053 52-35-02/02 256
```

raidscan

Display port status

Syntax `raidscan`

```
raidscan {-h -CLI | -p port[hgrp] | -s Seq# | -t targ | -l lun / -pd
raw_device
| -f[xfgde] | -find | -find conf [MU#] [-g name] | -find inst
| -find [op] [MU#] | -find sync [MU#] [-g name] | -find verify | -pi
strings / -q | -z | -zx | -m [MU#] }
```

Arguments	-h	Displays Help/Usage and version information.
	-CLI	Specifies structured output for Command Line Interface parsing. The column data is aligned in each row. The delimiters between columns are either a space or “-”. If you specify the -CLI option, raidscan will not display the cascading mirror (MU1-4).
	-f [<i>x</i>]	<i>x</i> displays the LDEV number in hexadecimal.
	-f [<i>f</i>]	Specifies the volume type in the output (for example, OPEN-3/8/9/K).
		If this option is specified, the -f [<i>g</i>] and -f [<i>d</i>] options are invalid.
	-f [<i>g</i>]	<i>g</i> displays the group name in the output. This option is used to search a group in the configuration definition file (local instance) and display a <i>group_name</i> when the scanned LDEV is contained in the group.
		If this option is specified, the -f [<i>f</i>] and -f [<i>d</i>] options are invalid.
	-f [<i>d</i>]	<i>d</i> displays the Device_File that was registered to the RM Group in the output, based on the LDEV (as defined in the local instance configuration definition file).

If this option is specified, the **-f[f]** and **-f[g]** options are invalid.

-f[e]

Displays the serial number and LDEV number of the external LUNs mapped to the LDEV.

If the external LUN mapped to the LDEV on a specified port does not exist, then this option will do nothing. If this option is specified, the **-f[f][g][d]** options are not allowed.

Example

```
# raidscan -p cl1-a-0 -fe -CLI
```

PORT#	/ALPA/C	TID#	LU#	Seq#	Num	LDEV#	P/S	Status	Fence	E-Seq#	E-LDEV#	
CL1-A-0	ef	0	0	48	62468	2	256	SMPL	-	-	30053	17
CL1-A-0	ef	0	0	49	62468	2	272	SMPL	-	-	30053	23
CL1-A-0	ef	0	0	50	62468	1	288	SMPL	-	-	30053	28

E-Seq#. Displays the production (serial) number of the external LUN.

E-LDEV#. Displays the LDEV# of the external LUN.

-find

(HP-UX, Linux, Solaris, Windows NT/2000/2003, and MPE/iX only) Used to display the Port, targetID, LUN (in RAID Manager notation) that was mapped to a LDEV using a special (raw device) file provided via STDIN.

If the target and LUN are unknown, this option can be used to discover the Port, targetID, LUN associated with a host device file so that the information can be included in a **horcm.conf** file.

This option can be used with the **-fx** option to display the LDEV numbers in hexadecimal format.

-find [*op*] [*MU#*] Used to execute the specified [*op*] using a raw device file provided by STDIN. See next entries.

-find conf [*MU#*] [**-g** *name*]
Used to display the port, target ID, and LUN in the **horcm.conf** file by using a special raw device file provided via STDIN.

If the target ID and LUN are unknown for the target device file, then you will have to start RM without a description for HORCM_DEV and HORCM_INST.

This option allows you to use the **-fx** option to display the LDEV numbers in hexadecimal format.

The **-g** *name* option is used to specify the name to be used for dev_group in the **horcm.conf** file. If this option is not specified, the group applies **VG** as the default.

-find sync [*MU#*] [**-g** *name*]

Flushes the system buffer of the logical drive corresponding to a **-g** *name* (dev_group) in the RM configuration file. The dev_group name is provided via STDIN through the KEY WORD(\$Volume,\$LETALL,\$Physical).

The **-g** *name* option is used to specify the name to be used for dev_group in the **horcm.conf** file. If this option is not specified, then the system buffers associated with all groups for the local instance are flushed.

If the logical drive corresponding to a **-g** *name* is not open for any application, then the logical drive system buffer is flushed and the drive is unmounted.

If the logical drive corresponding to a **-g** *name* is open for an application, then the logical drive system buffer is only flushed.

This option allows the system buffer to be flushed before a pairsplit without unmounting the PVOL (open state).

-find verify [MU#]

Used to verify the relationship between a Group in the configuration definition file and a Device_File registered to the LDEV map tables (based on the raw device file name provided via STDIN).

This option also allows you to use the **-fx** option to display the LDEV numbers in hexadecimal format. You can also use this in conjunction with the **-fd** option.

This option will be affected by the command execution environment (HORCC_MRCF).

If a device name is different in the DEVICE_FILE and Device_File fields, then an LDEV is being referenced by multiple device files. See the Examples section for an example of such a case.

-find inst

This option runs automatically at **/etc/horcm_startup** time. It is used to logically connect and register a device file name to all pertinent mirror descriptors [MU#s] in the LDEV map table. It allows RM to note permitted volumes.

Normally, the user does not need to run this command. RM gets the serial# and LDEV# from the disk array. Then, RM compares the inquiry result to the contents of the **Horcm.conf** file, and the result is displayed and stored within the RM instance. To minimize the time required, this option is terminated when the registration is finished based on the **horcm.conf** file.

This option can be used with the **-fx** option to display the LDEV numbers in hexadecimal format.

If the **-pi strings** option is also specified, then this option does not get its “strings” via STDIN. The *strings*

specified in the **-pi** option will, instead, be used as input.

-l *lun* Specifies a LUN for a specified SCSI/Fibre Channel target. Specifying a LUN without designating the target ID is not allowed.

If this option is not specified, the command applies to all LUNs.

If this option is specified, the **-t** option must also be used.

-p *port[hgrp]* Specifies the name of a port to be scanned by selecting it from CL1-A to CL1-R (excluding CL1-I and CL1-O), or CL2-A to CL2-R (excluding CL2-I and CL2-O).

For the XP1024, the expanded ports CL3-A up to CL3-R, or CL4-A up to CL4-R can also be selected.

For the XP12000, the expanded ports CL3-A up to CL3-R, or CLG-A up to CLG-R can also be selected.

Port specifications are not case sensitive (CL1-A=cl1-a= CL1-a= cl1-A).

This option must always be specified.

The [*hgrp*] option displays only the LDEVs mapped to a host group on a XP1024/XP12000 port.

-pd *raw_device* (*UNIX only*) Specifies a *raw_device* name.

(*Windows NT/2000/2003 only*) Specifies a physical device in this format:

\\.\PhysicalDriven

(*HP-UX, Linux, Solaris, Windows NT/2000/2003, and MPE/iX only*) Finds the *Seq#* and port name on the disk array and scans the port of the disk array (which

corresponds with the unit ID) and searches for the unit ID from *Seq#*.

If this option is specified, and then the **-s Seq#** option is invalid.

-pi strings

Used to explicitly specify a character string rather than receiving it from STDIN.

If this option is specified, then the **-find** option will be ignored; the *strings* specified in the **-pi** option will, instead, be used as input. The specified character *string* must be limited to 255 characters.

-q

Terminates interactive mode and exits this command.

-s seq#

Used to specify the serial number of the disk array on multiple disk array connections when you cannot specify the unit ID which is contained in the **-p port** option.

This option searches corresponding unit ID from *Seq#* and it scans the port which is specified by **-p port** option.

If this option is specified, the unit ID which is contained in **-p port** is ignored.

Example

If the unit ID#2 has been corresponding to seq#30053 in a multiple RM configuration, then you can specify the array in the following two ways:

-raidscan -p CL1-E2

(Unit ID which is contained in **-p port** is #2.)

-raidscan -p CL1-E -s 30053

-t targ

Specifies a SCSI/Fibre target ID. If this option is not specified, the command applies to all targets.

- z** Makes this command enter interactive mode.
- zx** (Not for use with MPE/iX or OpenVMS) Prevents using RM in interactive mode.
- m MU#** Displays the cascading mirror descriptor. If you specify the **-CLI** option, **raidscan** will not display the cascading mirror (MU1-4).

Description The **raidscan** command displays, for a given port, the target ID, LDEV (mapped for LUN, and the status of the LDEV), regardless of the configuration definition file.

Output Fields	Port#	The port name on the disk array.
	ALPA/C	Arbitrated loop physical address of the port on the disk array.
	TargetID# (TID#)	The SCSI/Fibre Channel target ID of specified port.
	LUN# (LU#)	The logical unit number of specified target ID in the disk array.
	Num (LDEV# . . .)	The number of LDEV and LDEV# that used LUSE volume on the disk array.
	P/S	The (P-VOL, S-VOL) attribute of a volume.
	Status	The status of the paired volume in the disk array.
	P-Seq#	Seq# (production serial number) of the pair partner array.
	P-LDEV#	LDEV# of the partner that becomes a pair in or among the disk array.
	Vol . Type	Type name of volume mapped to LUN.
	Group	The group name (<i>dev_group</i>) described in the configuration definition file.

PairVol	The paired volume name (dev_name) within the group defined in the configuration definition file.
M	The MU# defined in the configuration definition file. For CA, the MU# is shown as -. For BC, the MU# is shown as 0 , 1 , or 2 .
Device_File	The Device_File that is registered to the LDEV map tables within RM.
UID	The unit ID for multiple array configurations. If UID is displayed as -, a command device (HORCM-CMD) has not been found.
S/F	Shows whether a port is <u>S</u> CSI or <u>F</u> ibre Channel.
PORT	The port number.
TARG	The target ID (converted by the fibre conversion table).
LUN	The logical unit number (converted by the fibre conversion table).
SERIAL	The production (serial) number of the disk array.
LDEV	The LDEV number.
PRODUCT_ID	The product ID field from the STD inquiry page.

Examples `raidscan` using the `-CLI` option formats the display so that all the columns are aligned.

```
# raidscan -p CL1-C -CLI
Port# TargetID# Lun# Seq# Num LDEV# P/S Status Fence P-Seq# P-LDEV#
CL1-C 1 0 30053 1 274 SMPL - - -
CL1-C 2 2 30053 1 260 P-VOL PAIR NEVER 30053 268
CL1-C 2 3 30053 1 261 P-VOL PAIR NEVER 30053 269
```

If you specify the `-CLI` option, `raidscan` will not display the cascading mirror (MU1-4).

A **raidscan** on a Fibre Channel port displays ALPA data for the port instead of target ID number.

```
# raidscan -p CL2-P
PORT# /ALPA/C,TID#,LU#.Num(LDEV#..)..P/S, Status,LDEV#,P-Seq#,P-LDEV#
CL2-P / ef/0, 0, 0-1.0(58).....P-VOL PSUS 58, 35641 61
CL2-P / ef/0, 0, 1-1.0yp(59).....P-VOL PSUS 59, 35641 62
CL2-P / ef/0, 0, 2...0(61).....S-VOL SSUS 61, ----- 58
CL2-P / ef/0, 0, 3...0(62).....S-VOL SSUS 62, ----- 59
```

The following example uses the **-find** option.

```
# ls /dev/* | raidscan -find
DEVICE_FILE      UID  S/F PORT   TARG  LUN   SERIAL  LDEV  PRODUCT_ID
/dev/ldev101     0    S  CL1-M    0     2     31168   118  OPEN-3-CVS
/dev/ldev102     0    S  CL1-M    0     3     31168   121  OPEN-3-CVS
/dev/ldev105     -    -  CL1-M    -     -     31170   121  OPEN-3-CVS
```

The following example uses the **-find conf** option.

```
# cat /etc/horcmperm.conf | raidscan -find conf 0 -g ORA
HORCM_DEV
#dev_group      dev_name      port#      TargetID     LU#      MU#
# /dev/rdisk/c23t0d0  SER =      61456  LDEV = 192 [FIBRE FCTBL = 4 ]
ORA             ORA_000      CL2-J      0            0        0
# /dev/rdisk/c23t0d1  SER =      61456  LDEV = 193 [ FIBRE FCTBL = 4 ]
ORA             ORA_001      CL2-J      0            1        0
# /dev/rdisk/c23t0d2  SER =      61456  LDEV = 194 [ FIBRE FCTBL = 4 ]
ORA             ORA_002      CL2-J      0            2        0
# /dev/rdisk/c23t0d3  SER =      61456  LDEV = 195 [ FIBRE FCTBL = 4 ]
ORA             ORA_003      CL2-J      0            3        0
# ERROR [CMDDEV] /dev/rdisk/c23t0d7  SER =      61456  LDEV = 259 [ OPEN-3-CM ]
```

The target device is suppressed if:

- It is the command device:

```
# ERROR [CMDDEV] /dev/rdisk/c23t0d7  SER =      61456  LDEV = 259 [ OPEN-3-CM ]
```

- It shares an LDEV among multiple device files and an LDEV is already displayed by another target device:

```
# ERROR [LDEV LINK] /dev/rdisk/c24t0d3  SER =      61456  LDEV = 195 [FIBRE FCTBL = 4]
```

- It does not have a valid MU#:

```
# ERROR [INVALID MUN (2 < 1)] /dev/rdisk/c24t0d3 SER = 61456 LDEV = 195 [ OPEN-3 ]
```

- It mixes different RAID types:

```
# ERROR [MIXING RAID TYPE] /dev/rdisk/c24t0d3 SER = 61456 LDEV = 195 [ OPEN-3 ]
```

The following example flushes the system buffer associated with the ORB group through \$Volume. This example uses the echo \$Volume | **raidscan -find sync -g ORB** or **raidscan -pi \$Volume -find sync -g ORB** options.

```
[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}
```

The following example flushes the system buffer associated with all of groups for the local instance. This example uses the echo \$Volume | **raidscan -find sync** or **raidscan -pi \$Volume -find sync** options.

```
[SYNC] : ORA ORA_000[-] -> \Vol144\Dsk0 :
Volume{56e4954a-28d5-4824-a408-3ff9a6521e5d}
[SYNC] : ORA ORA_000[-] -> \Vol145\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}
```

The following example uses the **-find inst** option.

```
# ioscan -fun | grep rdsk | raidscan -find inst
DEVICE_FILE      Group      PairVol    PORT      TARG  LUN M   SERIAL  LDEV
/dev/rdsk/c0t3d0  oradb     oradev1    CL1-D     3     0 -   35013   17
/dev/rdsk/c0t3d0  oradb     oradev1    CL1-D     3     0 0   35013   17
/dev/rdsk/c0t3d0  oradb1    oradev2    CL1-D     3     0 1   35013   17
```

The following example uses the **-find verify** option.

```
# ioscan -fun | grep rdsk | raidscan -find verify
DEVICE_FILE      Group      PairVol    PORT      TARG      LUN M      SERIAL    LDEV
/dev/rdsk/c0t3d0  oradb     oradev1    CL1-D     3         0 0      35013    17
/dev/rdsk/c0t3d1  oradb     oradev2    CL1-D     3         1 0      35013    18
/dev/rdsk/c0t3d2  -         -          -         -         - 0      35013    19
```

The following example uses the `-find verify` and `-fd` options.

```
# ioscan -fun | grep rdsk | raidscan -find verify 1 -fd
DEVICE_FILE      Group      PairVol    Device_File  M      SERIAL    LDEV
/dev/rdsk/c0t3d0  oradb     oradev1    c0t3d0      1      35013    17
/dev/rdsk/c0t3d1  oradb     oradev2    Unknown     1      35013    18
/dev/rdsk/c0t3d2  -         -          -           1      35013    19
```

SCSI Port Specification

```
# raidscan -p cl1-r
Port#, TargetID#, Lun# Num(LDEV#...) P/S, Status, Fence, LDEV#, P-Seq# P-LDEV#
CL1-R, 15, 7 5(100,101...) P-VOL PAIR NEVER 100, 5678 200
CL1-R, 15, 6 5(200,201...) SMPL ---- ---- ---- ---- ----
```

```
# raidscan -p cl1-r -f
Port#, TargetID#, Lun# Num(LDEV#...) P/S, Status, Fence, LDEV#, Vol. Type
CL1-R, 15, 7 5(100,101...) P-VOL PAIR NEVER 100, OPEN-3
CL1-R, 15, 6 5(200,201...) SMPL ---- ---- ---- OPEN-3
```

```
# raidscan -pd /dev/rdsk/c0t15/d7 -fg
Port#, TargetID#, Lun# Num(LDEV#...) P/S, Status, Fence, LDEV#, Group
CL1-R, 15, 7 5(100,101...) P-VOL PAIR NEVER 100, oradb
CL1-R, 15, 6 5(200,201...) SMPL ---- ---- ---- oradb1
```

The specified device is LDEV# **0100**.

Fibre Channel Port Specification

The ALPA/C, TID# field displays the value that was converted using the Fibre Channel Address Conversion tables in “Fibre Channel address conversions”.

```
# raidscan -p cl1-r
PORT#/ALPA/C,TID#,LU#..Num(LDEV#...) P/S, Status,Fence, LDEV#, P-Seq# P-LDEV#
CL1-R/ ce/15, 15, 7..5(100,101...) P-VOL PAIR NEVER 100, 5678 200
CL1-R/ ce/15, 15, 6..5(200,201...) SMPL ---- ---- ---- ----
```

```
# raidscan -p cl1-r -f
PORT#/ALPA/C,TID#,LU#..Num(LDEV#...) P/S, Status,Fence, LDEV#, Vol.Type
CL1-R/ ce/15, 15, 7..5(100,101...) P-VOL PAIR NEVER 100, OPEN-3
CL1-R/ ce/15, 15, 6..5(200,201...) SMPL ---- ---- ---- OPEN-3
```

SCSI Specification with BC

When using BC, **raidscan** will display the MU# (and the status of MU# 0-2) under the LUN column (for example, 7-0 for LUN 7, MU 0).

```
# raidscan -p cl1-r
Port#, TargetID#, Lun# Num(LDEV#...) P/S, Status, LDEV#, P-Seq# P-LDEV#
CL1-R, 15, 7-0 5(100,101...) P-VOL PAIR 100, 5678 300
CL1-R, 15, 7-1 5(100,101...) P-VOL PAIR 100, 5678 301
CL1-R, 15, 7-2 5(100,101...) P-VOL PAIR 100, 5678 302
CL1-R, 15, 6-0 5(200,201...) SMPL ---- ---- ---- ----
CL1-R, 15, 6-1 5(200,201...) SMPL ---- ---- ---- ----
CL1-R, 15, 6-2 5(200,201...) SMPL ---- ---- ---- ----
CL1-R, 15, 7-0 5(400,101...) S-VOL PAIR 400, 5678 100
CL1-R, 15, 7-1 5(400,101...) SMPL ---- ---- ---- ----
CL1-R, 15, 7-2 5(400,101...) SMPL ---- ---- ---- ----
```

Windows NT does not support the LDM volume. The user must specify \$LETALL instead of \$Volume as follows.

```
raidscan -pi $LETALL -find sync -g ORA
[SYNC] : ORA ORA_000[-] -> F:\Dsk1\p1 : F:
```

This option cannot specify the device object name as shown below:

```
D:\Vol (Dms, Dmt, Dmr) X\DskY, \Vol (Dms, Dmt, Dmr) X\DskY
```

Related Information

For STDIN file specification information, see Appendix D, “[STDIN file formats](#)”.

Command Options for Windows NT/2000/2003

RM provides the following commands specific to Windows NT/2000/2003. These commands are built into the RM commands and are executed by using the **-x** option with any general RM command. For instance, enter:

raidscan -x <command> <arg>

Any general command (not just **raidscan**) can be used; the **-x** option overrides the normal operation of the RM command.

It is not necessary to have an RM instance running to execute these command options when only the subcommand is to be executed.

If you execute one of these Windows NT/2000/2003 commands from a UNIX command line, a syntax error will be returned.

drivescan

Display disk drive and connection information

Windows NT/2000/2003 only

Syntax *RM_command* -x **drivescan** *stringx,y*

Arguments

<i>RM_command</i>	Any general RM command.
<i>string</i>	Any alphabetic character string; provided for readability.
<i>x,y</i>	Specifies a range of disk drive numbers.

Description The **drivescan** command displays the relationship between hard disk numbers on Windows NT/2000/2003 and the actual physical drives.

Output Fields

<i>harddiskn</i>	The hard disk number.
<i>Port</i>	The port number on the device adapter.
<i>PhId</i>	The BUS number in the device adapter port.
<i>TId</i>	The target ID of the hard disk that connects to the device adapter port. For information about Fibre Channel address conversion, see Appendix , “Fibre Channel addressing” .
<i>LUN</i>	The logical unit number of the hard disk that connects to the device adapter port. This item shows the LDEV# of the LUN that completes a pair.

Example This example shows **drivescan** executed from the **raidscan** command, and displays the connection of the actual physical drive for disk drive number 0 to 10.

```
raidscan -x drivescan harddisk0,10  
Harddisk 0..Port[ 1] PhId[ 0] Tid[ 0] Lun[ 0] [HITACHI] [DK328H-43WS]  
Harddisk 1..Port[ 2] PhId[ 4] Tid[ 29] Lun[ 0] [HITACHI] [OPEN-3]  
Port[CL1-J] Ser#[ 30053] LDEV#[ 9(0x009)]  
HORC = P-VOL HOMRCF[MU#0 = SMPL MU#1 = SMPL MU#2 = SMPL]  
Harddisk 2..Port[ 2] PhId[ 4] Tid[ 29] Lun[ 1] [HITACHI] [OPEN-3]  
Port[CL1-J] Ser#[ 30053] LDEV#[10(0x00A)]  
HORC = S-VOL HOMRCF[MU#0 = SMPL MU#1 = SMPL MU#2 = SMPL]  
Harddisk 3..Port[ 2] PhId[ 4] Tid[ 29] Lun[ 2] [HITACHI] [OPEN-3 ]  
Port[CL1-J] Ser#[ 30053] LDEV#[11(0x00B)]  
HORC = P-VOL HOMRCF[MU#0 = SMPL MU#1 = SMPL MU#2 = SMPL]  
Harddisk 4..Port[ 2] PhId[ 4] Tid[ 29] Lun[ 3] [HITACHI] [OPEN-3 ]  
Port[CL1-J] Ser#[ 30053] LDEV#[12(0x00C)]  
HORC = S-VOL HOMRCF[MU#0 = SMPL MU#1 = SMPL MU#2 = SMPL]  
Harddisk 5..Port[ 2] PhId[ 4] Tid[ 29] Lun[ 6] [HITACHI] [OPEN-3-CM]  
Port[CL1-J] Ser#[ 30053] LDEV#[15(0x00F)]
```

Related Information Appendix , “Fibre Channel addressing” .

env

Display environment variable

Windows NT/2000/2003 only

Syntax *RM_command* **-x env**

Argument *RM_command* Any general RM command.

Description The **env** command displays an environment variable within a RAID Manager command.

Example This example displays the current value of the HORCC_MRCF environment variable.

```
raidscan -x env HORCC_MRCF  
1
```

findcmddev

Search for a command device

Windows NT/2000/2003 only

Syntax *RM_command* -x **findcmddev** *stringx,y*

Arguments

<i>RM_command</i>	Any general RM command.
<i>string</i>	Any alphabetic character string; provided for readability.
<i>x,y</i>	Specifies a range of disk drive numbers.

Restriction The **findcmddev** command is used when a command device name to be described in the configuration definition file is unknown. RM must not be running when this command is used.

Description The **findcmddev** command searches to see if a command device exists within the range of the specified disk drive numbers. When the command device exists, the command displays the command device in the format described in the RM configuration definition file.

This command searches for a command device as a physical drive, a Logical drive, and a Volume{GUID} for Windows 2000/2003.

If a command device is specified as a logical drive in addition to a Physical Drive, then a drive letter is assigned to the command device. This drive letter should be deleted from the list of those available to general users.

The Volume{GUID} must be made by creating a partition, using the disk manager without the filesystem format option, and is used to keep as the same command device even though the physical drive numbers are changed on every reboot in a SAN environment.

Example This example executes **findcmddev**, searching device numbers 0 to 20.

```
raidscan -x findcmddev hdisk0, 20
cmddev of Ser# 62496 = \\.\PhysicalDrive0
cmddev of Ser# 62496 = \\.\E:
cmddev of Ser# 62496 = \\.\Volume{b9b31c79-240a-11d5-a37f-00c00d003b1e}
```

mount

Mount and display a device

Windows NT/2000/2003 only

Syntax *RM_command* **-x mount**

Windows NT:

RM_command **-x mount** *D: hdisk# [partition#] . . .*

Windows 2000/2003:

RM_command **-x mount** *D: volume#*

RM_command **-x mount** *D: [\directory] volume#]*

Arguments	<i>RM_command</i>	Any general RM command.
	<i>D</i>	Specifies the logical drive letter.
	<i>hdisk#</i>	(Windows NT only) The disk drive (hard disk) number to be mounted.
	<i>partition#</i>	(Windows NT only) The partition number within the drive.
	<i>directory</i>	(Windows 2000/2003 only) The directory mount point on the logical drive.
	<i>volume#</i>	(Windows NT only) The volume name and number to be mounted.
		(Windows 2000/2003 only) Volume# must be specified in LDM format: '\Vol#' or '\Dms#' or '\Dmt#' or '\Dmr#'

Description The **mount** command allocates the specified logical drive letter to the specified partition on the disk drive (hard disk). If no arguments are specified, this option displays a list of mounted devices.

Restrictions The partition on the specified disk drive (hard disk) must be recognized on Windows NT/2000/2003.

RAID Manager supports the mount command specifying the device object name (such as “\Device\Harddiskvolume X”). However, Windows 2003 will change the device number for the device object name when it recovers from a failure of the PhysicalDrive. So, the mount command specifying the device object name may fail due to this change.

To overcome this, specify a Volume{GUID} as well as the device object name. If a Volume{GUID} is specified, it will be converted to a device object name during execution. You can discover the Volume{GUID}s by using **inqraid \$Vol -fv** command.

Example

```
C:\HORCM\etc>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
                CL2-D      62496    256   -   -       -       -          OPEN-3-CVS-CM
```

Issuing a mount using DefineDosDevice() allows you to force a dismount of the mounted volume by logging off Windows 2000/2003.

Example

```
C:\HORCM\etc>raidscan -x mount E: Volume{cec25efe-d3b8-11d4-aead-00c00d003b1e}
E: <+> HarddiskVolume3
```

Issuing a mount using a Directory mount prevents a forced dismount due to logging off Windows 2000/2003.

Example

```
C:\HORCM\etc>raidscan -x mount E:\ Volume{cec25efe-d3b8-11d4-aead-00c00d003b1e}
E:\ <+> HarddiskVolume3
```

Output Fields	Drive	The logical drive.
	FS_name	The type of file system on the logical drive.
	VOL_name	The volume label of the logical drive.
	Device Partition	The name of device and partition mounted to the logical drive.

Port	Port number, path ID, target ID, and LUN on the device
PathID	adapter mounted to the logical drive. For information
Targ	on Fibre Channel connection on the port, see
Lun	Appendix , “Fibre Channel addressing” .

Examples

Windows NT This Windows NT example executes **mount** from the **pairsplit** command option, mounting the F:\ drive to partition 1 on disk drive 2, and mounting the G:\ drive to partition 1 on disk drive 1. Then a list of mounted devices is displayed.

```
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			

Windows 2000/2003 This Windows 2000/2003 example shows the specification of a directory mount point on the logical drive.

```
pairsplit -x mount D:\hd1 \Vol8
D:\hd1 <+> HarddiskVolume8
pairsplit -x mount D:\hd2 \Vol9
D:\hd2 <+> HarddiskVolume9
```

This Windows 2000/2003 example executes the **mount** command from a sub-command option of **pairsplit**. It mounts the F:\ drive to the harddiskvolume2, then displays the mounted devices. When the command is executed without specifying a partition#, the drive is mounted as HarddiskVolume# for Windows 2000/2003.

```
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]			

portscan

Display devices on designated ports

Windows NT/2000/2003 only

Syntax *RM_command* **-x portscan** *stringx,y*

-x portscan *port0,n*

Arguments

<i>RM_command</i>	Any general RM command.
<i>string</i>	Any alphabetic character string; provided for readability.
<i>x,y</i>	Specifies a range of port numbers.

Description The **portscan** command displays the physical devices that are connected to the designated port.

Output Fields

<i>Port</i>	The port number on the Windows NT/2000/2003 device adapter.
<i>IID</i>	The initiator ID on the device adapter port.
<i>PhId</i>	The BUS number in the device adapter port.
<i>TId</i>	The target ID of the hard disk that connects on the device adapter port. For information about Fibre Channel address conversion, see See Appendix , “Fibre Channel addressing” .
<i>LUN</i>	The logical unit number of the hard disk connected to the device adapter port. This item shows LDEV# of the partner that becomes a pair in the disk array.

Example This example executes **portscan** from the **raidscan** command option, and displays the connection of the physical device from port number 0 to 20.

```
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 ] [OPEN-3     ] ...Claimed  
  PhId[ 0] Tid[ 5] Lun[ 1] [HITACHI ] [OPEN-3     ] ...Claimed  
  PhId[ 0] Tid[ 5] Lun[ 2] [HITACHI ] [OPEN-3     ] ...Claimed  
  PhId[ 0] Tid[ 6] Lun[ 0] [HITACHI ] [3390-3A    ] ...Claimed
```

setenv

Set environment variable

Windows NT/2000/2003 only

Syntax *RM_command* **-x setenv** *variable value*

Arguments *RM_command* Any general RM command.

variable Specifies the environment variable to be set or deleted.

value Specifies the value or character string of the environment variable to be set.

Description The **setenv** command sets an environment variable within a RAID Manager command.

Restrictions Set environment variable prior to starting RM, unless you are using interactive mode.

Changing an environment variable after an execution error of a RAID Manager command is invalid.

Example This example changes the execution environment from **HORC** to **HOMRCF** by using **raidscan** to change the **HORCC_MRCF** environment variable.

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

Related Information [usetenv](#) (page 224)

sleep

Suspend execution

Windows NT/2000/2003 only

Syntax *RM_command -x sleep time*

Arguments *RM_command* Any general RM command.

time Specifies the sleep time in seconds.

Description The **sleep** command suspends execution for a specified period of time.

sync

Write data to drives

Windows NT/2000/2003 only

Syntax *RM_command -x sync A: B: C: ...*

RM_command -x sync all

RM_command -x sync drive#...

Windows 2000/2003 only:

RM_command -x sync volume#...

RM_command -x sync D:[\directory|\directory pattern]...

Arguments *RM_command* Any general RM command.

A:B:C: [\directory|\directory pattern] ...

Data is flushed to the specified logical (and the corresponding physical) drives.

If the specified logical drive has directory mount volumes, then SYNC is executed for all of the volumes on the logical drive.

[\directory|\directory pattern] (*Windows 2000/2003 only*) Specifies the directory mount point on the logical drive.

If *directory* is specified, then SYNC is executed for the specified directory mounted volume only.

If a *directory pattern* is specified then SYNC is executed for the directory mounted volumes identified by *directory pattern*.

all Data is flushed to all logical drives (and the physical drives corresponding to the logical drives assuming that they are hard disks), excluding the logical drive used by

RM and the logical drive supporting the current Windows directory.

D Data is flushed to the specified logical (and the corresponding physical) drive.

Volume#... (*Windows 2000/2003 only*) The LDM Volumes to be flushed. Volume# must be specified in LDM format: '\Vol# or \Dms# or \Dmt# or \Dmr#'

Description The **sync** command writes unwritten data remaining on the Windows NT/2000/2003 system to the logical and physical drives.

If the logical drives designated as the objects of the sync command is not opened to any applications, then sync flushes the system buffer to a drive and performs a dismount.

If the logical drives designated as the objects of the sync command are already opened to any applications, then sync only flushes the system buffer to a drive.

The **sync** command will accept a Volume{GUID} as well as the device object name. If you specify a Volume{GUID}, then RM will convert the Volume{GUID} to a device object name on execution.

Example

```
C:\HORCM\etc>raidscan -x sync Volume{cec25efe-d3b8-11d4-aead-00c00d003b1e}
[SYNC] Volume{cec25efe-d3b8-11d4-aead-00c00d003b1e}
```

Examples The following example executes SYNC for all of the volumes on a logical drive.

```
pairsplit -x sync D:
[SYNC] D: HarddiskVolume2
[SYNC] D:\hd1 HarddiskVolume8
[SYNC] D:\hd2 HarddiskVolume9
```

The following example executes SYNC for specified directory mounted volume.

```
pairsplit -x sync D:\hd1
[SYNC] D:\hd1 HarddiskVolume8
```

The following example executes SYNC for the directory mounted volumes identified by the *directory pattern* “D:\h”.

```
pairsplit -x sync D:\h
[SYNC] D:\hd1 HarddiskVolume8
[SYNC] D:\hd2 HarddiskVolume9
```

The following example executes SYNC for all of the volumes on the logical drives with directory mount volumes.

```
pairsplit -x sync all
[SYNC] C: HarddiskVolume1
[SYNC] D:\hd1 HarddiskVolume8
[SYNC] D:\hd2 HarddiskVolume9
[SYNC] G: HarddiskVolume10
```

The following example flushes HarddiskVolumeX:

```
pairsplit -x sync \VolX
```

The following example executes **sync** from a sub-command option of **pairsplit**. After flushing remaining data to the logical drives C: and D:, Read/Write access to the secondary volume is enabled.

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

The following example executes **sync** from a sub-command option of **pairsplit**. After flushing remaining data to harddisk2 and harddisk3, Read/Write access to the secondary volume is enabled in simplex mode.

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

This following example flushes the system buffer before the pairsplit without unmounting the PVOL (open state), and provides a warning.

```
pairsplit -x sync C:
```

```
WARNING: Only flushed to [\\.\C:] drive due to be opening.
```

```
[SYNC] C: HarddiskVolume3
```

umount

Unmount a device

Windows NT/2000/2003 only

Syntax *RM_command* -x **umount** *D*:

Windows 2000/2003

RM_command -x **umount** *D*: [*directory*]

Arguments

<i>RM_command</i>	Any general RM command.
<i>D</i>	Specifies the logical drive letter to unmount
<i>directory</i>	(Windows 2000 /2003 only) The directory mount point on the logical drive.

Restriction Before issuing the **umount** command, all drive activity must be stopped, including system activity and user applications. If activity is not stopped, the unmount operation is not completed and a “device busy” error is reported.

Description The **umount** command unmounts a logical drive and deletes the drive letter. Before deleting the drive letter, the command automatically executes the **sync** command for the specified logical drive (flushes unwritten buffer data to the disk).

Output Fields

Drive	Displays the logical drive on Windows NT/2000/2003.
FS_name	Displays the file system name of the logical drive on Windows NT/2000/2003.
VOL_name	Displays the volume label name of the logical drive on Windows NT/2000/2003.
Device Partition	Displays the name of the device and partition mounted to the logical drive on Windows NT/2000/2003.
Port PathID	Displays the port number, path ID, target ID, and LUN on the device adapter mounted to the logical drive on

Targ Windows NT/2000/2003.
Lun

For information on Fibre Channel connection on the port, see Appendix , “Fibre Channel addressing” .

Examples

Windows 2000/2003 This Windows 2000/2003 example shows the specification of a directory mount point on the logical drive.

```
pairsplit -x umount D:\hd1  
D:\hd1 <-> HarddiskVolume8  
pairsplit -x umount D:\hd2  
D:\hd2 <-> HarddiskVolume9
```

This example executes **umount** from the **pairsplit** command option, after unmounting the F:\ drive and G:\ drive. Read/Write access to the secondary volume is enabled, and mounted devices are displayed.

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

usetenv

Delete environment variable

Windows NT/2000/2003 only

Syntax *RM_command* **-x usetenv** *variable*

Arguments *RM_command* Any general RM command.
variable Specifies the environment variable to be deleted.

Description The **usetenv** command deletes an environment variable within a RAID Manager command.

Restrictions Changing an environment variable after an execution error of a RAID Manager command is invalid.

Example This example changes the execution environment from **HORC** to **HOMRCF** by using **raidscan** to change the **HORCC_MRCF** environment variable.

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

Related Information [setenv](#) (page 217)

Data Integrity Check Commands

To set and verify the validation check parameters for Data Integrity Check, RM provides the following commands.

raidvchkset

Integrity checking command

Data Integrity Check only

Syntax `raidvchkset { -h | -q | -z | -g group | -d pair_vol -d[g] raw_device [MU#] | -d[g] seq# LDEV# [MU#] | -nomsg | -vt [type] | -vs bsize [SLBA ELBA] [-vg [type] [rtime]] }`

Arguments	-h	Displays Help/Usage and version information.
	-q	Terminates interactive mode and exits this command.
	-z	This option makes this command enter interactive mode.
	-zx	<i>(Not for use with MPE/iX or OpenVMS)</i> This option prevents using RM in interactive mode.
	-g <i>group</i>	Specifies a group name from the configuration definition file. The command is executed for the specified group unless the -d <i>pair_vol</i> option is specified.
	-d <i>pair_vol</i>	Specifies a paired logical volume name from the configuration definition file. The command is executed only for the specified paired logical volume.
	-d[g] <i>raw_device</i> [MU#]	Searches the RM configuration file (local instance) for a volume that matches the specified raw device. If a volume is found, the command is executed on the paired volume (-d) or group (-dg). This option is effective without specification of the -g <i>group</i> option. If the volume is contained in two groups, the command is executed on the first volume encountered. If MU# is not specified, it defaults to 0.

-d[g] *seq# LDEV# [MU#]*

Searches the RM instance configuration file (local instance) for a volume that matches the specified sequence # and LDEV. If a volume is found, the command is executed on the paired logical volume (**-d**) or group (**-dg**).

This option is effective without specifying the **-g** *group* option.

If the volume is contained in two groups, the command is executed on the first volume encountered only.

seq# LDEV# can be specified in hexadecimal by addition of **0x**, or decimal.

-nomsg

Used to suppress messages when this command is executed from a user program.

This option must be specified at the beginning of the command arguments.

-vt [*type*]

Specifies the data type of the target volumes as an Oracle database. If *type* is not specified, this option disables all checking.

Valid values for *type*:

redo8

Sets the parameter for validation checking as Oracle redo log files (including archive logs) prior to Oracle9i. This option sets *bsize* to 1 (512 bytes) for Solaris or 2 (1024 bytes) for HP-UX.

data8

Sets the parameter for validation checking as Oracle data files prior to Oracle9i.

redo9

Sets the parameter for validation checking as Oracle redo log files for Oracle9iR2 or later. This option sets *bsize* to 1 (512 bytes) for Solaris or 2 (1024 bytes) for HP-UX.

data9

Sets the parameter for validation checking as Oracle data files (including control files) for Oracle9iR2 later.

-vs *bsize* [*SLBA ELBA*]

Specifies the data block size of Oracle I/O and a region on a target volume for validation checking.

bsize is used for specifying the data block size of Oracle I/O, in units of 512 bytes. *bsize* is able to specify between 1 (512 bytes) and 128 (64 kilobytes), but the effective size for Oracle is between 1 (512 bytes) and 64 (32 kilobytes).

If the **-vs** option is also used for redo log volumes to specify *SLBA ELBA*, *bsize* must be set to **2** for HP-UX or **1** for Solaris.

SLBA ELBA specifies a region defined between *Start_LBA* and *End_LBA* on a target volume for checking, in units of 512 bytes. The effective region is from 1 to end-of-LU.

SLBA ELBA can be specified in hexadecimal by addition of **0x**, or decimal.

If this option is not specified, then a region for a target volume is set as all blocks (*SLBA*=0; *ELBA*=0).

-vg [*type*][*rtime*]

Specifies the following guard types to the target volumes for HP StorageWorks LUN Security XP Extension.

If [*type*] is not specified, then this option disables all guarding. If no guard type has been specified, then the volume will be unguarded (read and write operations from the host as well as use as an S-VOL will be allowed).

If [*type*] has been specified previously to set a guard level and the time specified in [*rtime*] has not elapsed,

the guard characteristics of the target volumes will not be changed.

If [type] has been specified previously to set a guard level and the time specified in [rtime] has elapsed, then not specifying [type] will disable all guarding for the target volumes.

NOTE: Once a volume has guard attribute set, write access for that volume cannot be restored by the customer until [rtime] has expired. If a volume has been set to a guarded state by accident, contact HP support for recovery of the volume. Valid values for *type*:

inv

Conceals the target volumes from the SCSI Inquiry command by responding with “unpopulated volume.”

Sz0

The target volumes reply with “SIZE 0” through the SCSI read capacity command.

rwd

Disables the target volumes from reading and writing.

wtd

Disables the target volumes from writing. The volumes cannot be used as an S-VOL or written by a host.

svd

Disables the target volumes so they cannot become an S-VOL. Read and Write operations from hosts are still allowed.

[rtime]

Specifies the data retention time, in days. If [rtime] is not specified, then the data retention time never expires. Disk array microcode versions 21-06-xx and 21-07-xx ignore this option and always set the retention time to never expire.

If [ptime] is not specified, then the default time defined by the microcode version will be used. The default time is “infinite” in microcode version 21-06-xx or 21-07-xx. The default time is “zero” in microcode version 21-08-xx.

Description The **raidvchkset** command sets the parameters for integrity checking to the specified volumes and can also be used to turn off all integrity checking without specifying *type* when the time is specified in [ptime] and when the integrity checking that was originally set (or later extended) has elapsed.

The unit for the protection checking is based on a group in the RAID Manager configuration file.

When enabling DIC using **raidvchkset**, if there are redundant paths to the same LUN (for example, when using HP StorageWorks Auto Path or LVM pv-links), it is not necessary to enable **raidvchkset** on each path. Enable DIC on only one path, usually the path specified in the RM **horcm.conf** configuration file.

Returned values Return values in **exit()** allow you to check execution results from a user program. Normal termination returns **0**.

Examples This example sets the volumes for the oralog group as redo log file prior to Oracle9i.

```
raidvchkset -g oralog -vt redo8
```

This example sets the volumes for the oradat group as data file, where the Oracle block size is 8 kilobytes.

```
raidvchkset -g oradat -vt data8 -vs 16
```

This example sets to the volumes for the oradat group as data file, where the Oracle block size is 16 kilobytes.

```
raidvchkset -g oradat -vt data8 -vs 32
```

This example disables all volume checking for the oralog group.

```
raidvchkset -g oralog -vt
```

This example disables all writing to volumes for the oralog group:

```
raidvchkset -g oralog -vg wtd
```

This example disables all writing and retention time for the oralog group:

```
raidvchkset -g oralog -vg wtd 365
```

This example disables guarding for the oralog group:

```
raidvchkset -g oralog -vg
```

This example disables writing for the oralog group.

```
raidvchkset -g oralog -vg wtd
```

This example disables writing and sets as retention time of 365 days.

```
raidvchkset -g oralog -vg wtd 365
```

This example releases all guarding for the oralog group.

```
raidvchkset -g oralog -vg
```

Error codes This command is rejected with EX_ERPERM by connectivity checking between RAID Manager and the disk array.

The **raidvchkset -vg** option returns the following error code as well as generic errors:

Category	Error Code	Error Message	Value
Volume Status Unrecoverable	EX_EPRORT	Mode changes denied due to retention time	208

This means that the target volume mode cannot be changed, because retention time prevents it. Confirm the retention time for the target volume by using **raidvchkscan -v gflag**.

Flags The command sets the following four flags each for the guarding types:

Type	INQ	RCAP	READ	WRITE
Inv	1	1	1	1
Sz0	0	1	1	1
Rwd	0	0	1	1
Wtd	0	0	0	1

raidvchkdsp

Integrity checking confirmation command

Data Integrity Check only

Syntax `raidvchkdsp { -h | -q | -z | -zx | -g group | -d pair_vol -d[g] raw_device [MU#] | -d[g] seq# LDEV# [MU#] | -f[xde] | -v operation }`

Arguments	-h	Displays Help/Usage and version information.
	-q	Terminates interactive mode and exits this command.
	-z	This option makes this command enter interactive mode.
	-zx	<i>(Not for use with MPE/iX or OpenVMS)</i> This option prevents using RM in interactive mode.
	-g group	Specifies a group name from the configuration definition file. 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 from the configuration definition file. The command is executed only for the specified paired logical volume.
	-d[g] raw_device [MU#]	Searches the RM configuration file (local instance) for a volume that matches the specified raw device. If a volume is found, the command is executed on the paired volume (-d) or group (-dg). This option is effective without specification of the -g group option. If the volume is contained in two groups, the command is executed on the first volume encountered. If MU# is not specified, it defaults to 0.

-d[g] seq# LDEV# [MU#]

Searches the RM instance configuration file (local instance) for a volume that matches the specified sequence # and LDEV. If a volume is found, the command is executed on the paired logical volume (**-d**) or group (**-dg**).

This option is effective without specifying the **-g group** option.

If the volume is contained in two groups, the command is executed on the first volume encountered only.

seq# LDEV# can be specified in hexadecimal by addition of **0x**, or decimal.

-f[xde]

-fx displays the LDEV/STLBA/ENLBA number in hexadecimal.

-fd displays the relationship between the Device_File and the paired volumes, based on the group (as defined in the local instance configuration definition file). If the Device_File column shows “unknown” to either the local or the remote host (instance), then the volume is not recognized on the current host, and the command is rejected in protection mode.

Example

```
raidvchkdsp -g vg01 -fd -v cflag
```

Group	PairVol	Device_File	Seq#	LDEV#	BR-W-E-E	MR-W-B	BR-W-B	SR-W-B-S
vg01	oradb1	c4t0d2	2332	2	D E B R	D D D	D E E	D E D D
vg01	oradb2	c4t0d3	2332	3	D E B R	D D D	D E E	D E D D

-fe displays the serial numbers and LDEV numbers of the external LUNs mapped to the LDEV for the target volume.

This option displays the information above by adding to last column, and then ignores the 80-column format.

Example:

```
# raidvchkdsp -g horc0 -v gflag -fe
```

Group	...	TID	LU	Seq#	LDEV#	GI-C-R-W-S	PI-C-R-W-S	R-Time	EM	E-Seq#	E-LDEV#
horc0	...	0	20	63528	65	E E E E E	E E E E E	0	-	-	-
horc0	...	0	20	63528	66	E E E E E	E E E E E	0	-	-	-

EM. Displays the external connection mode
H = mapped E-LUN as hidden from the host.
V = mapped E-LUN as visible to the host
- = unmapped to the E-LUN

E-Seq#. Displays the production (serial) number of the external LUN. If it is unknown, the number will be shown as “-”.

E-LDEV#. Displays the LDEV# of the external LUN. If it is unknown, the number will be shown as “-”.

-v operation

Specifies an *operation* that displays the each parameter for validation checking.

Valid values for *operation*:

-v cflag

Displays all flags for checking regarding data block validation for target volumes.

BR-W-E-E: Displays the flags for checking data block size.

R=Read → E=Enable and D=Disable

W=Write → E=Enable and D=Disable

E=Endian format → L=Little and B=Big

E=Not rejected when validation error → W=Write and

R=Read

MR-W-B: Displays the flags for checking block header information.

R=Read → E=Enable and D=Disable

W=Write → E=Enable and D=Disable

B=Block #0 → E=Enable and D=Disable

BR-W-B: Displays the flags for checking data block number information.

R=Read → E=Enable and D=Disable

W=Write → E=Enable and D=Disable

B=Data Block → E=Enable and D=Disable

SR-W-B-S: Displays the flags for checking data block checksum.

R=Read → E=Enable and D=Disable

W=Write → E=Enable and D=Disable

B=Block #0 → E=Enable and D=Disable

S=Checksum → E=Enable and D=Disable

-v offset

Displays the range setting for data block size of Oracle I/O and a region on a target volume for validation checking.

Bsize: Displays the data block size of Oracle I/O, in units of bytes.

STLBA: Displays the Start of LBA on a target volume for checking, in units of LBAs.

ENLBA: Displays the End of LBA on a target volume for checking, in units of LBAs. If STLBA and ENLBA are both zero, this means to be checked as all blocks.

BNM: Displays whether this validation is disabled or enabled. If BNM is 0 then this validation is disabled.

-v gflag

Display the flags for guarding for target volumes.

Example

```
raidvchkdsp -g vg01 -fd -v gflag
```

Group	PairVol	Device_File	Seq#	LDEV#	GI-C-R-W-S	PI-C-R-W-S	R-Time
vg01	oradb1	c4t0d2	2332	2	E E D D E	E E D D E	365
vg01	oradb2	c4t0d3	2332	3	E E D D E	E E D D E	-

GI-C-R-W-S displays the protection flags for the target volume. The flags are “E” for enabled and “D” for disabled.

I. Inquiry command.
C. Read Capacity command.
R. Read command.
W. Write command.
S. Ability to become an S-VOL.

PI-C-R-W-S displays the permission flags, showing whether the permission flags can be changed to enable. “E” indicates that a flag can be changed to enable. “D” indicates that it cannot.

I. “I” flag permission.
C. “C” flag permission.
R. “R” flag permission.
W. “W” flag permission.
S. “S” flag permission.

-v R-Time

The retention time for write protection, in days. A hyphen (-) indicates that the retention time is “infinite.”

-v errcnt

Displays statistical information for errors counted on the target volumes. The error count is cleared when the individual flag for integrity checking is disabled.

CFEC: Error counter for checking of block size validation.

MNEC: Error counter for checking of block header validation.

SCEC: Error counter for checking of data block checksum validation.

BNEC: Error counter for checking of block number validation.

Description The **raidvchkdsp** command displays the parameters for protection checking of the specified volumes. The unit of checking for the protection is based on the group of RM configuration file.

A nonpermitted volume is shown without LDEV# information (LDEV# information is -).

Error codes This command is rejected with EX_ERPERM by connectivity checking between RAID Manager and the disk array.

Examples # raidvchkdsp -g vg01 -fd -v cflag

Group	PairVol	Device_File	Seq#	LDEV#	BR-W-E-E	MR-W-B	BR-W-B	SR-W-B-S
vg01	oradb1	Unknown	2332	-	- - -	- - -	- - -	- - -
vg01	oradb2	c4t0d3	2332	3	D E B R	D D D	D E E	D E D D

raidvchkdsp -g vg01 -fd -v offset

Group	PairVol	Device_File	Seq#	LDEV#	Bsize	STLBA	ENLBA	BNM
vg01	oradb1	c4t0d2	2332	2	1024	1	102400	9
vg01	oradb2	c4t0d3	2332	3	1024	1	102400	9

raidvchkdsp -g vg01 -fd -v cflag

Group	PairVol	Device_File	Seq#	LDEV#	BR-W-E-E	MR-W-B	BR-W-B	SR-W-B-S
vg01	oradb1	c4t0d2	2332	2	D E B R	D D D	D E E	D E D D
vg01	oradb2	c4t0d3	2332	3	D E B R	D D D	D E E	D E D D

raidvchkdsp -g vg01 -fd -v errcnt

Group	PairVol	Device_File	Seq#	LDEV#	CfEC	MNEC	SCEC	BNEC
vg01	oradb1	c4t0d2	2332	2	0	0	0	0
vg01	oradb2	c4t0d3	2332	3	0	0	0	0

raidvchkscan

Integrity checking confirmation command

Data Integrity Check only

Syntax `raidvchkdsp { -h | -q | -z[x] | -p port [hgrp] | -pd[g] raw_device |
-s seq# | -t target | -l LUN | -fx | -v operation }`

- Arguments**
- h** Displays Help/Usage and version information.
 - q** Terminates interactive mode and exits this command.
 - z** This option makes this command enter interactive mode.
 - zx** *(Not for use with MPE/iX or OpenVMS)* This option prevents using RM in interactive mode.
 - p port** Specifies the name of a port to be scanned by selecting it from CL1-A to CL1-R (excluding CL1-I and CL1-O), or CL2-A to CL2-R (excluding CL2-I and CL2-O). For the expanded port, specify CL3-a to CL3-r or CL4-a to CL4-r. Port names are not case sensitive.

This option always must be specified if **-pd** *raw_device* option is not specified.

[**hgrp**] is specified to display only the LDEVs mapped to a host group on a port for disk array.
 - pd[g] raw_device** Specifies a *raw_device* name.

Finds the *Seq#* and port name on the disk array and scans the port of the disk array (which corresponds to the unit ID) and searches for the unit ID from *Seq#*.

This option always must be specified if the **-find** or **-p** *port* option is not specified. If this option is specified, the **-s** *Seq#* option is invalid.

-pdg specifies the LUNs displayed in host view by locating a host group for XP 128 and XP 1024 arrays.

- s seq#** Specifies the serial number of the disk array on multiple disk array connections when you cannot specify the unit ID that is contained in the **-p port** option or the **-v jnl** option.
- This option searches corresponding unit ID from *Seq#* and it scans the port that is specified by **-p port** option.
- If this option is specified, the unit ID that contained in **-p port** is invalid. If this option is specified, the unit ID contained in **-v jnl** is invalid.
- t target** Specifies a SCSI/Fibre target ID of a specified port. If this option is not specified, the command applies to all targets.
- l LUN** Specifies the LUN of a specified SCSI/Fibre target. If this option is not specified, the command applies to all LUNs.
- A LUN-only specification without designating a target ID is invalid.
- fx** Displays the LDEV/STLBA/ENLBA number in hexadecimal.
- v operation** Specifies an *operation* that displays each parameter for validation checking.
- Valid values for *operation*:
- v cflag**
Displays all flags for checking regarding data block validation for target volumes.
- BR-W-E-E: Displays the flags for checking data block size.
R=Read → E=Enable and D=Disable
W=Write → E=Enable and D=Disable
E=Endian format → L=Little and B=Big
E=Not rejected when validation error → W=Write and R=Read

MR-W-B: Displays the flags for checking block header information.

R=Read → E=Enable and D=Disable

W=Write → E=Enable and D=Disable

B=Block #0 → E=Enable and D=Disable

BR-W-B: Displays the flags for checking data block number information.

R=Read → E=Enable and D=Disable

W=Write → E=Enable and D=Disable

B=Data Block → E=Enable and D=Disable

SR-W-B-S: Displays the flags for checking data block checksum.

R=Read → E=Enable and D=Disable

W=Write → E=Enable and D=Disable

B=Block #0 → E=Enable and D=Disable

S=Checksum → E=Enable and D=Disable

-v offset

Displays the range setting for data block size of Oracle I/O and a region on a target volume for validation checking.

-v Bsize

Displays the data block size of Oracle I/O, in units of bytes.

STLBA: Displays the Start of LBA on a target volume for checking, in units of LBAs.

ENLBA: Displays the End of LBA on a target volume for checking, in units of LBAs. If STLBA and ENLBA are both zero, this means to be checked as all blocks.

BNM: Displays whether this validation is disabled or enabled. If BNM is 0 then this validation is disabled.

-v errcnt

Displays the statistical information about errors on the target volumes. Statistical information is cleared when the individual flag for integrity checking is disabled.

CFEC: Displays the error counter for checking of block size validation.

MNEC: Displays the error counter for checking of block header validation.

SCEC: Displays the error counter for checking of data block checksum validation.

BNEC: Displays the error counter for checking of block number validation.

-v gflag

Display the flags for block data validation for target volumes.

Example

```
# raidvchkscan -p CL1-A -v gflag
```

PORT#	/ALPA/C	TID#	LU#	Seq#	Num	LDEV#	GI-C-R-W-S	PI-C-R-W-S	R-Time
CL1-A	/ ef/	0	0	2332	1	0	E E D D E	E E D D E	365
CL1-A	/ ef/	0	1	2332	1	1	E E D D E	E E D D E	-
CL1-A	/ ef/	0	2	2332	1	2	E E D D E	E E E E E	0

GI-C-R-W-S displays the protection flags for the target volume. The flags are “E” for enabled and “D” for disabled.

- I. Inquiry command.
- C. Read Capacity command.
- R. Read command.
- W. Write command.
- S. Ability to become an S-VOL.

PI-C-R-W-S displays the permission flags, showing whether the permission flags can be changed to

enable. “E” indicates that the “I” flag can be changed to enable. “D” indicates that it cannot.

- I. “I” flag permission.
- C. “C” flag permission.
- R. “R” flag permission.
- W. “W” flag permission.
- S. “S” flag permission.

R-Time is the retention time for write protection, in days. A hyphen (-) indicates that the retention time is “infinite.”

-v jnl unit#

Finds the journal volume lists setting and displays information for the journal volume.

Example

```
# raidvchkscan -v jnl 0
JID MU CTG JNLS AP U(%) Q-Marker Q-CNT D-SZ(BLK) Seq# Nnm LDEV#
001 0 1 PJNN 4 21 43216fde 30 512345 62500 2 265
002 1 2 PJNF 4 95 3459fd43 52000 512345 62500 3 270
003 0 3 PJSN 4 0 - - 512345 62500 1 275
004 0 4 PJSF 4 45 1234f432 78 512345 62500 1 276
005 0 5 PJSE 0 0 - - 512345 62500 1 277
006 - - SMPL - - - - 512345 62500 1 278
007 0 6 SMPL 4 5 345678ef 66 512345 62500 1 278
```

JID displays the journal group ID.

MU displays the mirror descriptions on CA-Journal.

CTG displays the CT group ID.

JNLS displays the following status for the journal group.

SMPL indicates that the journal volume doesn't have a pair or the pair has been deleted.

P(S)JNN = “P-(S)-VOL Journal Normal Normal.”

P(S)JSN = “P-(S-)VOL Journal Suspend Normal.”

PJNF = “P-(S-)VOL Journal Normal Full.”

P(S)JSF = “P-(S-)VOL Journal Suspend Full.”

P(S)JSE = “P-(S-)VOL Journal Suspend Error,”
with link failure.

AP (active path) displays the following two conditions, according to the pair status.

In the case of pair status **PJNL** or **SJNL** (except suspend state), this field shows the number of active paths on the initiator port in CA-Journal links. A “-” indicates “unknown.”

In case of pair status **SJNL** (suspend state), this field shows the result of the suspend operation and indicates whether or not all data on PJNL (P-VOL) were passed (synchronized) to S-JNL (S-VOL) completely. If AP is 1, all data were passed. If not, all data were not passed from S-JNL (S-VOL).

U(%) displays the usage rate of the journal data.

Q-Marker displays the sequence number of the journal group ID, called the Q-marker.

In the case of pair status **PJNL**, Q-Marker shows the latest sequence number on the PJNL volume.

In the of pair status **SJNL**, Q-Marker shows the latest sequence number on the cache (DFW).

Q-CNT displays the number of remaining Q-Markers in each journal volume.

D-SZ displays the capacity of the journal data on the journal volume.

Seq# displays the serial number of the array.

Num displays the number of LDEVs configured the journal volume.

LDEV# displays the first number of LDEV configured as the journal volume.

The table below shows the meanings of JNLS status when combined with other information.

JNLS		Other info		Meaning
P-JNL	S-JNL	QCNT	AP	
SMPL		0	-	Configured as journal volume, but NOT pair
		N	-	Deleting the journal volume
PJNN	SJNN	0	-	Normal state of the journal volume without data
PJNN	-	N	-	Normal state of the journal volume with data
	SJNN	N	N	Normal state of the journal volume with data
		N	0	Normal state of the journal volume at link failure
PJSN	SJSN	0	-	Suspended journal volume via operation
		N	-	Suspending the journal volume
PJNF	-	N	-	High water mark state
PJSF	SJSF	0	-	Suspended journal volume due to full journal
		N	-	Suspending the journal volume due to full journal
PJSE	-	0	-	Suspended journal volume due to failure/link failure
		N	-	Suspending the journal volume due to failure/link failure
-	SJSE	0	N	Suspended journal volume due to failure
			0	Suspended journal volume due to link failure
		N	N	Suspending the journal volume due to failure
			0	Suspending the journal volume due to link failure

Description The **raidvchkscan** command sets the parameters for protection checking to the specified volumes. The unit of checking for the protection is based on the **raidscan** command.

Error codes This command is rejected with EX_ERPERM by connectivity checking between RAID Manager XP and the disk array.

RAID Manager XP reports the following message to the syslog file as an integrity check error when each statistical information counted an error is updated.

HORCM_103 Detected a validation check error on this volume (*dev_group, dev_name, unit#X, ldev#Y*): CfEC=*n*, MNEC=*n*, SCEC=*n*, BNEC=*n*

Cause: A validation error occurred on the database volume, or validation parameters for this volume are invalid.

Action to be taken: Confirm the following items, and use the **raidvchkdsp -v operation** command for verifying the validation parameters.

Check whether the block size (**-vs size**) is an appropriate size.

Check whether the type for checking (**-vt type**) is an appropriate type.

Check whether the data validations are disabled for LVM configuration changes.

Check whether the data validations are not used based on the file system.

Check whether the redo log and data file are separated among the volumes.

Examples # `raidvchkscan -p CL1-A -v cflag`

PORT#	/ALPA/C	TID#	LU#	Seq#	Num	LDEV#	BR-W-E-E	MR-W-B	BR-W-B	SR-W-B-S
CL1-A	/ ef/	0	0	2332	1	0	D E B R	D D D	D E E	D E D D
CL1-A	/ ef/	0	1	2332	1	1	D E B R	D D D	D E E	D E D D
CL1-A	/ ef/	0	2	2332	1	2	D E B R	D D D	D E E	D E D D
CL1-A	/ ef/	0	3	2332	1	3	D E B R	D D D	D E E	D E D D
CL1-A	/ ef/	0	4	2332	1	4	D E B R	D D D	D E E	D E D D

raidvchkscan -p CL1-A -v offset

PORT#	/ALPA/C	TID#	LU#	Seq#	Num	LDEV#	Bsize	STLBA	ENLBA	BNM
CL1-A	/ ef/	0	0	2332	1	0	1024	1	102400	9
CL1-A	/ ef/	0	1	2332	1	1	1024	1	102400	9
CL1-A	/ ef/	0	2	2332	1	2	1024	1	102400	9
CL1-A	/ ef/	0	3	2332	1	3	1024	1	102400	9
CL1-A	/ ef/	0	4	2332	1	4	1024	1	102400	9

raidvchkscan -p CL1-A -v errcnt

PORT#	/ALPA/C	TID#	LU#	Seq#	Num	LDEV#	CfEC	MNEC	SCEC	BNEC
CL1-A	/ ef/	0	0	2332	1	0	0	0	0	0
CL1-A	/ ef/	0	1	2332	1	1	0	0	0	0
CL1-A	/ ef/	0	2	2332	1	2	0	0	0	0
CL1-A	/ ef/	0	3	2332	1	3	0	0	0	0
CL1-A	/ ef/	0	4	2332	1	4	0	0	0	0

Troubleshooting RAID Manager

This chapter lists RM errors and describes the problem, typical cause, and solution for each.

Error reporting

If you have a problem with RM, first make sure that the problem is not caused by the host or the connection to the disk array.

The tables in this chapter provide detailed troubleshooting information:

- “Operational notes” on page 251
- “Error codes” on page 254
- “Command return values” on page 256
- “Command errors” on page 259

If a failure occurs in CA or BC volumes, find the failure in the paired volumes, recover the volumes, and continue operation in the original system. If a CA command terminates abnormally, see the RM activation log file, error log file, and trace file to identify the cause.

RM monitors failures in the paired volumes at regular intervals. When it detects a failure, it sends an error message to the host syslog file. When a failure is detected and reported, collect the data in the RM error log file and trace data file (in all files under `$HORCM_LOG`) to determine the cause of the error.

Operational notes

Error	Solution
Coexistence of Logical Volume Manager (LVM) mirror and CA	<p>When the LVM mirror and CA volumes are used together, the LVM mirror handles write errors by switching LVM P-VOL volumes. Thus, the fence level of mirrored P-VOLs used by the LVM must be set to data.</p> <p>One instance of LVM must not be allowed to see both the P-VOL and S-VOL of the same BC or CA pair. This will cause an LVM error in that two volumes will contain the same LVM volume group ID.</p> <p>If you wish to split and mount an S-VOL on the same host as the P-VOL, you must first use the vgchgid command to give the S-VOL a new LVM volume group ID.</p>
Command device	<p>Each BC/CA command is executed by issuing a command to the command device. The BC/CA command is read from or written from or into a specific block area of the command device. Therefore, the command device cannot be used by the user. In addition, this device must not belong to an LVM volume group.</p>
Duplicated write error	<p>(<i>CA only</i>) Check the error notification command or the syslog file to identify the failed paired volume. Issue an RM command manually to the identified failed paired volume to try to recover it.</p> <p>If the secondary volume is the failed volume, issue the pairresync command to recover it.</p> <p>If the primary volume fails, delete or suspend the pair (pairsplit command) and use the secondary volume as the primary volume, and create another pair.</p>
RM command termination error (see the command log file and RM log file for error details.)	<p>If an RM command terminates abnormally because of a remote server failure, recover the machine from the failure, and then reexecute the RM command. If a hardware error occurs, read the log files and contact the HP support center.</p>
<i>(continued)</i>	

Error	Solution
horctakeover (swap-takeover)	When executing horctakeover on a standby server manually, I/O activity on the servers (for the pertinent CA volumes) must be stopped.
Host machines that can own opposite sides of a CA pair	Host machines must be running the same operating system and the same architecture.
New RM installations	After a new host system has been constructed, a RM failure to start can occur due to an improper environmental setting or an inaccurate configuration definition file. Use the RM activation log file for RM error definitions.
Host failure	<p>If a failure occurs on host A (with application failover software installed), host B detects the failure and issues a takeover command making the secondary volumes usable. If the secondary volumes can continue processing, host B takes over processing from host A.</p> <p>While host B continues to process (swap-takeover) the command, the volumes are swapped so that the secondary volumes become the primary volumes and the primary volumes become the secondary volumes. When host A has recovered from the failure, it can take back ownership and control through another swap-takeover command.</p>
Secondary volume failure	<p>If the primary volume detects a failure in the secondary volume, pair writing is suspended. The primary volume changes the paired volume status to PSUE. (The fence level determines whether host A continues processing, that is, writing, or host B takes over from host A.) RM detects the change in status and sends a message to the syslog.</p> <p>If host A had initiated a monitoring command, a message appears on host A. When the secondary volume recovers, host A updates the S-VOL data by running the pairsplit -S, paircreate -vl, or pairresync command.</p>
Startup failure	When the P-VOL server boots up, the secondary volume can be updated. If the secondary volume is used by the LVM, the volume group of the LVM must be deactivated. The secondary volume must only be mounted to a host when the volume is in PSUS state or in SMPL mode. The secondary volume must not be mounted automatically in any host boot sequence.
<i>(continued)</i>	

Error	Solution
SCSI alternating path restrictions	If the primary and secondary volumes are on the same server, alternate pathing, for example, pvlink, cannot be used (from primary volume to secondary volume). Use of SCSI alternative pathing to a volume pair is limited to one side of a pair. The hidden S-VOL option can avoid undesirable alternate pathing.

Error codes

Error Code	Problem	Cause	Solution
HORCM_001	The RM log file cannot be opened.	The file cannot be created in the RM directory.	Create space on the root disk.
HORCM_002	The RM trace file cannot be opened.	The file cannot be created in the RM directory.	Create space on the root disk.
HORCM_003	The RM daemon could not produce enough processes to complete the request.	The RM daemon attempted to create more processes than the maximum allowable number.	Close any unnecessary programs or increase the maximum number of allowed processes.
HORCM_004	RM failed, resulting in a fatal internal error.	An unidentifiable RM error occurred.	Restart the system, and call the HP support center.
HORCM_005	RM failed to create the end point for remote communication.	RM failed to create a socket, or an error exists in the RM configuration file \$HORCM_CONF .	See the RM startup log to identify the cause of the error.
HORCM_006	RM memory allocation failed.	RM memory could not be secured.	Increase the virtual memory of the system, or close any unnecessary programs.
HORCM_007	An error exists in the RM setup file.	An error exists in the RM setup file.	See the startup log and edit the parameters.
HORCM_008	The RM configuration file parameters could not be read.	An error exists in the format or parameters of the RM configuration file \$HORCM_CONF .	See the RM startup log to identify the cause of the error.
<i>(continued)</i>			

Error Code	Problem	Cause	Solution
HORCM_009	CA/RM connection to RM failed.	System devices are improperly connected, or an error exists in the RM configuration file \$HORCM_CONF .	See the RM startup log to identify the cause of the error.
HORCM_101	CA/RM and RM communication failed.	A system I/O error occurred or an error exists in the RM configuration file \$HORCM_CONF .	See the RM startup log to identify the cause of the error.
HORCM_102	The volume is suspended.	The pairing status was suspended.	Call the HP support center.
HORCM_103	A validation check error was detected on the volume.	A validation error occurred on the database volume or validation parameters for this volume are illegal.	<p>Confirm the following items, and use the raidvchkdsp -v <op> command to verify the validation parameters.</p> <p>The Block size(-vs <size>) is appropriate.</p> <p>The type for checking (-vt <type>) is appropriate.</p> <p>Data validation is disabled for LVM configuration changes.</p> <p>Data validation is not used based on the File system.</p> <p>The redo log and data file are on separate volumes.</p>

Command return values

For error descriptions, see [“Error codes” on page 254](#).

Return Value	Command Error	Error Message
211	EX_ERPERM	RAID permission denied.
212	EX_ENQSIZ	Unmatched pairing volume size.
213	EX_ENPERM	LDEV permission denied.
214	EX_ENQCTG	Unmatched CTGID.
215	EX_ENXCTG	No such CT group (Open Systems volume)
216	EX_ENTCTG	Extended CT group across disk arrays.
217	EX_ENOCTG	Not enough CT groups in the disk array.
218	EX_ENQSER	Unmatched Serial# / disk array unit ID
219	EX_ENOUNT	Specified disk array unit does not exist.
220	EX_INVNUM	Invalid MU number used with CA or BC.
221	EX_CMDRJE	An order of the control command was rejected.
222	EX_INVVOL	Invalid volume status.
223	EX_VOLCRE	Local and/or remote volume currency error.
224	EX_VOLCUE	Local volume currency error.
225	EX_VOLCUR	S-VOL currency error.
226	EX_INVRCD	Invalid return code.
227	EX_ENLDEV	Invalid logical device defined.
228	EX_INVSTP	Invalid pair status.
229	EX_INCSTG	Inconsistent status in group.
230	EX_UNWCMD	Unknown command.
<i>(continued)</i>		

Return Value	Command Error	Error Message
231	EX_ESTMON	RM monitoring has stopped.
232	EX_EWSLTO	Local host timeout error.
233	EX_EWSTOT	Timeout error.
234	EX_EWSUSE	Pairsplit –E.
235	EX_EVOLCE	Pair volume combination error.
236	EX_ENQVOL	Group volume matching error occurred.
237	EX_CMDIOE	Command I/O error.
238	EX_UNWCOD	Unknown function code.
239	EX_ENOGRP	Specified group is not defined.
240	EX_INVCMD	Invalid disk array command.
241	EX_INVMOD	Invalid disk array command.
242	EX_ENORMT	No available remote host.
243	EX_ENAMLG	Specified file name is too long.
244	EX_ERANGE	Resulting value is too large.
245	EX_ENOMEM	Insufficient memory.
246	EX_ENODEV	Specified device does not exist.
247	EX_ENOENT	Specified group or device does not exist.
248	EX_OPTINV	Specified option is invalid.
249	EX_INVNAM	Specified name is invalid.
250	EX_ATTDBG	Cannot attached to a Debug layer.
251	EX_ATHOR	RM software error.
252	EX_UNWOPT	Unknown option.
253	EX_INVARG	Invalid argument.
254	EX_REQARG	Required argument list is not specified.
<i>(continued)</i>		

Return Value	Command Error	Error Message
255	EX_COMERR	Cannot communicate with RM.
256	EX_ENOSUP	SVOL denied due to disabling
257	EX_EPRORT	Mode changes denied due to retention time.

Command errors

Command Error	Problem	Action
EX_ATTDBG	This command failed to communicate with RM, or a log directory file could not be created.	Verify that RM is functioning properly.
EX_ATTHOR	Connection could not be made with RM.	Verify that RM has started and that the correct HORCMINST value has been defined.
EX_CMDIOE	The request to the command device either failed or was rejected.	Check to see whether the syslog file of the host reports an Illegal Request (0x05) Sense Key, if so, then verify: <ul style="list-style-type: none"> • The BC/CA functions are installed on the disk array; • The ESCON RCP and LCP ports are set properly; • The CU paths have been established; • The target volume is available.
EX_CMDRJE	The request to the command device either failed or was rejected.	Verify the following: <ul style="list-style-type: none"> • The BC/CA functions are installed on the disk array. • The ESCON RCP and LCP ports are set properly. • The CU paths have been established. • The target volume is available.
EX_COMERR	This command failed to communicate with RM.	Verify that RM is running.
EX_ENAMLG	Undefined error.	Call the HP support center.
<i>(continued)</i>		

Command Error	Problem	Action
EX_ENLDEV	A device defined in the configuration file does not have an assigned LUN, port, or target ID.	Verify that the configuration file is correct and that all devices are defined correctly.
EX_ENOCTG	Not enough CT groups. Could not register because 15 CTs (XP256), 63 CTs (XP512), 127 CTs (XP1024), or 255 CTs (XP12000) are already in use.	Decrease the number of CTs in use, or use the pairvolchk command to display the CTs in use; then use paircreate with -f async CTGID or -mgrp CTGID to specifically assign new pairs to existing CTs.
EX_ENODEV	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.
EX_ENOENT	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.
EX_ENOGRP	The designated device or group name does not exist in the configuration file, or the network address for remote communication does not exist for the specified group name.	Verify the device or group name and add it to the configuration file of the remote and local hosts.
EX_ENOMEM	Insufficient memory.	Increase the virtual memory of the system, or close any unnecessary programs.
EX_ENORMT	A timeout error occurred.	Verify that the local and remote servers are properly communicating, and increase the timeout value in the configuration file.
EX_ENOSUP	S-VOL error	Verify the microcode version by using the raidqry -l command
EX_ENOUNT	The disk array unitID that was designated as a command argument does not exist in the configuration file.	Verify the disk array unitID and add it to the HORCM_CMD section of the local host configuration file.
<i>(continued)</i>		

Command Error	Problem	Action
EX_ENPERM	A device mentioned in the configuration file does not have permission for a pair operation.	Use the pairdisplay or raidscan –find verify command to confirm that a pair operation is permitted for the device.
EX_ENQCTG	The CT group in a group does not match the CTGID number.	Confirm the CTGID by using the pairvolchk command.
EX_ENQSER	The group that was designated by paircreate (for BC) does not have the same disk array unit, or the unitID is not identical to the unitID of the same serial# of the disk array.	Confirm the serial# by using the pairdisplay or the raidqry –r command.
EX_ENQSIZ	Unmatched pairing volume size.	Use the raidscan -f command to confirm the volume size or number of the LUSE volume, and change the SIZE of the volumes to match.
EX_ENQVOL	The attributes or the fence level of the primary and secondary volumes do not match.	Confirm the attributes and fence level settings using the pairdisplay command and reset the volume attributes and fence levels.
EX_ENXCTG	An available CT group for an Open Systems volume does not exist for asynchronous CA or BC.	Confirm whether all CT groups are already used by other volumes.
EX_EPRORT	Mode changes denied due to retention time	Verify the retention time for a target volume using the raidvchkscan -v gflag command.
EX_ERANGE	The argument or the result of the argument exceeds the maximum command value.	Re-issue the command, making sure to correctly define all of the command arguments.
EX_ERPERM	RAID permission denied.	Use the inraid -CLI and raidqry -h commands to confirm the type of RAID permitted for RM.
<i>(continued)</i>		

Command Error	Problem	Action
EX_ESTMON	RM monitoring is prohibited.	Verify the poll value defined in the configuration file.
EX_EVOLCE	The chosen primary and secondary volumes cannot be paired.	Confirm the status of each volume using the pairedisplay command.
EX_EWSLTO	The command timed out because the remote host did not respond.	Verify that the remote server is functioning properly.
EX_EWSTOT	The command has timed out.	Change the timeout value and re-issue the command.
EX_ENOMEM	Insufficient memory.	Increase the virtual memory of the system, or close any unnecessary programs.
EX_EWSUSE	A paired volume has failed and become suspended.	Issue the pairresync command to try to recover the failed pair. If the pairresync command does not restore the pair, call the HP support center.
EX_EXTCTG	A CA volume is defined in the configuration file HORCM_CONF as a group extended across disk arrays.	Confirm serial number or unit ID of the volumes by using the pairedisplay command.
EX_INCSTG	The status of a volume in the group is not consistent with the pair status.	Verify the pair status using the pairedisplay command.
EX_INVARG	An option or arguments of the command is incorrect.	Reissue the command, making sure to correctly define all of the command arguments.
EX_INVCMD	Disk array error.	Call the HP support center.
EX_INVMOD	Disk array error.	Call the HP support center.
EX_INVMUN	An invalid MU number has been defined.	Confirm the MU number of the specified group using the pairedisplay command.
EX_INVNAM	An invalid name is defined in the command argument.	Reissue the command, making sure to correctly define all of the command arguments.
<i>(continued)</i>		

Command Error	Problem	Action
EX_INVRCDD	Incorrect return code.	Call the HP support center.
EX_INVSTP	The target volume is not accessible because of an invalid volume status.	Verify the volume status using the pairdisplay command.
EX_INVVOL	The target volume is not accessible because of an invalid volume status.	Verify the volume status using the pairdisplay command.
EX_OPTINV	Disk array error.	Call the HP support center.
EX_REQARG	All the necessary command arguments have not been provided.	Reissue the command, making sure to define all of the command arguments.
EX_UNWCMD	An unknown command has been defined.	Verify the command name and re-issue the command.
EX_UNWCOD	Disk array reporting error.	Call the HP support center.
EX_UNWERR	Undefined error.	Call the HP support center.
EX_UNWOPT	An unknown option has been defined.	Reissue the command, making sure to use only defined command arguments.
EX_VOLCRE	Swap-takeover volume specification error.	Verify the pair status using the pairdisplay command.
EX_VOLCUE	S-VOL specification error.	Verify the pair status using the pairdisplay command.
EX_VOLCUR	The currency of the S-VOL data cannot be verified.	Verify the pair status using the pairdisplay command.


A

Configuration file examples

This appendix presents examples of RM configuration files.

Configuration definition for cascading volumes

RAID Manager is capable of keeping track of up to four MU pair associations per LDEV (one for CA, three for BC). The following figure shows this configuration.



Correspondence between a configuration file and mirror descriptors

The following table shows how MU usage can indicate that a pair is CA, BC, or either.

Leaving MU blank means “0, and usable for either a CA or BC pair.” An explicit **0** (or **1** or **2**) means BC only.

HORCM_DEV entries can be in random order.

No	MU designations in configuration file						MU#0		BC only	
							CA	BC	MU#1	MU#2
1	HORCM_DEV #dev_group dev_name port# TargetID LU# MU# Oradb oradev1 CL1-D 2 1	oradev1	oradev1							
2	HORCM_DEV #dev_group dev_name port# TargetID LU# MU# Oradb oradev1 CL1-D 2 1 Oradb1 oradev11 CL1-D 2 1 Oradb2 oradev21 CL1-D 2 1	oradev1	oradev1	oradev11	oradev21					
3	HORCM_DEV #dev_group dev_name port# TargetID LU# MU# Oradb oradev1 CL1-D 2 1 Oradb1 oradev11 CL1-D 2 0 Oradb2 oradev21 CL1-D 2 1 Oradb3 oradev31 CL1-D 2 1	oradev1	oradev11	oradev21	oradev31					
4	HORCM_DEV #dev_group dev_name port# TargetID LU# MU# Oradb oradev1 CL1-D 2 1		oradev1							
5	HORCM_DEV #dev_group dev_name port# TargetID LU# MU# Oradb oradev1 CL1-D 2 1 Oradb1 oradev11 CL1-D 2 1 Oradb2 oradev21 CL1-D 2 1		oradev1	oradev11	oradev21					

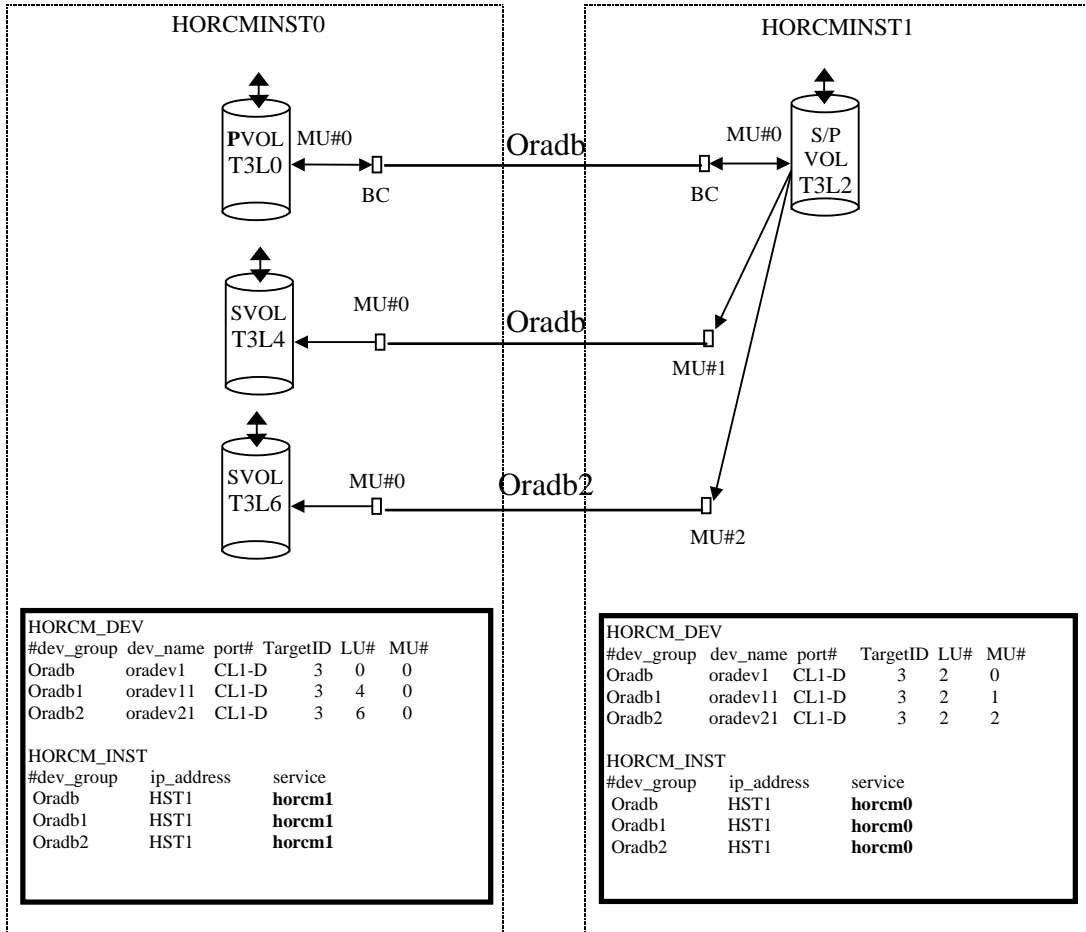
Cascading connection configuration files

The following are examples of configuration files and the corresponding (`pairedisplay`) outputs.

Cascading BC

You should use two configuration files to describe a cascaded (tiered) BC configuration, as shown in the preceding figure.

Instance 0, in this case, describes the root (and all leaf) volumes (as if the normal diagram had been folded over from right to left). Instance 1 describes the intermediate S-VOL/P-VOLs.



The instance 0 configuration file in the figure above specifies that:

- Three BC pairs are recognized.
- None of the BC pairs are an intermediate S-VOL/P-VOLs in a cascade, because each has a different TID/LUN combination, with an explicit MU# of 0.

The instance 1 configuration file in the preceding figure specifies that:

- Three BC pairs are recognized.
- The BC pairs are intermediate S-VOL/P-VOLs, because the TID/LUN combinations are all the same.

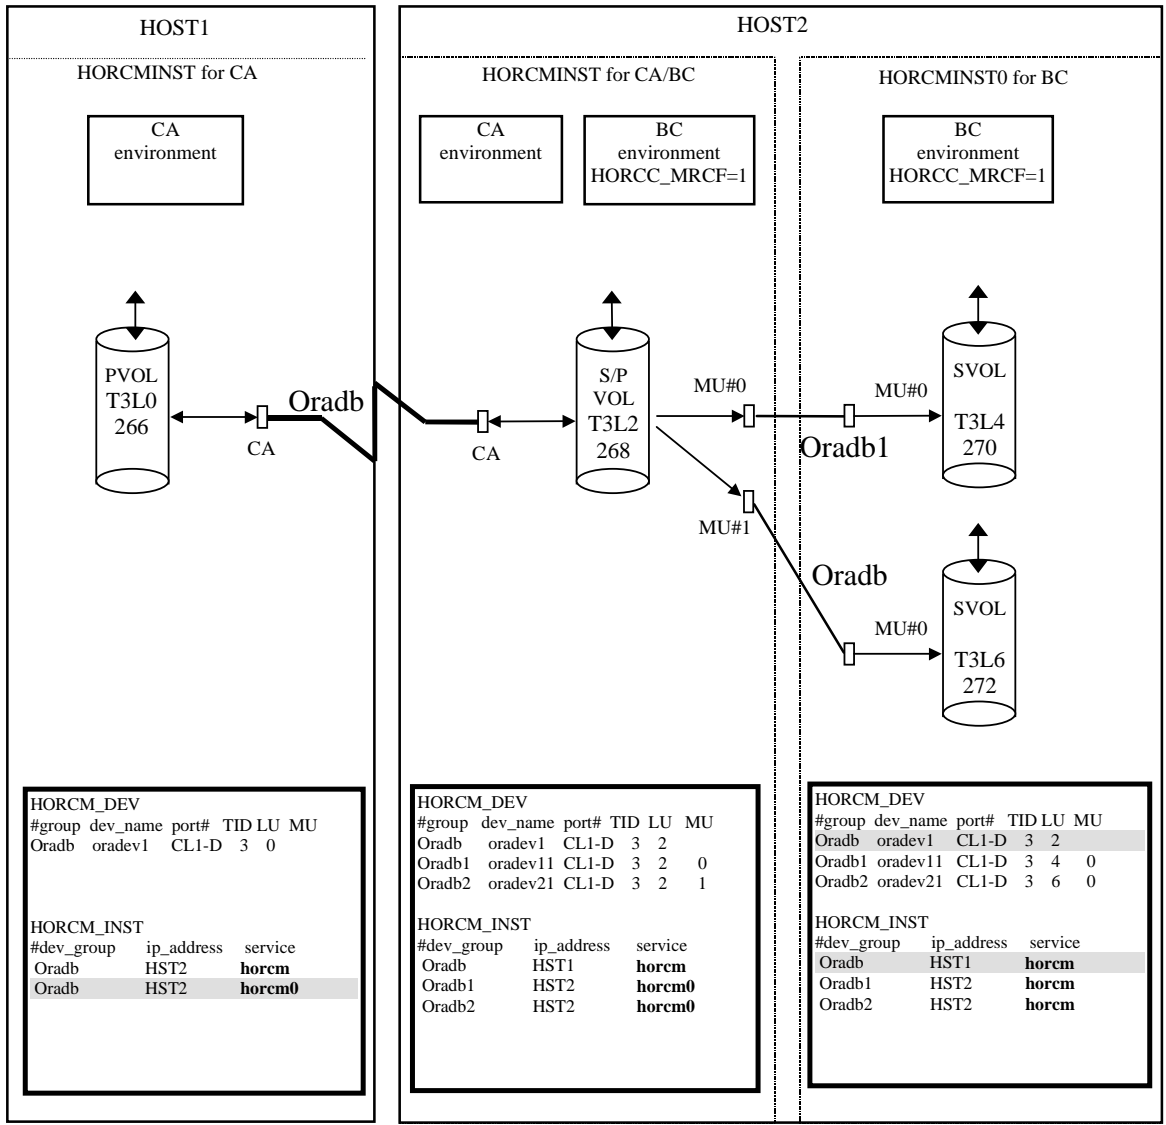
Connecting CA and BC

You can use three configuration files to describe a CA/BC cascaded configuration, as shown in the following figure.

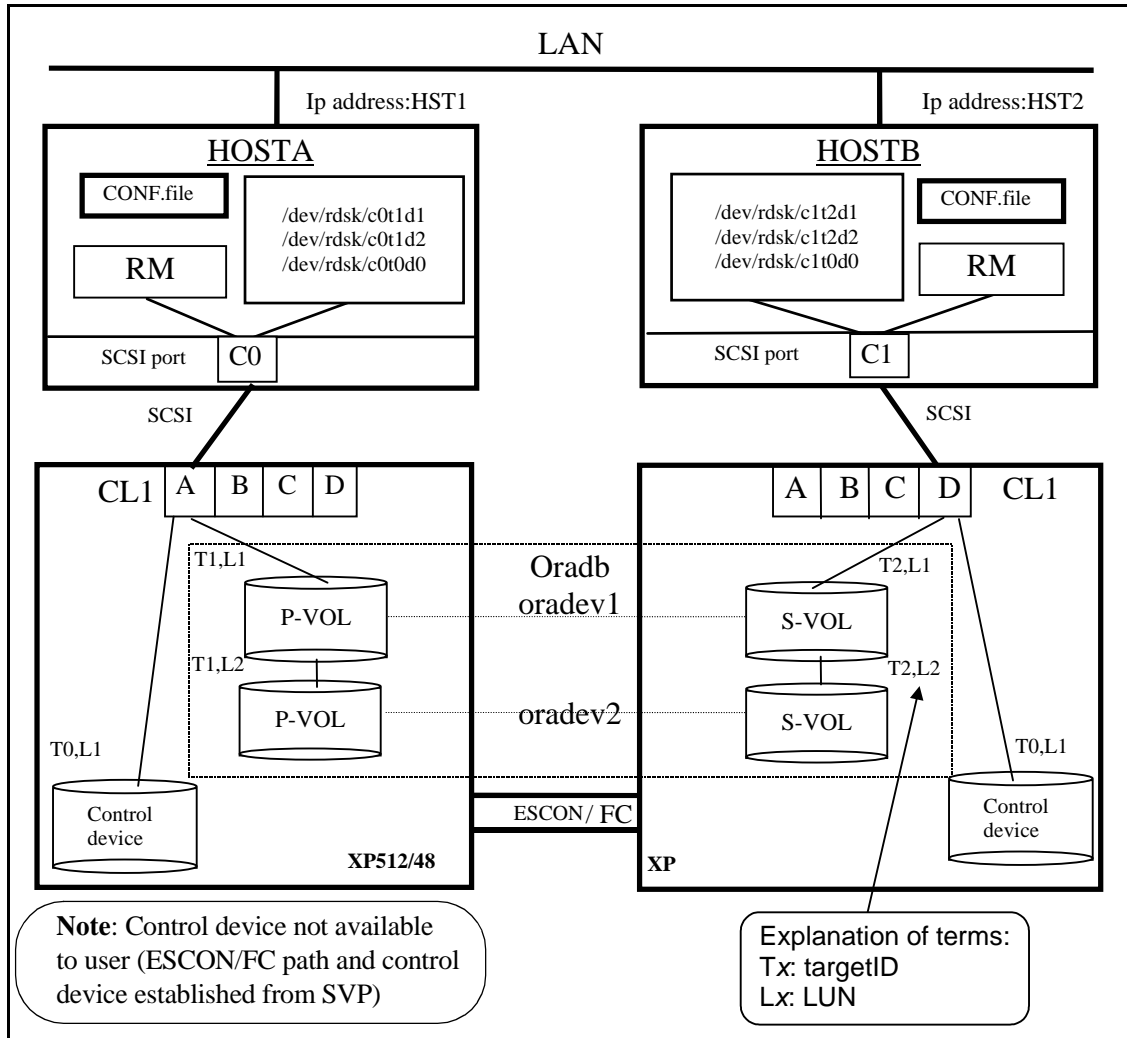
The configuration file in the center (HORCMINST for CA/BC) can be interpreted as follows:

- The first line is ambiguous as to whether the pair is CA or BC because the MU# is blank.
- The second line shows that the MU# is 0, meaning BC. The Port/TID/LUN is the same as on the first line, so you know that the top line refers to a CA pair.
- Since all three lines use the same Port/TID/LUN, you know that the pair is the intermediate part of a CA, cascaded BC configuration.

Another hint that **oradb** is a CA pair is that its remote pair serving host is a host other than HOST2, which serves this RM instance.



CA configuration (remote CA, two hosts)



Configuration file for HOSTA (/etc/horcm.conf) on page 271

```
HORCM_MON
#ip_address  service  poll(10ms)  timeout(10ms)
HST1        horcm    1000        3000
```

```
HORCM_CMD
#dev_name
/dev/xxx (See "Note" on page 272)
```

```
HORCM_DEV
#dev_group  dev_name  port#  TargetID  LU#
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) on page 271

```
HORCM_MON
#ip_address  service  poll(10ms)  timeout(10ms)
HST2        horcm    1000        3000
```

```
HORCM_CMD
#dev_name
/dev/xxx (See "Note" on page 272)
```

```
HORCM_DEV
#dev_group  dev_name  port#  TargetID  LU#
Oradb       oradev1  CL1-D  2         1
Oradb       oradev2  CL1-D  2         2
```

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

Note

There must be at least one command device described in the configuration definition for every instance. Up to 16 instances can use the same command device via the same port. Instances beyond 16 must use a different SCSI path.

The following shows an example of the (raw) control device file format that must be used. $HOST_x = HOSTA, HOSTB, etc...$

- HP-UX
HORCM_CMD for $HOST_x$... **/dev/rdisk/c0t0d1**
- Solaris
HORCM_CMD for $HOST_x$... **/dev/rdisk/c0t0d1s2**
- AIX
HORCM_CMD for $HOST_x$... **/dev/rhdiskNN**
Where NN is the device number assigned automatically by AIX.
- Digital UNIX
HORCM_CMD for $HOST_x$... **/dev/rrzbNNc**
Where NN is device number (BUS number \times 8 + target ID) defined by Digital UNIX.
- DYNIX/ptx
HORCM_CMD for $HOST_x$... **/dev/rdisk/sdNN**
Where NN is the device number assigned automatically by DYNIX/ptx.
- Windows NT/2000/2003
HORCM_CMD for $HOST_x$... **\\.\PhysicalDriveN**
Where N is the device number assigned automatically by Windows NT/2000/2003.
- Linux, xLinux
HORCM_CMD for $HOST_x$... **/dev/sdN**
Where N is the device number assigned automatically by Linux/xLinux.

CA (remote CA, two host) command examples

Commands from HOSTA in the figure on [page 271](#)

The following examples employ CA commands from HOSTA.

- Designate a group name (**Oradb**) and a local host P-VOL:

```
# paircreate -g Oradb -f never -v1
```

This command begins a pair coupling between the volumes designated as **Oradb** in the configuration definition file and begins copying the two pairs (in the example configuration).

- Designate a volume name (**oradev1**) and a local host P-VOL:

```
# paircreate -g Oradb -d oradev1 -f never -v1
```

This command begins a pair coupling between the volumes designated as **oradev1** in the configuration definition file.

In the example configuration, this pairs **CL1-A, T1, L1** and **CL1-D, T2, L1**

- Designate a group name and confirm pair volume state:

```
# pairdisplay -g Oradb
```

Group	PairVol (L/R)	(P,T#,L#)	Seq#	LDEV#	..P/S	Status	Fence	Seq#	P-LDEV#	M
oradb	oradev1 (L)	(CL1-A, 1,1)	30053	18	..P-VOL	COPY	NEVER	30054	19	-
oradb	oradev1 (R)	(CL1-D, 2,1)	30054	19	..S-VOL	COPY	NEVER	-----	18	-
oradb	oradev2 (L)	(CL1-A, 1,2)	30053	20	..P-VOL	COPY	NEVER	30054	21	-
oradb	oradev2 (R)	(CL1-D, 2,2)	30054	21	..S-VOL	COPY	NEVER	-----	20	-

Commands from HOSTB in the figure on [page 271](#)

The following examples employ CA commands from HOSTB.

- Designate a group name and a remote host P-VOL:

```
# paircreate -g Oradb -f never -vr
```

This command begins a pair coupling between the volumes designated as **Oradb** in the configuration definition file and begins copying the two pairs (in the example configuration).

- Designate a volume name (**oradev1**) and a remote host P-VOL:

```
# paircreate -g Oradb -d oradev1 -f never -vr
```

This command begins a pair coupling between the volumes designated as **oradev1** in the configuration definition file.

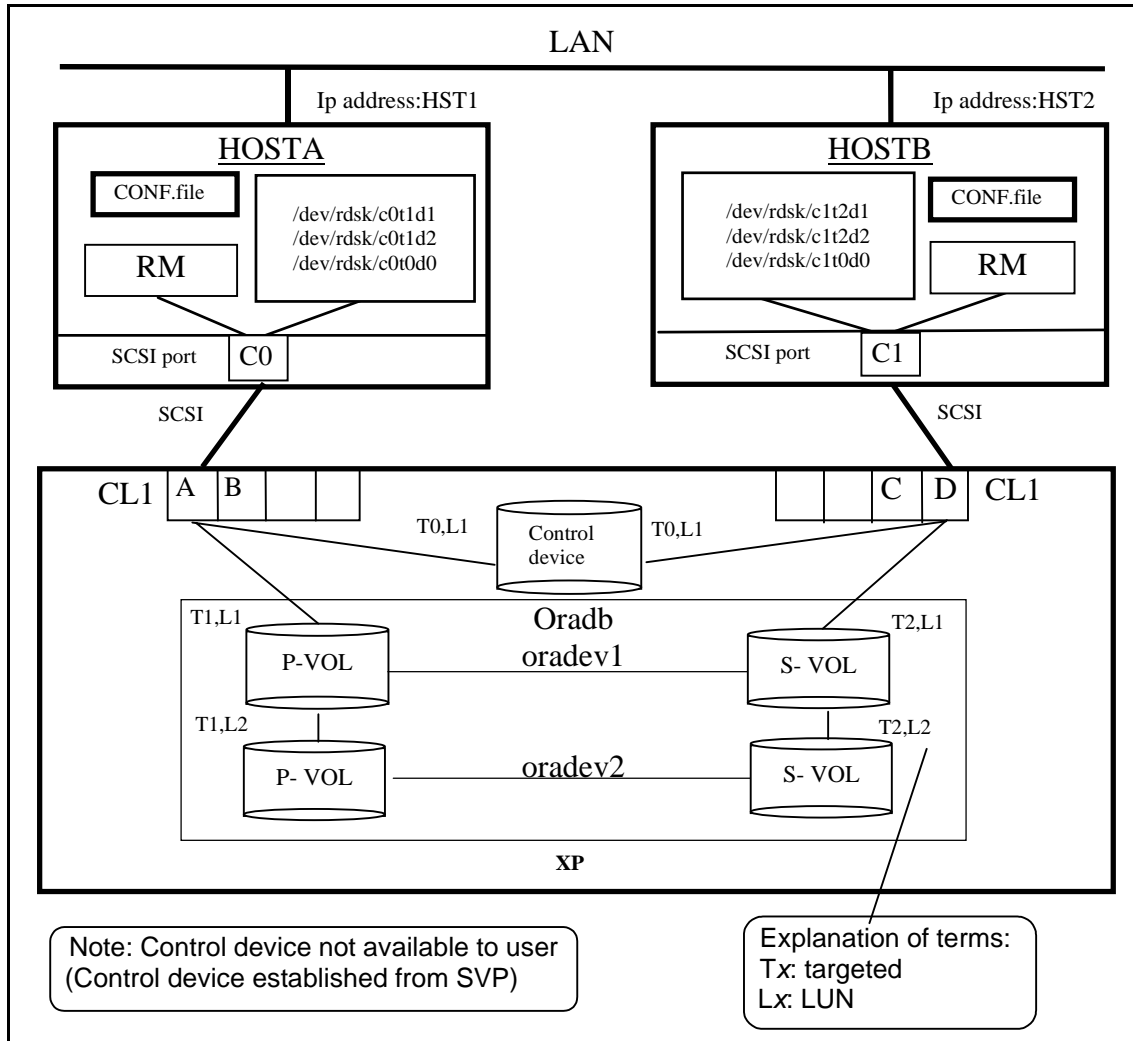
In the example configuration, this pairs **CL1-A, T1, L1** and **CL1-D, T2, L1**

- Designate a group name and confirm pair volume state:

```
# pairdisplay -g Oradb
```

Group	PairVol (L/R)	(P,T#,L#)	Seq#	LDEV#..P/S	Status	Fence	Seq#	P-LDEV#	M
oradb	oradev1 (L)	(CL1-D, 2,1)	30054	19	..S-VOL	COPY	NEVER, -----	18	-
oradb	oradev1 (R)	(CL1-A, 1,1)	30053	18	..P-VOL	COPY	NEVER, 30054	19	-
oradb	oradev2 (L)	(CL1-D, 2,2)	30054	21	..S-VOL	COPY	NEVER, -----	20	-
oradb	oradev2 (R)	(CL1-A, 1,2)	30053	20	..P-VOL	COPY	NEVER, 30054	21	-

CA configuration (local loopback, two hosts)



Configuration file for HOSTA on [page 276 \(/etc/horcm.conf\)](#)

```
HORCM_MON
#ip_address  service  poll(10ms)  timeout(10ms)
HST1         horcm      1000        3000
```

```
HORCM_CMD
#dev_name
/dev/xxx (See "Note" on page 272)
```

```
HORCM_DEV
#dev_group  dev_name  port#  TargetID  LU#
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 on [page 276 \(/etc/horcm.conf\)](#)

```
HORCM_MON
#ip_address  service  poll(10ms)  timeout(10ms)
HST2         horcm      1000        3000
```

```
HORCM_CMD
#dev_name
/dev/xxx (See "Note" on page 272)
```

```
HORCM_DEV
#dev_group  dev_name  port#  TargetID  LU#
Oradb       oradev1  CL1-D  2         1
Oradb       oradev2  CL1-D  2         2
```

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

CA (local loopback, two hosts) command examples

Commands from HOSTA in the figure on [page 276](#)

The following examples employ RM commands from HOSTA.

- Designate a group name (**Oradb**) and a local host P-VOL:

```
# paircreate -g Oradb -f never -v1
```

This command begins a pair coupling between the volumes designated as **Oradb** in the configuration definition file and begins copying the two pairs (in the example configuration).

- Designate a volume name (**oradev1**) and a local host P-VOL:

```
# paircreate -g Oradb -d oradev1 -f never -v1
```

This command begins a pair coupling between the volumes designated as **oradev1** in the configuration definition file.

In the example configuration, this pairs **CL1-A, T1, L1** and **CL1-D, T2, L1**

- Designate a group name and confirm pair volume state:

```
# pairdisplay -g Oradb
```

Group	PairVol (L/R)	(P,T#,L#)	Seq#	LDEV#	..P/S	Status	Fence	Seq#	P-LDEV#	M
oradb	oradev1 (L)	(CL1-A, 1,1)	30053	18	..P-VOL	COPY	NEVER	30053	19	-
oradb	oradev1 (R)	(CL1-D, 2,1)	30053	19	..S-VOL	COPY	NEVER	-----	18	-
oradb	oradev2 (L)	(CL1-A, 1,2)	30053	20	..P-VOL	COPY	NEVER	30053	21	-
oradb	oradev2 (R)	(CL1-D, 2,2)	30053	21	..S-VOL	COPY	NEVER	-----	20	-

Commands from HOSTB in the figure on [page 276](#)

The following examples employ RM commands from HOSTB.

- Designate a group name and a remote host P-VOL:

```
# paircreate -g Oradb -f never -vr
```

This command begins a pair coupling between the volumes designated as **Oradb** in the configuration definition file and begins copying the two pairs (in the example configuration).

- Designate a volume name (**oradev1**) and a remote host P-VOL:

```
# paircreate -g Oradb -d oradev1 -f never -vr
```

This command begins a pair coupling between the volumes designated as **oradev1** in the configuration definition file.

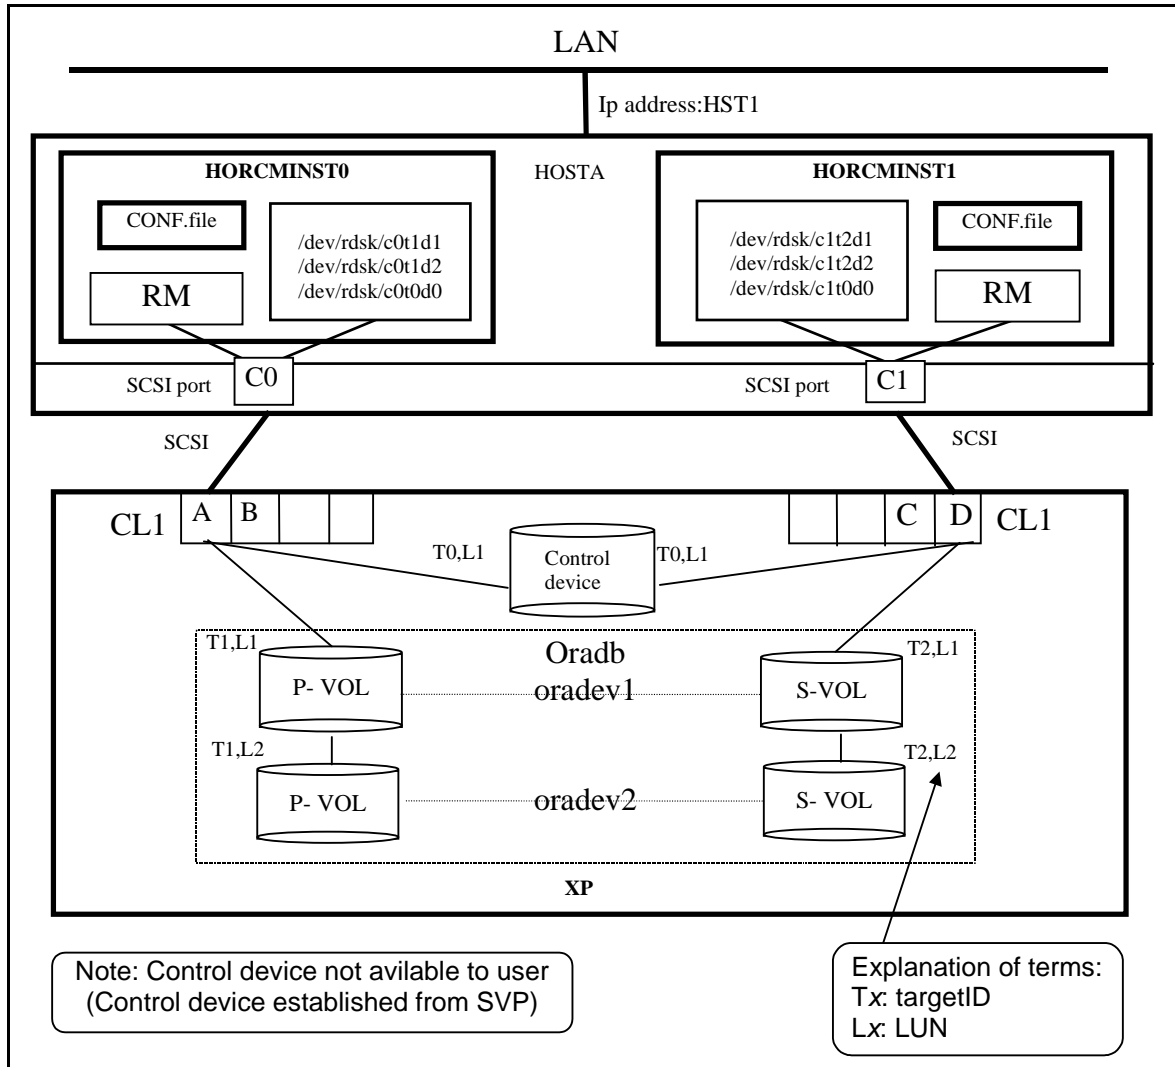
In the example configuration, this pairs **CL1-A, T1, L1** and **CL1-D, T2, L1**

- Designate a group name and confirm pair volume state:

```
# pairdisplay -g Oradb
```

Group	PairVol (L/R)	(P,T#,L#)	Seq#	LDEV#..P/S	Status	Fence	Seq#	P-LDEV#	M
oradb	oradev1 (L)	(CL1-D, 2,1)	30053	19..	S-VOL	COPY	NEVER, -----	18	-
oradb	oradev1 (R)	(CL1-A, 1,1)	30053	18..	P-VOL	COPY	NEVER, 30053	19	-
oradb	oradev2 (L)	(CL1-D, 2,2)	30053	21..	S-VOL	COPY	NEVER, -----	20	-
oradb	oradev2 (R)	(CL1-A, 1,2)	30053	20..	P-VOL	COPY	NEVER, 30053	21	-

CA configuration (two RM instances, one host)



Configuration file for HOSTA, Instance 0 shown on [page 280](#) (/etc/horcm0.conf)

```
HORCM_MON
#ip_address      service      poll(10ms)    timeout(10ms)
HST1             horcm0       1000          3000
```

```
HORCM_CMD
#dev_name
/dev/xxx (See "Note" on page 272)
```

```
HORCM_DEV
#dev_group      dev_name      port#          TargetID      LU#
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 HOSTA, Instance 1 shown on [page 280](#) (/etc/horcm1.conf)

```
HORCM_MON
#ip_address      service      poll(10ms)    timeout(10ms)
HST1             horcm1       1000          3000
```

```
HORCM_CMD
#dev_name
/dev/xxx (See "Note" on page 272)
```

```
HORCM_DEV
#dev_group      dev_name      port#          TargetID      LU#
Oradb           oradev1      CL1-D         2             1
Oradb           oradev2      CL1-D         2             2
```

```
HORCM_INST
#dev_group      ip_address    service
Oradb           HST1         horcm0
```

CA (two RM instances, one host) command examples

Commands from HOSTA, Instance 0 in the figure on [page 280](#)

The following examples employ RM commands from HOSTA, Instance 0.

- Set the instance number. (If C shell)

```
# setenv HORCMINST 0
```

(Windows NT/2000/2003) `set HORCMINST=0`

- Designate a group name (**Oradb**) and a local instance P-VOL:

```
# paircreate -g Oradb -f never -v1
```

This command begins a pair coupling between the two pairs of volumes designated as **Oradb** in the configuration definition file.

- Designate a volume name (**oradev1**) and a local instance P-VOL:

```
# paircreate -g Oradb -d oradev1 -f never -v1
```

In the example configuration, this pairs **CL1-A, T1, L1** and **CL1-D, T2, L1**

- Designate a group name and confirm pair volume state:

```
# pairdisplay -g Oradb
```

Group	PairVol (L/R)	(P,T#,L#)	Seq#	LDEV#	..P/S	Status	Fence	Seq#	P-LDEV#	M
oradb	oradev1 (L)	(CL1-A, 1,1)	30053	18	..P-VOL	COPY	NEVER	30053	19	-
oradb	oradev1 (R)	(CL1-D, 2,1)	30053	19	..S-VOL	COPY	NEVER	-----	18	-
oradb	oradev2 (L)	(CL1-A, 1,2)	30053	20	..P-VOL	COPY	NEVER	30053	21	-
oradb	oradev2 (R)	(CL1-D, 2,2)	30053	21	..S-VOL	COPY	NEVER	-----	20	-

Commands from HOSTA, Instance 1 in the figure on page 280

The following examples employ RM commands from HOSTA, Instance 1.

- Set the instance number. (If C shell)

```
# setenv HORCMINST 1
```

(Windows NT/2000/2003) `set HORCMINST=1`

- Designate a group name and a remote instance P-VOL:

```
# paircreate -g Oradb -f never -vr
```

This command begins a pair coupling between the two pairs of volumes designated as **Oradb** in the configuration definition file.

- Designate a volume name (**oradev1**) and a remote instance P-VOL:

```
# paircreate -g Oradb -d oradev1 -f never -vr
```


In the example configuration, this pairs **CL1-A, T1, L1** and **CL1-D, T2, L1**

- Designate a group name and confirm pair volume state:

```
# pairdisplay -g Oradb
```

Group	PairVol (L/R)	(P,T#,L#),	Seq#,	LDEV#..P/S,	Status,	Fence,	Seq#,	P-LDEV#	M
oradb	oradev1 (L)	(CL1-D, 2,1)	30053	19	..S-VOL	COPY	NEVER, -----	18	-
oradb	oradev1 (R)	(CL1-A, 1,1)	30053	18	..P-VOL	COPY	NEVER, 30053	19	-
oradb	oradev2 (L)	(CL1-D, 2,2)	30053	21	..S-VOL	COPY	NEVER, -----	20	-
oradb	oradev2 (R)	(CL1-A, 1,2)	30053	20	..P-VOL	COPY	NEVER, 30053	21	-

BC configuration



Configuration file for HOSTA shown on [page 284](#) (/etc/horcm.conf)

```

HORCM_MON
#ip_address      service      poll(10ms)      timeout(10ms)
HST1             horcm        1000             3000

HORCM_CMD
#dev_name
/dev/xxx (See "Note" on page 272)

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

Oradb1         oradev1-1     CL1-A           1             1             1
Oradb1         oradev1-2     CL1-A           1             2             1

Oradb2         oradev2-1     CL1-A           1             1             2
Oradb2         oradev2-2     CL1-A           1             2             2

HORCM_INST
#dev_group      ip_address      service
Oradb           HST2           horcm
Oradb1         HST3           horcm
Oradb2         HST4           horcm

```

Configuration file for HOSTB shown on [page 284](#) (/etc/horcm.conf)

```

HORCM_MON
#ip_address      service      poll(10ms)      timeout(10ms)
HST2             horcm        1000             3000

HORCM_CMD
#dev_name
/dev/xxx (See "Note" on page 272)

HORCM_DEV
#dev_group      dev_name      port#           TargetID      LU#           MU#
Oradb           oradev1       CL2-B           2             1
Oradb           oradev2       CL2-B           2             2

HORCM_INST
#dev_group      ip_address      service
Oradb           HST1           horcm

```

Configuration file for HOSTC shown on [page 284](#) (/etc/horcm.conf)

```
HORCM_MON
#ip_address      service      poll(10ms)      timeout(10ms)
HST3             horcm        1000             3000

HORCM_CMD
#dev_name
/dev/xxx (See "Note" on page 272)

HORCM_DEV
#dev_group      dev_name      port#           TargetID  LU#  MU#
Oradb1          oradev1-1    CL2-C           2         1
Oradb1          oradev1-2    CL2-C           2         2

HORCM_INST
#dev_group      ip_address    service
Oradb1          HST1         horcm
```

Configuration file for HOSTD shown on [page 284](#) (/etc/horcm.conf)

```
HORCM_MON
#ip_address      service      poll(10ms)      timeout(10ms)
HST4             horcm        1000             3000

HORCM_CMD
#dev_name
/dev/xxx (See "Note" on page 272)

HORCM_DEV
#dev_group      dev_name      port#           TargetID  LU#  MU#
Oradb2          oradev2-1    CL2-D           2         1
Oradb2          oradev2-2    CL2-D           2         2

HORCM_INST
#dev_group      ip_address    service
Oradb2          HST1         horcm
```

BC command examples

Commands from HOSTA shown on page 284 (group Oradb)

- Set the **HORCC_MRCF** environment variable. (If C shell)

```
# setenv HORCC_MRCF 1
```

(Windows NT/2000/2003) **set HORCC_MRCF=1**

- Designate a group name (**Oradb**) and a local host P-VOL:

```
# paircreate -g Oradb -v1
```

This command begins a pair coupling between the two pairs of volumes designated as **Oradb** in the configuration definition file.

- Designate a volume name (**oradev1**) and a local host P-VOL:

```
# paircreate -g Oradb -d oradev1 -v1
```

In the example configuration, this pairs **CL1-A, T1, L1** and **CL1-B, T2, L1**

- Designate a group name and confirm pair volume state:

```
# 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)	30053	18..P-VOL	COPY	30053	20	-
oradb	oradev1 (R)	(CL2-B, 2, 1-0)	30053	20..S-VOL	COPY	-----	18	-
oradb	oradev2 (L)	(CL1-A, 1, 2-0)	30053	19..P-VOL	COPY	30053	21	-
oradb	oradev2 (R)	(CL2-B, 2, 2-0)	30053	21..S-VOL	COPY	-----	19	-

Commands from HOSTB shown on page 284 (group Oradb)

- Set the **HORCC_MRCF** environment variable. (If C shell)

```
# setenv HORCC_MRCF 1
```

(Windows NT/2000/2003) **set HORCC_MRCF=1**

- Designate a group name and a remote host P-VOL:

```
# paircreate -g Oradb -vr
```

This command begins a pair coupling between the two pairs of volumes designated as **Oradb** in the configuration definition file.

- Designate a volume name (**oradev1**) and a remote host P-VOL:

```
# paircreate -g Oradb -d oradev1 -vr
```

In the example configuration, this pairs **CL1-A, T1, L1** and **CL1-B, T2, L1**

- Designate a group name and confirm pair volume state:

```
# pairdisplay -g Oradb
```

Group	PairVol (L/R)	(Port#, TID, LU-M)	Seq#	LDEV#..P/S	Status	Seq#	P-LDEV#	M
oradb	oradev1 (L)	(CL2-B, 2, 1-0)	30053	20..S-VOL	COPY	-----	18	-
oradb	oradev1 (R)	(CL1-A, 1, 1-0)	30053	18..P-VOL	COPY	30053	20	-
oradb	oradev2 (L)	(CL2-B, 2, 2-0)	30053	21..S-VOL	COPY	-----	19	-
oradb	oradev2 (R)	(CL1-A, 1, 2-0)	30053	19..P-VOL	COPY	30053	21	-

Commands from HOSTA shown on [page 284](#) (group Oradb1)

- Set the **HORCC_MRCF** environment variable. (If C shell)

```
# setenv HORCC_MRCF 1
```

(Windows NT/2000/2003) **set HORCC_MRCF=1**

- Designate a group name (**Oradb1**) and a local host P-VOL:

```
# paircreate -g Oradb1 -v1
```

This command begins a pair coupling between the two pairs of volumes designated as **Oradb1** in the configuration definition file.

- Designate a volume name (**oradev1-1**) and a local host P-VOL:

```
# paircreate -g Oradb1 -d oradev1-1 -v1
```

In the example configuration, this pairs **CL1-A, T1, L1** and **CL2-C, T2, L1**

- Designate a group name and confirm pair volume state:

```
# pairdisplay -g Oradb1
```

Group	PairVol (L/R)	(Port#, TID, LU-M)	Seq#	LDEV#..P/S	Status	Seq#	P-LDEV#	M
oradb	oradev1-1 (L)	(CL1-A, 1, 1-1)	30053	18..P-VOL	COPY	30053	22	-
oradb	oradev1-1 (R)	(CL2-C, 2, 1-0)	30053	22..S-VOL	COPY	-----	18	-
oradb	oradev2-2 (L)	(CL1-A, 1, 2-1)	30053	19..P-VOL	COPY	30053	23	-
oradb	oradev2-2 (R)	(CL2-C, 2, 2-0)	30053	23..S-VOL	COPY	-----	19	-

Commands from HOSTC shown on [page 284](#) (group Oradb1)

- Set the **HORCC_MRCF** environment variable. (If C shell)

```
# setenv HORCC_MRCF 1
```

(Windows NT/2000/2003) **set HORCC_MRCF=1**

- Designate a group name and a remote host P-VOL:

```
# paircreate -g Oradb1 -vr
```

This command begins a pair coupling between the two pairs of volumes designated as **Oradb1** in the configuration definition file.

- Designate a volume name (**oradev1-1**) and a remote host P-VOL:

```
# paircreate -g Oradb1 -d oradev1-1 -vr
```

In the example configuration, this pairs **CL1-A, T1, L1** and **CL2-C, T2, L1**

- Designate a group name and confirm pair volume state:

```
# pairdisplay -g Oradb1
```

Group	PairVol (L/R)	(Port#, TID, LU-M)	Seq#	LDEV#..P/S	Status	Seq#	P-LDEV#	M
oradb	oradev1-1 (L)	(CL2-C, 2, 1-0)	30053	22..S-VOL	COPY	-----	18	-
oradb	oradev1-1 (R)	(CL1-A, 1, 1-1)	30053	18..P-VOL	COPY	30053	22	-
oradb	oradev1-2 (L)	(CL2-C, 2, 2-0)	30053	23..S-VOL	COPY	-----	19	-
oradb	oradev1-2 (R)	(CL1-A, 1, 2-1)	30053	19..P-VOL	COPY	30053	23	-

Commands from HOSTA shown on [page 284](#) (group Oradb2)

- Set the **HORCC_MRCF** environment variable. (If C shell)

```
# setenv HORCC_MRCF 1
```

(Windows NT/2000/2003) **set HORCC_MRCF=1**

- Designate a group name (**Oradb2**) and a local host P-VOL:

```
# paircreate -g Oradb1 -v1
```

This command begins a pair coupling between the two pairs of volumes designated as **Oradb2** in the configuration definition file.

- Designate a volume name (**oradev2-1**) and a local host P-VOL:

```
# paircreate -g Oradb2 -d oradev2-1 -v1
```

In the example configuration, this pairs **CL1-A, T1, L1** and **CL2-D, T2, L1**

- Designate a group name and confirm pair volume state:

```
# pairdisplay -g Oradb2
```

Group	PairVol (L/R)	(Port#,TID,LU-M)	Seq#	LDEV#..P/S	Status	Seq#	P-LDEV#	M
oradb	oradev2-1 (L)	(CL1-A, 1, 1-2)	30053	18..P-VOL	COPY	30053	24	-
oradb	oradev2-1 (R)	(CL2-D, 2, 1-0)	30053	24..S-VOL	COPY	-----	18	-
oradb	oradev2-2 (L)	(CL1-A, 1, 2-2)	30053	19..P-VOL	COPY	30053	25	-
oradb	oradev2-2 (R)	(CL2-D, 2, 2-0)	30053	25..S-VOL	COPY	-----	19	-

Commands from HOSTD shown on [page 284](#) (group Oradb2)

- Set the **HORCC_MRCF** environment variable. (If C shell)

```
# setenv HORCC_MRCF 1
```

(Windows NT/2000/2003) **set HORCC_MRCF=1**

- Designate a group name and a remote host P-VOL:

```
# paircreate -g Oradb2 -vr
```

This command begins a pair coupling between the two pairs of volumes designated as **Oradb2** in the configuration definition file.

- Designate a volume name (**oradev2-1**) and a remote host P-VOL:

```
# paircreate -g Oradb2 -d oradev2-1 -vr
```

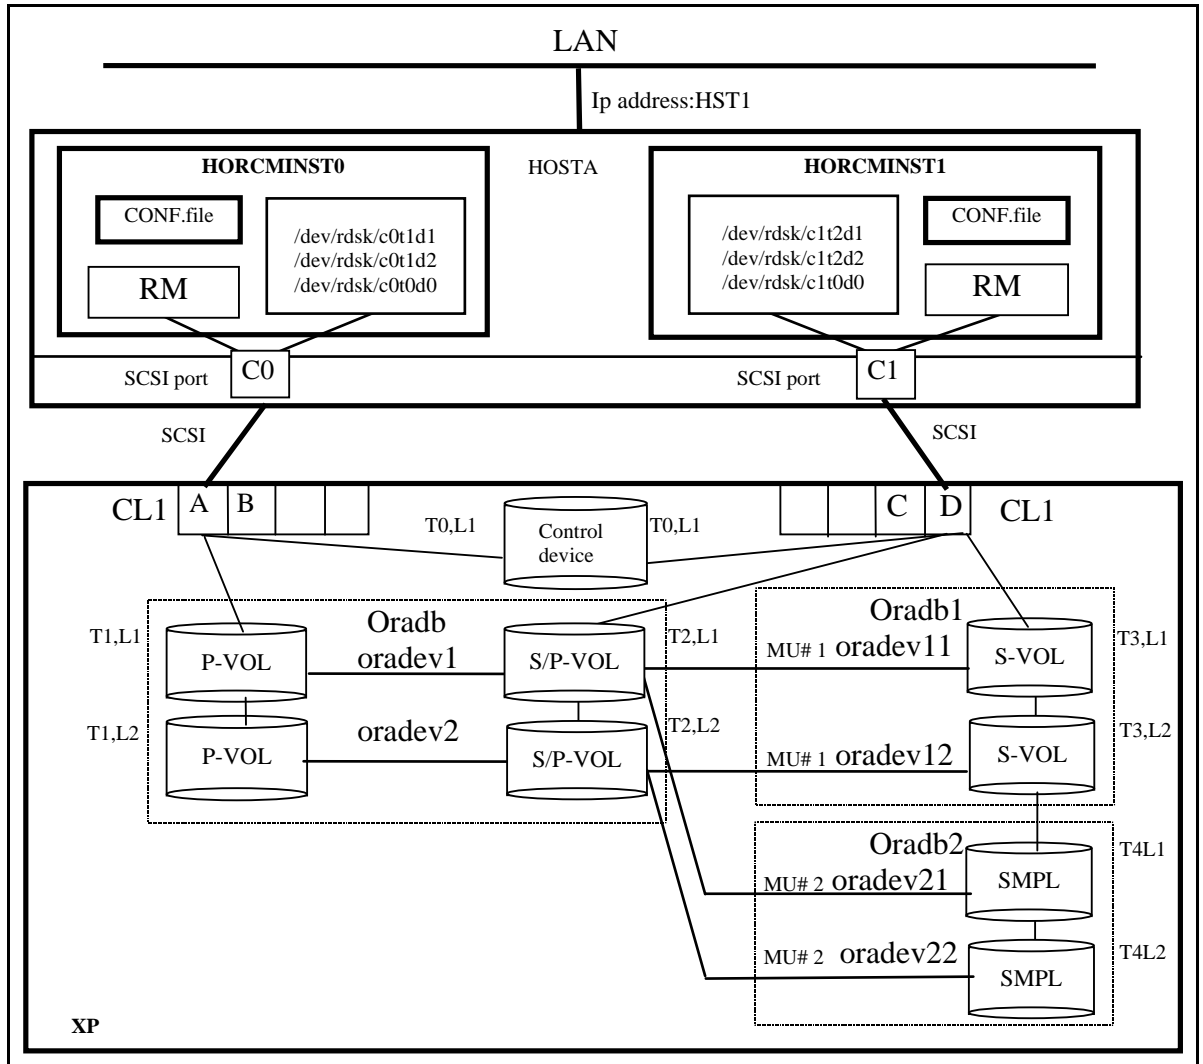
In the example configuration, this pairs **CL1-A, T1, L1** and **CL2-D, T2, L1**

- Designate a group name and confirm pair volume state:

```
# pairdisplay -g Oradb2
```

Group	PairVol (L/R)	(Port#, TID, LU-M)	Seq#	LDEV#..P/S	Status	Seq#	P-LDEV#	M
oradb	oradev2-1 (L)	(CL2-D, 2, 1-0)	30053	24..S-VOL	COPY	-----	18	-
oradb	oradev2-1 (R)	(CL1-A, 1, 1-2)	30053	18..P-VOL	COPY	30053	24	-
oradb	oradev2-2 (L)	(CL2-D, 2, 2-0)	30053	25..S-VOL	COPY	-----	19	-
oradb	oradev2-2 (R)	(CL1-A, 1, 2-2)	30053	19..P-VOL	COPY	30053	25	-

Configuration for a BC cascaded connection



Configuration file for HOSTA shown on page 293 (/etc/horcm0.conf)

```

HORCM_MON
#ip_address      service      poll(10ms)      timeout(10ms)
HST1             horcm0        1000             3000

HORCM_CMD
#dev_name
/dev/xxx (See "Note" on page 272)

HORCM_DEV
#dev_group      dev_name      port#           TargetID      LU#      MU#
Oradb           oradev1      CL1-A           1             1        0
Oradb           oradev2      CL1-A           1             2        0
Oradb1          oradev11     CL1-D           3             1        0
Oradb1          oradev12     CL1-D           3             2        0
Oradb2          oradev21     CL1-D           4             1        0
Oradb2          oradev22     CL1-D           4             2        0

HORCM_INST
#dev_group      ip_address      service
Oradb           HST1           horcm1
Oradb1          HST1           horcm1
Oradb2          HST1           horcm1

```

Configuration file for HOSTA shown on page 293 (/etc/horcm1.conf)

```

HORCM_MON
#ip_address      service      poll(10ms)      timeout(10ms)
HST1             horcm1        1000             3000

HORCM_CMD
#dev_name
/dev/xxx (See "Note" on page 272)

HORCM_DEV
#dev_group      dev_name      port#           TargetID      LU#      MU#
Oradb           oradev1      CL1-D           2             1        0
Oradb           oradev2      CL1-D           2             2        0
Oradb1          oradev11     CL1-D           2             1        1
Oradb1          oradev12     CL1-D           2             2        1
Oradb2          oradev21     CL1-D           2             1        2
Oradb2          oradev22     CL1-D           2             2        2

HORCM_INST
#dev_group      ip_address      service
Oradb           HST1           horcm0
Oradb1          HST1           horcm0
Oradb2          HST1           horcm0

```

BC cascaded connection command examples

Commands from HOSTA, Instance 0 shown on [page 293](#)

The following examples employ RM commands from HOSTA, Instance 0.

- When the command execution environment is not set, set the instance number. (If C shell)

```
# setenv HORCMINST 0
(Windows NT/2000/2003) set HORCMINST=0
```

- Set the **HORCC_MRCF** environment variable. (If C shell)

```
# setenv HORCC_MRCF 1
(Windows NT/2000/2003) set HORCC_MRCF=1
```

- Designate group names (**Oradb** and **Oradb1**) and a local instance P-VOL:

```
# paircreate -g Oradb -vl
# paircreate -g Oradb1 -vr
```

This command begins a pair coupling between the four pairs of volumes designated as **Oradb** and **Oradb1** in the configuration definition file.

- Designate a group name and confirm pair states:

```
# pairdisplay -g oradb -m cas
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) 30053 266..P-VOL PAIR 30053 268 -
oradb  oradev1 (R) (CL1-D, 2, 1-0) 30053 268..S-VOL PAIR ----- 266 -
oradb1 oradev11 (L) (CL1-D, 2, 1-1) 30053 268..P-VOL PAIR 30053 270 -
oradb2 oradev21 (R) (CL1-D, 2, 1-2) 30053 268..SMPL ---- ----- ---- -
oradb  oradev2 (L) (CL1-A, 1, 1-0) 30053 267..P-VOL PAIR 30053 269 -
oradb  oradev2 (R) (CL1-D, 2, 2-0) 30053 269..S-VOL PAIR ----- 267 -
oradb1 oradev12 (L) (CL1-D, 2, 2-1) 30053 269..P-VOL PAIR 30053 271 -
oradb2 oradev22 (R) (CL1-D, 2, 2-2) 30053 269..SMPL ---- ----- ---- -
```

Commands from HOSTA, Instance 1 shown on page 293

The following examples employ RM commands from HOSTA, Instance 1.

- Set the instance number. (If C shell)

```
# setenv HORCMINST 1
(Windows NT/2000/2003) set HORCMINST=1
```

- Set the **HORCC_MRCF** environment variable. (If C shell)

```
# setenv HORCC_MRCF 1
(Windows NT/2000/2003) set HORCC_MRCF=1
```

- Designate group names (**Oradb** and **Oradb1**) and a remote instance P-VOL:

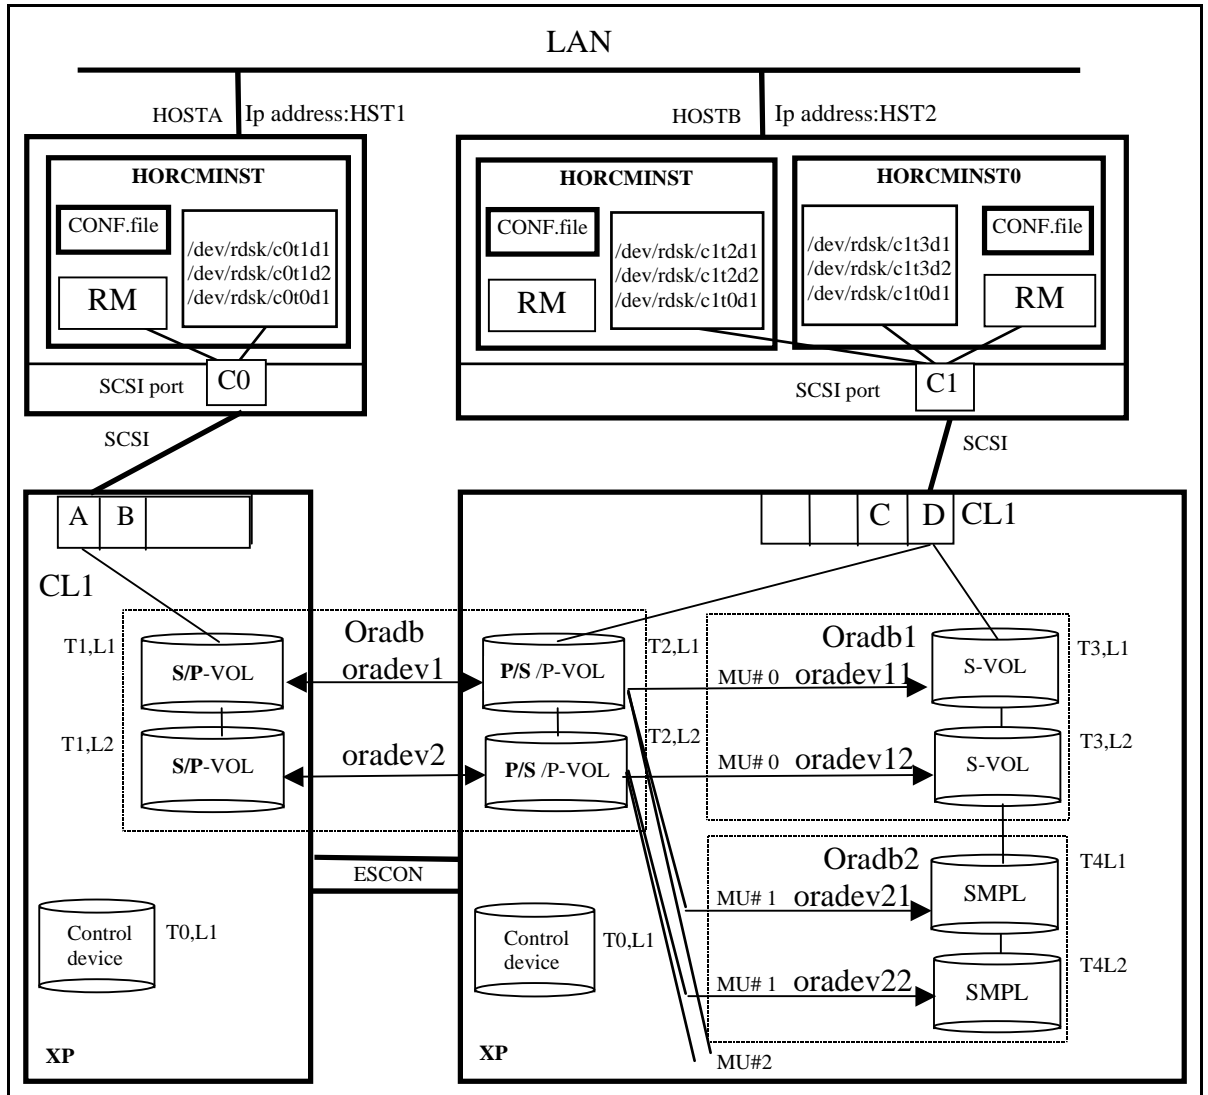
```
# paircreate -g Oradb -vr
# paircreate -g Oradb1 -v1
```

This command begins a pair coupling between the four pairs of volumes designated as **Oradb** and **Oradb1** in the configuration definition file.

- Designate a group name and confirm pair states:

```
# pairdisplay -g oradb -m cas
Group PairVol (L/R) (Port#,TID,LU-M), Seq#, LDEV#..P/S, Status, Seq#, P-LDEV# M
oradb oradev1 (L) (CL1-D, 2, 1-0) 30053 268..P-VOL PAIR ----- 266 -
oradb oradev11 (R) (CL1-D, 2, 1-1) 30053 268..S-VOL PAIR 30053 270 -
oradb1 oradev21 (L) (CL1-D, 2, 1-2) 30053 268..P-VOL ---- - - - - -
oradb2 oradev1 (R) (CL1-A, 1, 1-0) 30053 266..SMPL PAIR 30053 268 -
oradb oradev2 (L) (CL1-D, 2, 2-0) 30053 269..P-VOL PAIR ----- 267 -
oradb1 oradev12 (R) (CL1-D, 2, 2-1) 30053 269..S-VOL PAIR 30053 271 -
oradb2 oradev22 (L) (CL1-D, 2, 2-2) 30053 269..P-VOL ---- - - - - -
oradb oradev2 (R) (CL1-A, 1, 2-0) 30053 267..SMPL PAIR 30053 269 -
```


Configuration for a CA/BC cascaded connection



Configuration file for HOSTA shown on page 297 (/etc/horcm.conf)

```

HORCM_MON
#ip_address      service      poll(10ms)      timeout
HST1             horcm        1000            3000

HORCM_CMD
#dev_name
/dev/xxx (See "Note" on page 272)

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

HORCM_INST
#dev_group      ip_address    service
Oradb           HST2         horcm
Oradb           HST2         horcm0

```

Configuration file for HOSTB shown on page 297 (/etc/horcm.conf)

```

HORCM_MON
#ip_address      service      poll(10ms)      timeout(10ms)
HST2             horcm1       1000            3000

HORCM_CMD
#dev_name
/dev/xxx (See "Note" on page 272)

HORCM_DEV
#dev_group      dev_name      port#           TargetID      LU#    MU#
Oradb           oradev1       CL1-D           2             1      1
Oradb           oradev2       CL1-D           2             2
Oradb1         oradev11      CL1-D           2             1      0
Oradb1         oradev12      CL1-D           2             2      0
Oradb2         oradev21      CL1-D           2             1      1
Oradb2         oradev22      CL1-D           2             2      1

HORCM_INST
#dev_group      ip_address    service
Oradb           HST1         horcm
Oradb1         HST2         horcm0
Oradb2         HST2         horcm0

```

Configuration file for HOSTB shown on [page 297](#) (`/etc/horcm0.conf`)

```
HORCM_MON
#ip_address    service    poll(10ms)    timeout(10ms)
HST2          horcm0      1000          3000

HORCM_CMD
#dev_name
/dev/xxx      (See "Note" on page 272)

HORCM_DEV
#dev_group    dev_name    port#         TargetID    LU#    MU#
Oradb       oradev1    CL1-D        2          1
Oradb       oradev2    CL1-D        2          2
Oradb1       oradev11    CL1-D         3           1      0
Oradb1       oradev12    CL1-D         3           2      0
Oradb2       oradev21    CL1-D         4           1      0
Oradb2       oradev22    CL1-D         4           2      0

HORCM_INST
#dev_group    ip_address    service
Oradb         HST1         horcm
Oradb1       HST2         horcm
Oradb2       HST2         horcm
```

CA/BC cascaded connection command examples

Commands from HOSTA and HOSTB shown on [page 297](#)

The following examples employ RM commands from HOSTA and HOSTB.

- Set the **HORCC_MRCF** environment variable. (If C shell)

```
# setenv HORCC_MRCF 1
```

(Windows NT/2000/2003) **set HORCC_MRCF=1**
- Designate a group name (**Oradb**) on the CA environment of HOSTA:

```
# paircreate -g Oradb -v1
```

- Designate a group name (**Oradb1**) on the BC environment of HOSTB:

```
# paircreate -g Oradb1 -v1
```

This command begins a pair coupling between the four pairs of volumes designated as **Oradb** and **Oradb1** in the configuration definition file.

- Designate a group name and confirm pair volume state on HOSTA:

```
# pairdisplay -g oradb -m cas
```

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)	30053	266..SMPL	----	-----	----	-
oradb	oradev1 (L)	(CL1-A, 1, 1)	30053	266..P-VOL	COPY	30053	268	-
oradb1	oradev11 (R)	(CL1-D, 2, 1-0)	30053	268..P-VOL	COPY	30053	270	-
oradb2	oradev21 (R)	(CL1-D, 2, 1-1)	30053	268..SMPL	----	-----	----	-
oradb	oradev1 (R)	(CL1-D, 2, 1)	30053	268..S-VOL	COPY	-----	266	-
oradb	oradev2 (L)	(CL1-A, 1, 2-0)	30053	267..SMPL	----	-----	----	-
oradb	oradev2 (L)	(CL1-A, 1, 2)	30053	267..P-VOL	COPY	30053	269	-
oradb1	oradev12 (R)	(CL1-D, 2, 2-0)	30053	269..P-VOL	COPY	30053	271	-
oradb2	oradev22 (R)	(CL1-D, 2, 2-1)	30053	269..SMPL	----	-----	----	-
oradb	oradev2 (R)	(CL1-D, 2, 2)	30053	269..S-VOL	COPY	-----	267	-

Commands from HOSTB shown on page 297

The following examples employ RM commands from HOSTB.

- Set the **HORCC_MRCF** environment variable. (If C shell)

```
# setenv HORCC_MRCF 1
```

(Windows NT/2000/2003) **set HORCC_MRCF=1**

- Designate a group name (**Oradb**) on the CA environment of HOSTB:

```
# paircreate -g Oradb -vr
```

- Designate a group name (**Oradb1**) on the BC environment of HOSTB:

```
# paircreate -g Oradb1 -vl
```

This command begins a pair coupling between the four pairs of volumes designated as **Oradb** in the configuration definition file.

- Designate a group name and confirm pair volume state on the CA environment of HOSTB:

```
# pairdisplay -g oradb -m cas
```

Group	PairVol (L/R)	(Port#,TID,LU-M)	Seq#	LDEV#..P/S	Status	Seq#	P-LDEV#	M
oradb1	oradev11 (L)	(CL1-D, 2, 1-0)	30053	268..P-VOL	PAIR	30053	270	-
oradb2	oradev21 (L)	(CL1-D, 2, 1-1)	30053	268..SMPL	----	-----	----	-
oradb	oradev1 (L)	(CL1-D, 2, 1)	30053	268..S-VOL	PAIR	-----	266	-
oradb	oradev1 (R)	(CL1-A, 1, 1-0)	30053	266..SMPL	----	-----	----	-
oradb	oradev1 (R)	(CL1-A, 1, 1)	30053	266..P-VOL	PAIR	30053	268	-
oradb1	oradev12 (L)	(CL1-D, 2, 2-0)	30053	269..P-VOL	PAIR	30053	271	-
oradb2	oradev22 (L)	(CL1-D, 2, 2-1)	30053	269..SMPL	----	-----	----	-
oradb	oradev2 (L)	(CL1-D, 2, 2)	30053	269..S-VOL	PAIR	-----	267	-
oradb	oradev2 (R)	(CL1-A, 1, 2-0)	30053	267..SMPL	----	-----	----	-
oradb	oradev2 (R)	(CL1-A, 1, 2)	30053	267..P-VOL	PAIR	30053	269	-

- Designate a group name and confirm BC pair states from HOSTB:

```
# pairdisplay -g oradb1 -m cas
```

Group	PairVol (L/R)	(Port#, TID, LU-M)	Seq#	LDEV#..P/S	Status	Seq#	P-LDEV#	M
oradb1	oradev11 (L)	(CL1-D, 2, 1-0)	30053	268..P-VOL	PAIR	30053	270	-
oradb2	oradev21 (L)	(CL1-D, 2, 1-1)	30053	268..SMPL	----	-----	----	-
oradb	oradev1 (L)	(CL1-D, 2, 1)	30053	268..S-VOL	PAIR	-----	266	-
oradb1	oradev11 (L)	(CL1-D, 3, 1-0)	30053	270..S-VOL	PAIR	-----	268	-
oradb1	oradev12 (L)	(CL1-D, 2, 2-0)	30053	269..P-VOL	PAIR	30053	271	-
oradb2	oradev22 (L)	(CL1-D, 2, 2-1)	30053	269..SMPL	----	-----	----	-
oradb	oradev2 (R)	(CL1-D, 2, 2)	30053	269..S-VOL	PAIR	-----	267	-
oradb1	oradev12 (R)	(CL1-D, 3, 2-0)	30053	271..S-VOL	PAIR	-----	269	-

- Designate a group name and confirm BC pair states from HOSTB, Instance 0:

```
# pairdisplay -g oradb1 -m cas
```

Group	PairVol (L/R)	(Port#, TID, LU-M)	Seq#	LDEV#..P/S	Status	Seq#	P-LDEV#	M
oradb1	oradev11 (L)	(CL1-D, 3, 1-0)	30053	270..S-VOL	PAIR	-----	268	-
oradb1	oradev11 (R)	(CL1-D, 2, 1-0)	30053	268..P-VOL	----	30053	270	-
oradb2	oradev21 (R)	(CL1-D, 2, 1-1)	30053	268..SMPL	PAIR	-----	----	-
oradb	oradev1 (R)	(CL1-D, 3, 1)	30053	268..S-VOL	PAIR	-----	266	-
oradb1	oradev12 (L)	(CL1-D, 3, 2-0)	30053	271..S-VOL	PAIR	-----	269	-
oradb1	oradev12 (R)	(CL1-D, 2, 2-0)	30053	269..P-VOL	----	30053	271	-
oradb2	oradev22 (R)	(CL1-D, 2, 2-1)	30053	269..SMPL	PAIR	-----	----	-
oradb	oradev2 (R)	(CL1-D, 3, 2)	30053	269..S-VOL	PAIR	-----	267	-

Two-host BC configuration

These two RM configuration files illustrate how to configure a two-host BC. Each host will run one instance of RM.

File 1

This is the RaidManager Configuration file for host blue.
It will manage the PVOLs in the Business Copy pairing.

```
HORCM_MON
#local host      local service  poll    timeout
blue             horcm0      1000    3000

HORCM_CMD
/dev/rdisk/c4t14d0

HORCM_DEV
#group  disk-name      interface  target    lun  mirror
Group1  disk_1_g1         CL1-A     2         0
Group1  disk_2_g1         CL1-A     2         1

HORCM_INST
#group  remote host      remote service name
Group1  yellow             horcm1
```

File 2

This is the Raid Manager Configuration file for host yellow.
It will manage the SVOLs in the Business Copy pairing.

```
HORCM_MON
#local host      local service  poll    timeout
yellow           horcm1        1000    3000

HORCM_CMD
/dev/rdisk/c10t14d0

HORCM_DEV
#group  disk-name      interface  target    lun  mirror
Group1  disk_1_g1         CL1-E     3         3
Group1  disk_2_g1         CL1-E     3         4

HORCM_INST
#group  remote host      remote service name
Group1  blue               horcm0
```

The RM configuration files show one RM group defined. The group, **Group1**, contains two disks. The comments note that system **blue** is defining the P-VOLs and system **yellow** is defining the S-VOLs. However, the P-VOL/S-VOL relationship is set when the **paircreate** command is issued. The set of disks that becomes the P-VOL or S-VOL depends on two conditions:

- the RM instance to which the command is issued
- the option specified in the **paircreate** command

The instance that the command is issued to becomes the local instance. If the option passed to the **paircreate** command is **-vl**, then the volumes defined in the local instance become the P-VOLs. If the option is **-vr**, the volumes defined in the remote instance become the P-VOLs.

Two BC mirror configuration

These two RM configuration files illustrate how to configure two BC mirrors of the same P-VOLs.

File 1

This is the Raid Manager Configuration file for host blue.
It will manage the PVOLs in the Business Copy pairing.

```
HORCM_MON
#local host      local service  poll    timeout
blue            horcm0      1000    3000

HORCM_CMD
/dev/rdisk/c4t14d0

HORCM_DEV
#group  disk-name      interface  target    lun  mirror
Group1-0  disk_1_g1-0    CL1-A      2         0    0
Group1-0  disk_2_g1-0    CL1-A      2         1    0
Group1-1  disk_1_g1-1    CL1-A      2         0    1
Group1-1  disk_2_g1-1    CL1-A      2         1    1

HORCM_INST
#group  remote host      remote service name
Group1-0  blue            horcm1
Group1-1  blue            horcm1
```

File 2

This is the Raid Manager Configuration file for host blue.
It will manage the SVOLs in the Business Copy pairing.

```
HORCM_MON
#local host      local service  poll    timeout
blue            horcm1      -1      3000

HORCM_CMD
/dev/rdisk/c4t14d0

HORCM_DEV
#group  disk-name      interface  target    lun
Group1-0  disk_1_g1-0    CL1-A      5         5
Group1-0  disk_2_g1-0    CL1-A      5         6
Group1-1  disk_1_g1-1    CL1-A      6         0
Group1-1  disk_2_g1-1    CL1-A      6         1

HORCM_INST
#group  remote host      remote service name
Group1-0  blue            horcm0
Group1-1  blue            horcm0
```

A one-host configuration differs from a two-host configuration as follows:

- The host names for the local and remote are the same.
- The poll value under the **HORCM_MON** section for the S-VOL configuration file is **-1**.

When creating more than one BC of the same P-VOL, the mirror unit column in the **HORCM_DEV** section must be filled in for the P-VOL configuration. Do not fill it in for the S-VOL configuration. If the mirror unit column is not filled in, the default value is 0.

Three-host BC configuration

These three RM configuration files illustrate how to configure a three-host BC. Each host will run one instance of RM.

File 1

This is the Raid Manager configuration file for host blue.
#It will manage the PVOLs in the Business Copy pairing.

```
HORCM_MON
#local host      local service  poll    timeout
blue             horcm0      1000    3000

HORCM_CMD
/dev/rdisk/c4t14d0

HORCM_DEV
#group  disk-name      interface  target    lun  mirror
Group1  disk_1_g1        CL1-A      2         0
Group1  disk_2_g1        CL1-A      2         1
Group2  disk_1_g2        CL1-A      3         0
Group2  disk_2_g2        CL1-A      4         0
Group2  disk_3_g2        CL1-A      4         1

HORCM_INST
#group  remote host      remote service name
Group1  yellow           horcm1
Group2  green            horcm0
```

File 2

This is the Raid Manager Configuration file for host yellow.
It will manage the SVOLs in the Business Copy pairing.

```
HORCM_MON
#local host      local service  poll    timeout
yellow           horcm1        -1      3000

HORCM_CMD
/dev/rdisk/c10t14d0

HORCM_DEV
#group  disk-name      interface  target    lun  mirror
Group1  disk_1_g1        CL1-E      3         3
Group1  disk_2_g1        CL1-E      3         4

HORCM_INST
#group  remote host      remote service name
Group1  blue             horcm0
```

File 3

This is the Raid Manager Configuration file for host green.
It will manage the SVOLs in the Business Copy pairing.

HORCM_MON

```
#local host    local service  poll    timeout
green          horcm0         -1      3000
```

HORCM_CMD

```
/dev/rdisk/c10t14d0
```

HORCM_DEV

#group	disk-name	interface	target	lun	mirror
Group2	disk_1_g2	CL1-F	3	3	
Group2	disk_2_g2	CL1-F	3	4	
Group2	disk_2_g2	CL1-F	3	5	

HORCM_INST

#group	remote host	remote service name
Group2	blue	horcm0

Device group configuration

This RM configuration file shows how to configure two device groups that belong to different unit IDs (disk arrays).

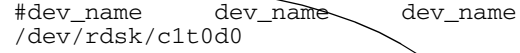
File 1

```
HORCM_MON
#ip_address      service      poll(10ms)    timeout(10ms)
HST1             horcm        1000          3000

HORCM_CMD
#unitID 0... (seq#30014)
#dev_name        dev_name      dev_name
/dev/rdsk/c0t0d0
#unitID 1... (seq#30015)
#dev_name        dev_name      dev_name
/dev/rdsk/c1t0d0

HORCM_DEV
#dev_group      dev_name      port#          TargetID      LU#           MU#
oradb           oradb1        CL1-A         3              0
oradb           oradb2        CL1-A         3              1
oralog         oralog1       CL1-A1        5              0
oralog         oralog2       CL1-A1        5              1
oralog         oralog3       CL1-A1        5              2

HORCM_INST
#dev_group      ip_address    service
oradb           HST2         horcm
oradb           HST3         horcm
oralog         HST3         horcm
```



B

HA Failover and failback

This appendix covers high availability (HA) failover and failback sequences.

Using RAID Manager in HA environments

When using HA software (such as MC/ServiceGuard or Cluster Extension XP), application packages can be transferred to the takeover host node at any time. If the application package transfer operation is performed in an environment where CA is used, you may need to switch the CA secondary volumes to primary volumes. The **horctakeover** command provides this function.

The **horctakeover** command provides macro functions to determine the data consistency of the S-VOL and to perform these takeover functions:

- takeover-switch
- swap-takeover
- SVOL-takeover
- PVOL-takeover

The **horctakeover** command is not available for BC.

HA control script state transitions

The table on page 311 lists volume states and state transitions resulting from the execution of either **pairvolchk** or **horctakeover** in HA control scripts from either Data Center 1 (DC1) or Data Center 2 (DC2).

When a HA failover/failback control script is activated, it will:

- retrieve the state of the CA paired volume accessible to the local host (via **pairvolchk -s**)
- retrieve the state of the remote side of the pair (via **pairvolchk -s -c**)
- select the proper action (for example, failover, failback, ask for operator intervention, etc.)

state No	Volume Attributes and Pair Status			Execution Results of pairvolchk and horctakeover from DC1(DC2)					
	DC1(DC2)		DC2(DC1)		pairvolchk -s (local volume)	pairvolchk -s -c (remote volume)	PAIR STATUS	Horctakeover result	
1	SMPL or SVOL-PSUS (SSWS)		SMPL		SMPL or SVOL-PSUS	SMPL	SMPL	EX_VOLCRE	
2			P-VOL	COPY		PVOL_XXX	XXX	Nop	
3				PAIR/PFUL					
4				PSUS					
4-1				PFUS					
5				PSUE					
6			PDUB						
8			S-VOL				EX_EVOLCE		EX_EVOLCE
9			Unknown				EX_ENORMT or EX_CMDIOE		(EX_ENORMT) (EX_CMDIOE)
10	P-VOL		SMPL		PVOL_XXX	SMPL	XXX	EX_VOLCRE	
11	P-VOL		P-VOL			EX_EVOLCE		EX_EVOLCE	
12	data or status && PSUE or PDUB Other		S-VOL			SVOL_YYY	XXX	PVOL-PSUE → 12 or PVOL-SMPL → 8 Nop	
13	data or status && PSUE or PDUB Other		Unknown			EX_ENORMT or EX_CMDIOE	XXX	PVOL-PSUE → 13 or PVOL-SMPL → 9 Nop	
14	S-VOL		SMPL			EX_EVOLCE		EX_EVOLCE	
15			P-VOL	COPY		PVOL_XXX	XXX	SVOL_E* → 4,5 SVOL_E*	
16				PAIR/PFUL				Swap → 12	
17				PSUS				SVOL_E → 4	
18				PFUS				SVOL → 4-1	
				PSUE PDUB	data			SVOL → 5,6	
			status				SVOL_E → 5,6		
never						SVOL_E → 5,6			
			async			SVOL → 5,6			
21	S-VOL			EX_EVOLCE		EX_EVOLCE			
22	COPY		Unknown		SVOL_YY	EX_ENORMT or EX_CMDIOE	YYY	SVOL_E * → 4,5 SVOL_E*	
23	PAIR/ PFUL	data				SVOL → 4			
		status				SVOL → 4			
		never				SVOL_E → 4			
		async				SVOL → 4			
24	PSUS					SVOL_E → 4			
	PFUS				SVOL → 4-1				
25	PSUE PDUB	data			SVOL → 5,6				
		status			SVOL_E → 5,6				
		never			SVOL_E → 5,6				
		async			SVOL → 5,6				

Table terms:

XXX	Pair status of P-VOL that was returned by the pairvolchk -s or pairvolchk -s -c command.
YYY	Pair status of S-VOL that was returned by the pairvolchk -s or pairvolchk -s -c command.
PAIR STATUS	Since the P-VOL controls status, PAIR STATUS is reported as PVOL_XXX (except when the P-VOLs status is Unknown).
PVOL-PSUE	PVOL-PSUE-takeover
PVOL-SMPL	PVOL-SMPL-takeover
Nop	Nop-takeover
Swap	swap-takeover

When the **horctakeover** command execution succeeds, the state transitions to that of the shown (→) number.

XP256 microcode 52-47-xx and under **XP512/48 microcode 10-00-xx and under**

With older firmware, a **horctakeover** used to result in a SMPL S-VOL, which necessitated a full copy at failback time. See [“Swap-takeover function”](#) .

SVOL	SVOL-SMPL takeover
SVOL_E	Execute SVOL-SMPL takeover and return EX_VOLCUR.
SVOL_E*	Execute SVOL-SMPL takeover and return EX_VOLCUR.

XP256 microcode 52-47-xx and over
XP512/48 microcode 10-00-xx and over
XP1024/XP128
XP12000

With newer firmware, a **horctakeover** results in a SSWS state S-VOL so that a delta copy is all that is required at failback. This functionality is known as “fast failback” and is accomplished via the **–swaps|p** option to **pairresync**.

SVOL	SVOL-SSUS takeover or swap-takeover In case of a host failure, this function will execute swap-takeover. In case of an ESCON/FC link or P-VOL site failure, this function will execute SVOL-SSUS-takeover.
SVOL_E	Execute SVOL-SSUS takeover and return EX_VOLCUR
SVOL_E*	Return EX_VOLCUR

When the **horctakeover** command execution succeeds, the state transitions (→) to that of the shown line number.

For instance, if the HA control script sees **svol_pair** at the local volume and **pvol_pair** at the remote volume (as in State 16), it performs a swap-takeover that results in a State 12 situation.

Failback after SVOL-SMPL takeover

This failover situation occurs, for instance, when:

- The original P-VOL status is unavailable
- The S-VOL is changed to SMPL and unable to failback.

The Host B (DC1) sequence illustrated by the following figures is required to change the SMPL volume to pvol_pair and make it suitable for failback.

From the Data Center 1 (DC1) side, the required steps are:


1. `pairsplit -S`
2. `paircreate -vl`
3. `pairevtwait` (wait for PAIR)

From the Data Center 2 (DC2) side, the required steps (not shown) would be:

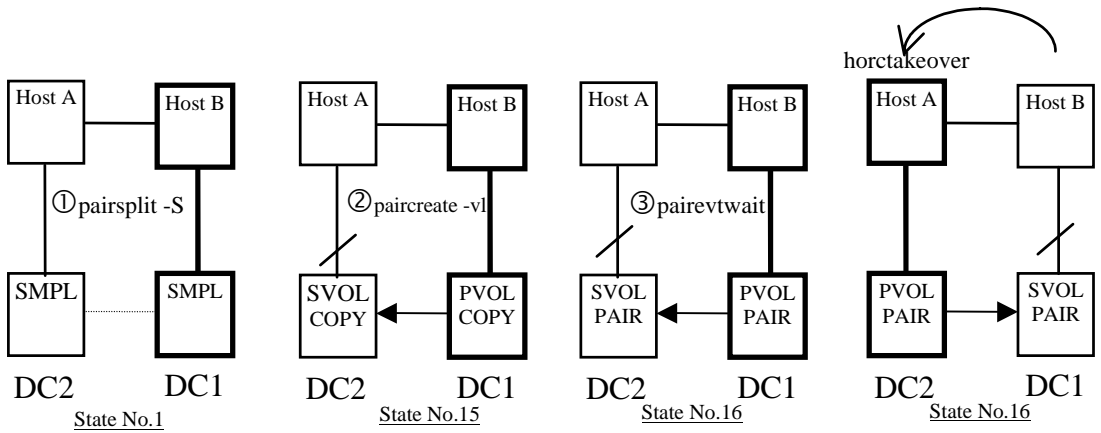
1. `pairsplit -S`
2. `paircreate -vr`
3. `pairevtwait` (wait for PAIR)

Refer to the state definitions in the table under the heading [“HA control script state transitions”](#) on page 312.


Initial state:




- When the DC2 volume becomes a svol_pair, it executes a swap-takeover to become a pvol_pair:



- If DC2 attempts a failback while the DC1 volume is still SMPL, it is a State 10 situation. The takeover operation returns an EX_VOLCRE (local/remote vol currency) error.

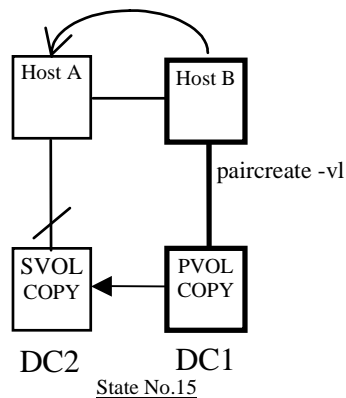


- If a takeover operation is attempted while both volumes are SMPL (State 1), an EX_VOLCRE error results. If **pairvolchk** is executed during a volume group split, it would likely return an EX_ENQVOL error, indicating that the statuses of the volumes in the group do not match.




- If a takeover is needed during State 15 (copy), the HA script could either run **pairevtwait** to wait for PAIR state, or prompt for system administrator intervention. If you choose to continue, an SVOL-takeover and an EX_VOLCUR error results.


An attempt to execute **pairvolchk** in the middle of a group **paircreate** returns an EX_ENQVOL error on the DC2 side.



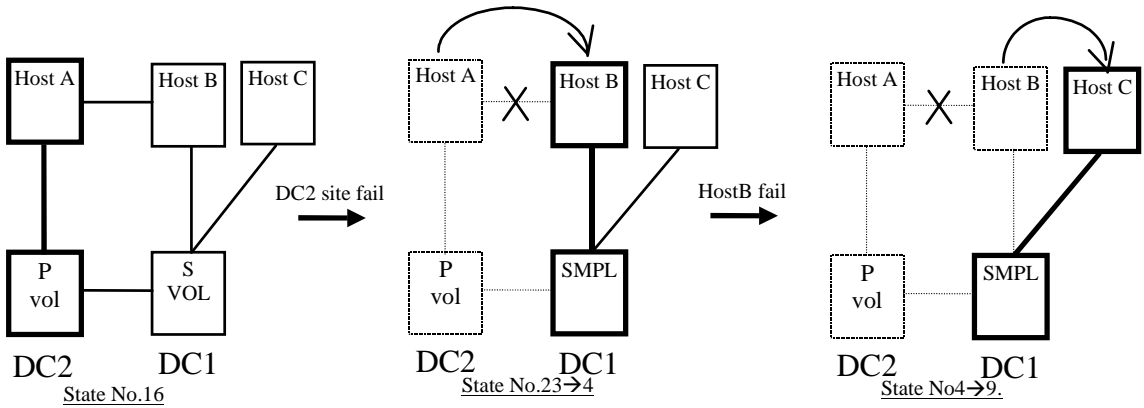
- An attempt to do a takeover prior to all group volumes reaching PAIR state (svol_copy) results in a SVOL-takeover and an EX_VOLCUR error.



- The HA script should prompt you for a decision before attempting a takeover in SVOL_PSUS (stale data) State 17, because it will result in a SVOL-takeover and an EX_VOLCUR error.



- The **horctakeover** command will fail with an EX_ENORMT error in the following nested failure case (State No. 4 → 9). Therefore, the HA Control Script should prompt you for a decision and not change the volume state on the DC1 side.

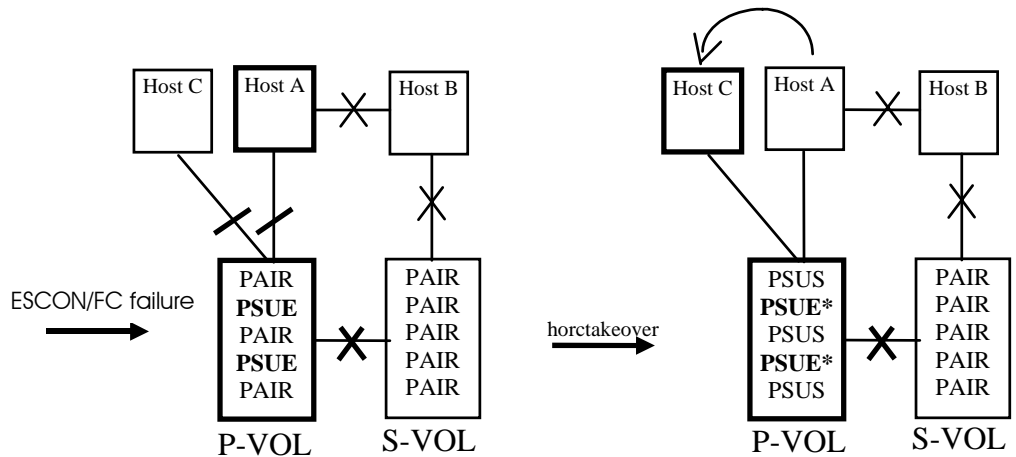


PVOL-PSUE takeover

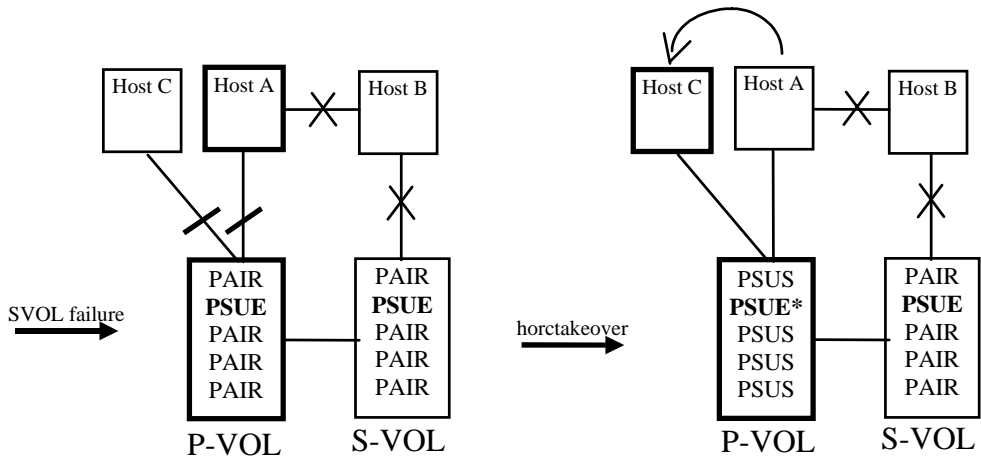
The **horctakeover** command executes a PVOL-PSUE-takeover when the primary volume cannot report status or refuses writes (for example, data fence).

- PSUE (or PDUB) and the **horctakeover** command returns a PVOL-PSUE-takeover value at **exit()**.
- A PVOL-PSUE-takeover forces the primary volume to the suspend state (PSUE or PDUB → PSUE*, PAIR → PSUS), which permits WRITES to all primary volumes of the group.

The following illustrates how volumes in the same volume group may be of different status. Only the volumes that were active at the time of link failure would immediately be PSUE.



- Even if connected to the ESCON/FC link, PVOL-PSUE-takeover changes only the active P-VOL/S-VOLs to suspend state.

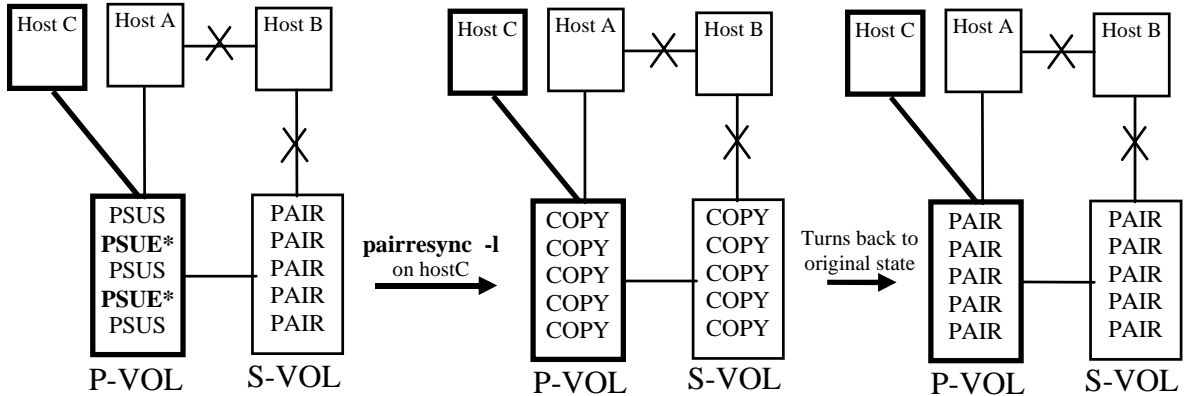


P-VOL group status

The result of the pvol_psue takeover is that PSUE and PSUS status is intermingled within the group. If any volumes are PSUE (or PDUB in the case of LUSE volumes), **pairvolchk** will give them precedence over PSUS as the returned group status.

Recovery after PVOL-PSUE-takeover

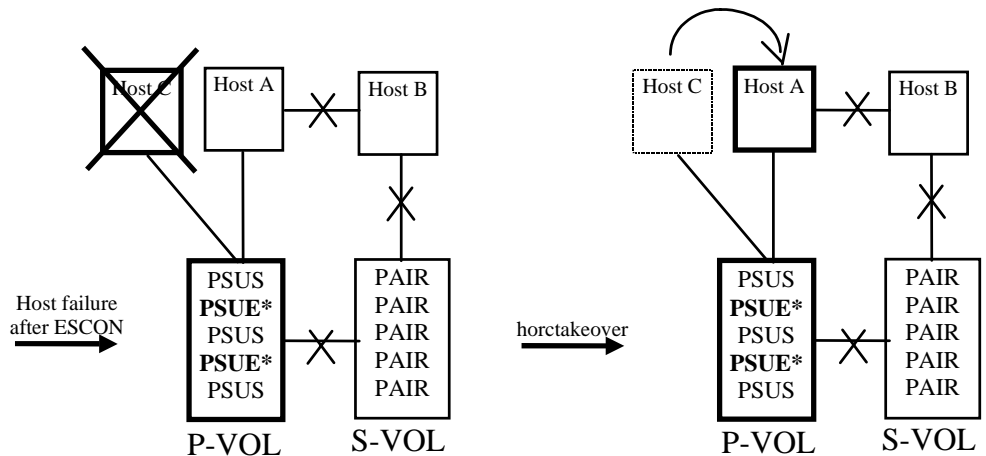
The special **PSUE*** state can be turned back to **PAIR** state by issuing the **pairresync** command (after the recovery of the ESCON/FC link) instead of the **horctakeover** command.



If the **pairresync** command fails because the ESCON/FC link is not yet restored, the **PSUE*** state is not changed.

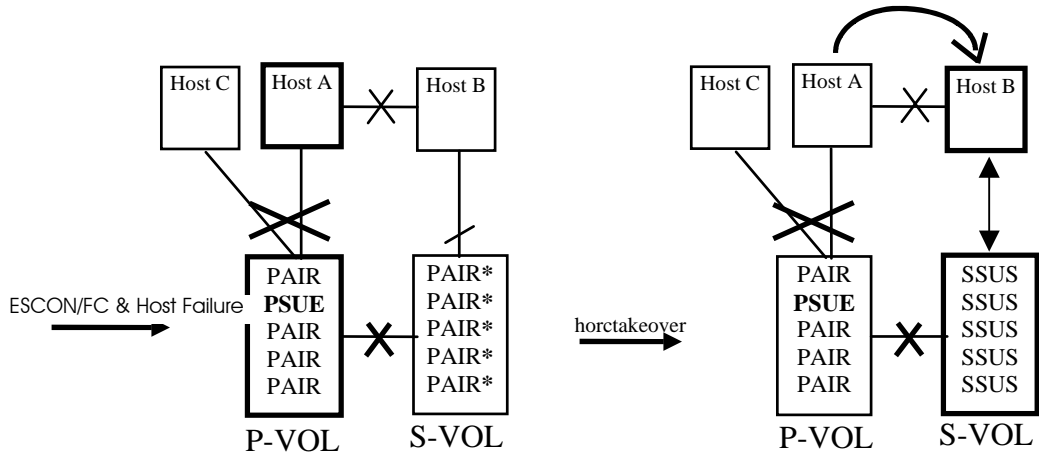
Another PVOL-Takeover Case

The **horctakeover** from Host A (after a Host C failure) performs a PVOL-PSUE-takeover in this nested (multiple) failure case.



SVOL-SSUS-takeover in the case of ESCON/FC link and host failure

An SVOL-Takeover executes an SVOL-SSUS-takeover so that S-VOL writing is enabled without going to SMPL state. An SVOL-SSUS-takeover changes the secondary volume to suspend (PAIR, PSUE → SSUS) state, which permits WRITE and delta data maintenance (BITMAP) for all secondary volumes of the group.



Terms:

PAIR* Equivalent to PAIR for sync-CA.
 Equivalent to PAIR → PSUE for async-CA.

SSUS Equivalent to SVOL_PSUS.

Group status after a SVOL-SSUS-takeover

After an SVOL-SSUS-takeover completes:

- The S-VOL status is displayed as SSUS by the **pairdisplay** command.
- The **pairvolchk** command returns the S-VOL status as SVOL_PSUS.

Also, this special state is displayed as SSWS by using the **-fc** option of the **pairdisplay** command. This special state (PVOL_PSUE and

SVOL_PSUS) between P-VOL and S-VOL may need to be handled by HA control scripts.

Async-CA specific behavior

Before the S-VOL is changed to SSUS, an SVOL-takeover will try to copy non-transmitted data, which remains in the FIFO queue (sidefile) of the P-VOL, to the S-VOL side.

In the case of an ESCON/FC link failure, this data synchronization operation may fail. Even so, the SVOL-takeover function performs the force split to SSUS, and enables usage of the secondary volume.

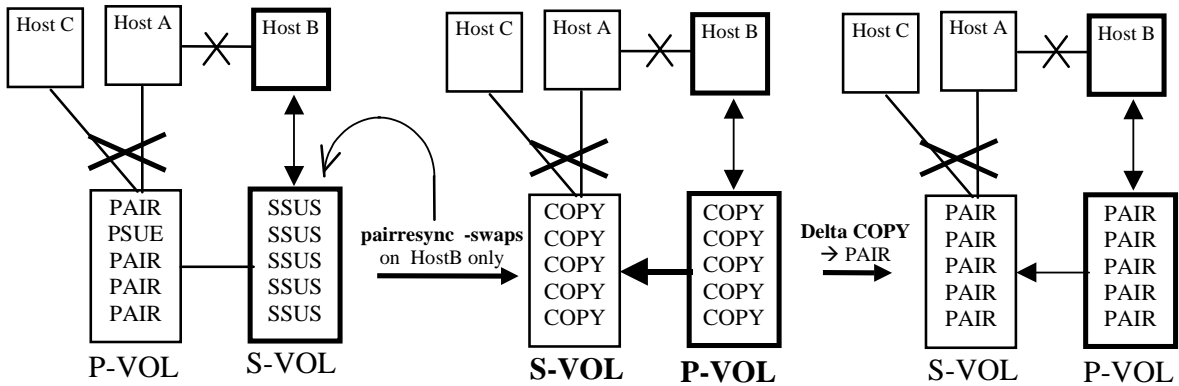
Non-transmitted data, which remains in the FIFO queue (sidefile) of the P-VOL, will be moved to a BITMAP to empty the FIFO queue, and the pair state will be set to PSUE.

Caution

*Ordering information regarding the non-transmitted data that is moved to a BITMAP will be lost. The data represented in the bitmap can be resynchronized (lost) as the new S-VOL by issuing **pairresync -swaps** for recovery from a SVOL-SSUS-takeover at the takeover site (Host B).*

Recovery from a SVOL-SSUS-takeover

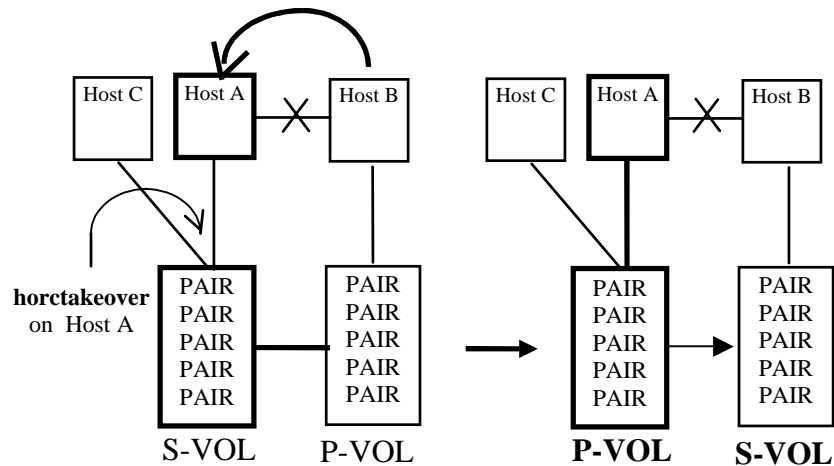
After the recovery of the ESCON/FC link, this special state (PVOL_PSUE and SVOL_PSUS) will be changed to COPY state. Thus, the original S-VOL becomes the NEW_PVOL and overwrites the NEW_SVOL, by issuing the **pairresync -swaps** command at the takeover site (Host B).



If the **pairresync -swaps** command fails because the ESCON/FC link is not yet restored, then the special state (PVOL_PSUE and SVOL_PSUS) is not changed.

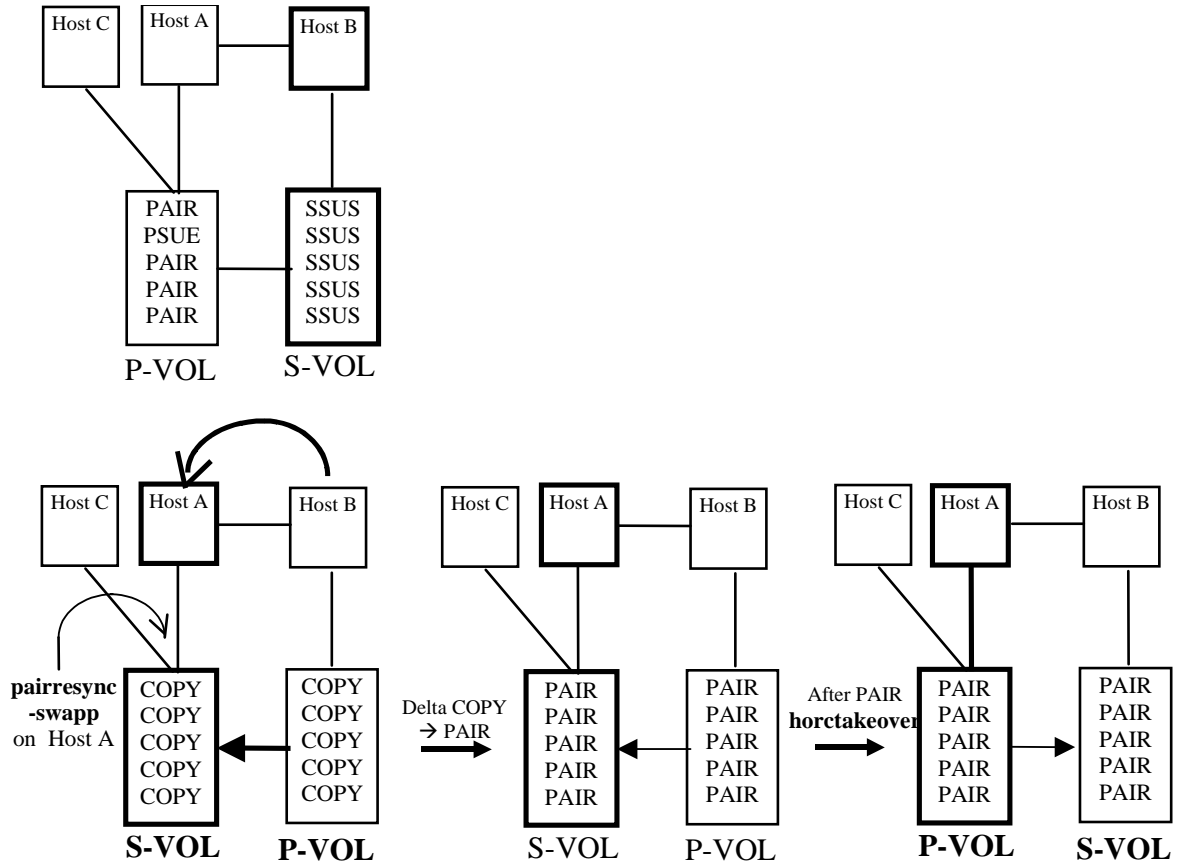
Failback after recovery on Host B

If you stop the application on Host B and restart the application on Host A after a recovery due to the execution of the **pairresync -swaps** command on Host B, then **horctakeover** does a swap-takeover, even though Host A cannot communicate with remote Host B.



Failback without recovery on Host B

The following procedure for recovery is necessary if, after host and ESCON/FC link recovery, you stop the application without executing the **pairresync -swaps** command on Host B and restart the application on the Host A. At that time, the **pairvolchk** command on Host A will return PVOL_PSUE and SVOL_PSUS as the state combination.



pairresync -swapp

This is used to swap the P-VOL/S-VOL designations during the P-VOL suspend state. The new S-VOL is based on the data of the new P-VOL. The target volume of the local host must have been the P-VOL, which was probably out of date.

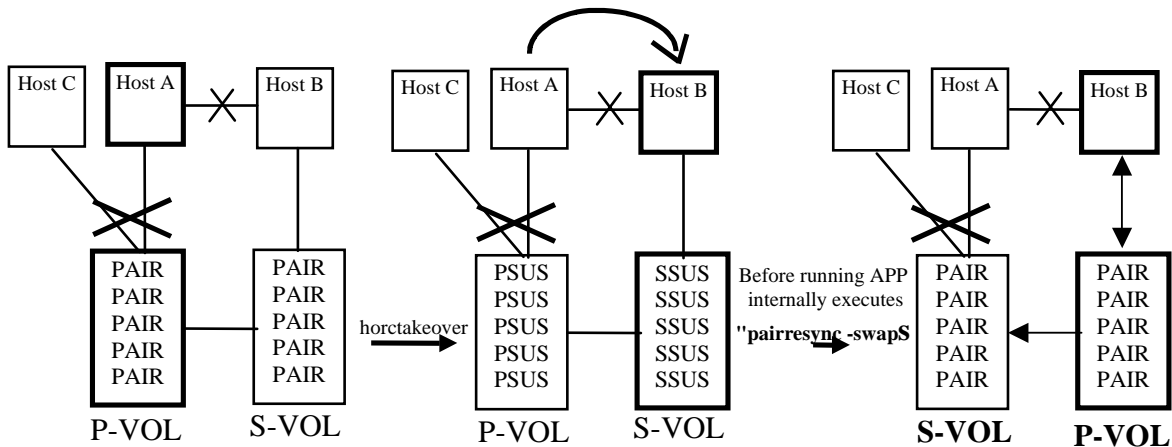
SVOL-takeover in the case of a host failure

After SVOL-takeover changes the S-VOL (only) to suspend (PAIR, PSUE → SSUS) state, the SVOL-takeover will automatically execute the **pairresync -swaps** command to copy data between the new P-VOL and the new S-VOL. The **horctakeover** command returns a swap-takeover.

Async-CA specific behavior

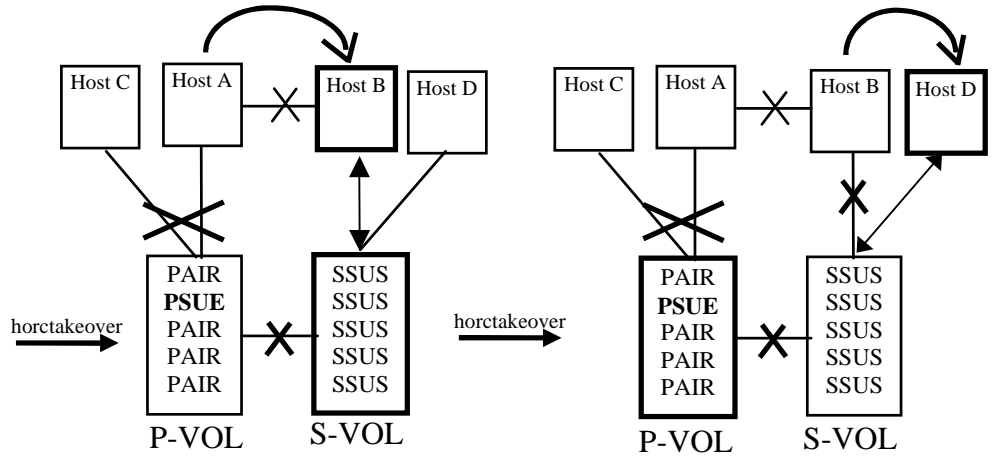
Before the S-VOL is changed to SSUS, the SVOL-takeover operation will copy non-transmitted data (which remains in the P-VOL sidefile) to the S-VOL. The SVOL-takeover operation waits to attempt to copy all P-VOL data to the S-VOL before a timeout (specified by the **-t timeout** option).

After all the P-VOL sidefile data has been successfully copied to the S-VOL, the SVOL-takeover operation splits the pairs and changes the S-VOL to SSUS state. The remainder of the operation is the same as for the non-asynchronous case.



Another case of SVOL-takeover

An SVOL-takeover from Host B to Host D will do nothing because the S-VOL was already in SSWS state.



S-VOL data consistency function

The consistency of the data within a pair is determined by the pair status and the fence level of the pair. The **paircurchk** command can be specified for each paired logical volume or each group. If this command is specified for a group, a data consistency check runs for all volumes in the group. Inconsistent volumes are printed in the message log and displayed.

See the table on the next page for details on this function. The terms used are defined in the list following the table.

Object volume				Currency SVOL_Takeover
Attribute	Status	Fence	paircurchk	
SMPL	—	—	Needs to be confirmed	—
P-VOL	—	—	Needs to be confirmed	—
S-VOL	COPY	data		Inconsistent
		status	Inconsistent (due to out-of-order copying)	
		never	Inconsistent	
		async	Inconsistent	
	PAIR	data	OK	OK
		status	OK	OK
		never	Must be analyzed	To be analyzed
		async	Must be analyzed	OK (Assumption)
	PFUL	async	To be analyzed	OK (Assumption)
	PSUS	data	suspect	suspect
		status	suspect	suspect
		never	suspect	suspect
		async	suspect	suspect
	PFUS	async	suspect	OK (Assumption)
	PSUE	data	OK	OK
		status	suspect	suspect
	PDUB	never	suspect	suspect
		async	suspect	OK (Assumption)

(continued)

Object volume				Currency SVOL_Takeover
Attribute	Status	Fence	paircurchk	
	SSWS	data	suspect	—
		status	suspect	—
		never	suspect	—
		async	suspect	—

Terms:

Inconsistent Data in the volume is inconsistent because it is being copied.

Suspect The primary volume data and secondary volume data are not consistent (the same).

Must be analyzed It cannot be determined from the status of the secondary volume whether data is consistent. It is “OK” if the status of the primary volume is PAIR. It is “suspect” if the status is PSUS or PSUE.

Needs to be confirmed
It is necessary to manually check the volume.

When the S-VOL Data Consistency function is used, **paircurchk** sets either of the following returned values in `exit()`, which allows users to check the execution results with a user program.

normal termination
0 (OK. Data is consistent.)

abnormal termination
Other than 0. (For the error cause and details, see the execution logs.)

Takeover-switch function

The takeover command, when activated manually or by a control script, checks the attributes of volumes on the local and remote disk array to determine the proper takeover action. The table below shows the takeover actions.

Local node (Takeover)		Remote node		Takeover action
Volume attribute	Fence and status	Volume attribute	P-VOL status	
SMPL	—	SMPL	—	NG
		P-VOL	—	Nop-Takeover**
		S-VOL	—	Volumes unconformable
		Unknown	—	NG
P-VOL (primary)	Fence == Data or status and Status == PSUE or PDUB	SMPL	—	NG
		P-VOL	—	Volumes unconformable
		S-VOL	—	PVOL-Takeover*
		Unknown Status For example: LAN down	—	PVOL-Takeover* *required to allow local writes
	Others	SMPL	—	NG
		P-VOL	—	Volumes unconformable
		S-VOL	—	Nop-Takeover**
		Unknown Status For example: LAN down	—	Nop-Takeover** **no action needed to allow local writes

(continued)

Local node (Takeover)		Remote node		Takeover action	
Volume attribute	Fence and status	Volume attribute	P-VOL status		
S-VOL (secondary)	Status == SSWS (After SVOL_SSUS- takeover)	Don't care	—	Nop-Takeover**	
	Others	SMPL	—	Volumes unconformable	
		P-VOL	PAIR or PFUL		Swap-Takeover*
			Others		SVOL-Takeover*
		S-VOL	—		Volumes unconformable
		Unknown	—		SVOL-Takeover*

Terms:

nop-takeover No operation is done, though the takeover command is accepted.

Volumes unconformable
A pair of volumes are not conformable to each other. The takeover command terminates abnormally.

NG The “no good” takeover command is rejected and the operation terminates abnormally.

PVOL-takeover Executed from the P-VOL side. Gives the P-VOL read/write capability even if the S-VOL is unavailable with a fence level of data or status.

SVOL-takeover Executed from the S-VOL side. Attempts to swap the P/S designations. If unable to swap the P/S designations, changes the SVOL to SVOL-SSUS mode. If unable to change the SVOL to SVOL-SSUS mode, changes the SVOL to SMPL mode to allow writes to the volume.

swap-takeover Swaps the primary and secondary volume designations.

Unknown	The attribute of the remote node is unknown. This means that the remote node system has failed or cannot communicate.
---------	---

Swap-takeover function

It is possible to swap the designations of the primary and secondary volumes when the P-VOL of the remote disk array is in the PAIR or PFUL (async-CA and over HWM) state and the mirror consistency of S-VOL data has been assured.

The takeover command carries out the commands internally to swap the designations of the primary and secondary volumes. You can specify swapping at the granularity of volume pair, CT group, or volume group.

Swap-takeover works differently according to microcode version:

XP256 microcode 52-47-xx and under

XP512/48 microcode 10-00-xx and under

1. The command splits the pair and puts each volume in the SMPL state.

If this step fails, the swap-takeover function is disabled and the **SVOL-takeover** command executes.

2. The local volumes of the takeover node are paired in “No Copy” mode and switched to be the primary volume.

If this step fails, step 1 repeats to cancel the operation of step 2, and the SVOL-takeover function is then executed.

If step 1 fails again, the swap-takeover fails.

XP256 microcode 52-47-xx and over
XP512/48 microcode 10-00-xx and over
XP1024/128
XP12000

The swap-takeover function no longer uses “Simplex” and “No Copy” mode for swapping. This assures greater mirror consistency. Moreover, it is included as a function of SVOL-takeover.

1. The command orders a “suspend for swapping” (SSWS) for the local volume (S-VOL).

If this step fails, the swap-takeover function is disabled and returns an error.

2. The command orders a “resync for swapping” to switch to the primary volume. The local volume (S-VOL) is swapped as the NEW_PVOL. The NEW_SVOL is resynchronized based on the NEW_PVOL.

If the remote host is known, the command will use the value of P-VOL specified at **paircreate** time for the number of simultaneous copy tracks. If the remote host is unknown, the command will use a default of **3** simultaneous copy tracks for “resync for swapping.”

If this step fails, the swap-takeover function is returned at SVOL-SSUS-takeover. The local volume (S-VOL) is maintained in the SSUS (PSUS) state which permits WRITE and maintenance of delta data (BITMAP) for the secondary volume. This special state is also displayed as the SSWS state, using the **-fc** option of the **pairdisplay** command.

Async-CA Extended Specific Behavior

XP256 microcode 52-47-xx and under
XP512/48 microcode 10-00-xx and under

1. The P-VOL side RM issues a **pairsplit** command to the P-VOL side disk array.
2. Non-transmitted data that remains in the FIFO queue (sidefile) of the P-VOL is copied to the S-VOL side.

3. The swap operation is performed.

The swap operation must copy non-transmitted P-VOL data within the timeout value specified by the `-t timeout` option.

4. The swap command returns after the synchronization between the P-VOL and S-VOL.

XP256 microcode 52-47-xx and over
XP512/48 microcode 10-00-xx and over
XP1024/128
XP12000

1. The S-VOL side RM issues a “suspend for swapping” to the S-VOL side disk array.
2. Non-transmitted data that remains in the FIFO queue (sidefile) of the P-VOL is copied to the S-VOL side.
3. A “resync for swapping” operation is performed.

The swap operation must copy non-transmitted P-VOL data within the timeout value specified by the `-t timeout` option.

SVOL-takeover function

This function enables the takeover node to have exclusive access to the S-VOL volume in SSUS (PSUS) state (reading and writing are enabled), except in COPY state, on the assumption that the remote node, controlling the P-VOL, is unavailable or unreadable.

The data consistency of the secondary volume is judged by its pair status and fence level. If the data consistency check fails, the SVOL-takeover function fails.

You can specify SVOL-takeover at the granularity of a paired logical volume or group.

If this check proves that the data is consistent, this function runs to switch to the primary volume using a Resync for Swapping. If this switch succeeds, this function returns with swap-takeover. Otherwise, this function

returns SVOL-SSUS-takeover as the return value of **horctakeover** command.

If there is a Host failure, this function returns as swap-takeover.

If an ESCON/FC link or P-VOL site failure occurs, this function returns as SVOL-SSUS-takeover.

If SVOL-takeover is specified for a group, the data consistency check executes for all volumes in the group. Inconsistent volumes are displayed in the execution log file.

Example

Group	Pair vol	Port	targ#	lun#	LDEV#...Volstat	Status	Fence	To be..
oradb1	/dev/dsk/hd001	CL1-A	1	5	145...S-VOL	PAIR	NEVER	Analyzed
oradb1	/dev/dsk/hd002	CL1-A	1	6	146...S-VOL	PSUS	STATUS	Suspected

Async-CA extended specific behavior

The S-VOL side RM issues a suspend for swapping operation to the S-VOL side of the disk array. The non-transmitted data of the primary volume copies to the S-VOL side and a resync for swapping operation runs after the copy process.

If there is a host failure, this data-synchronize operation runs and the SVOL-takeover function returns as swap-takeover by running a resync for swapping operation.

If an ESCON/FC link or P-VOL site failure occurs, this data-synchronize operation could fail. Even so, the SVOL-takeover function performs a suspend for swapping operation and enables the secondary volume to be used. This function returns as SVOL-SSUS-takeover. The non-transmitted data of the primary volume is not transmitted completely when SVOL-takeover returns SVOL-SSUS-takeover.

The SVOL-takeover operation is required to copy non-transmitted P-VOL data within a given timeout value specified by the **-t timeout** option.

If the timeout occurs before S-VOL takeover has completed all S-VOL changes to the SSWS state, the **horctakeover** command fails with EX_EWSTOT.

PVOL-takeover function

The PVOL-takeover function terminates the PAIR state of a pair or group. The takeover node is given unrestricted and exclusive access to the primary volume (reading and writing are enabled), on the assumption that the remote node (controlling the S-VOL) is unavailable or unreachable.

The PVOL-takeover function has two roles:

- PVOL-PSUE-takeover puts the P-VOL into PSUE state, which permits “WRITE” access to all primary volumes of that group.
- PVOL-SMPL-takeover puts the P-VOL into SMPL state.

PVOL-takeover first attempts to use PVOL-PSUE-takeover. If PVOL-PSUE-takeover fails, then PVOL-SMPL-takeover is executed.

You can specify PVOL-takeover with a granularity of logical volume or group.

P-VOLs (primary volumes) in DATA fence will not accept write commands after ESCON/FC link or remote array failures. PVOL-takeover can be used on these P-VOLs to allow the application to update the P-VOL if you choose to do so. However, none of those updates will be replicated or mirrored to the remote S-VOL.

Async-CA specific behavior

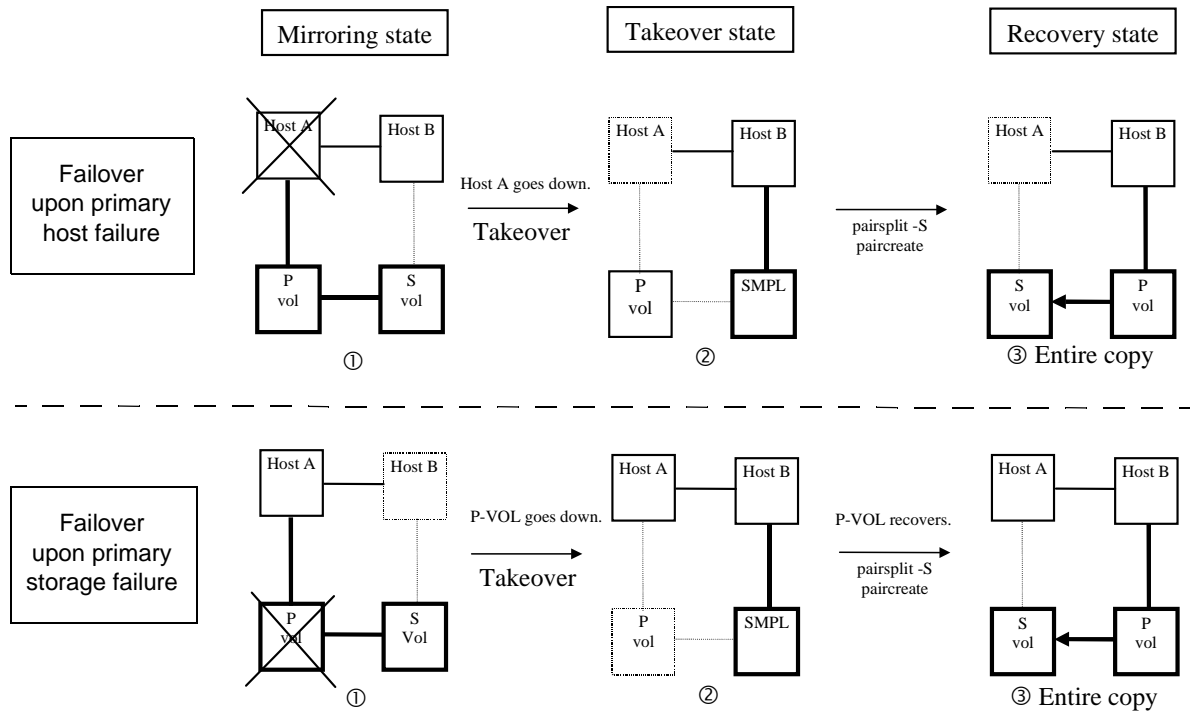
PVOL-takeover is not executed. It becomes a nop-takeover.

Recovery procedures of HA system configuration

After installing CA, the system administrator should conduct operation tests on the assumption that system failures may occur. In normal operation, service personnel obtain failure cause information from the SVP. However, the CA commands may also give error information.

XP256 microcode 52-47-xx and under
XP512/48 microcode 10-00-xx and under

The following figure shows a diagram of system failure and recovery.

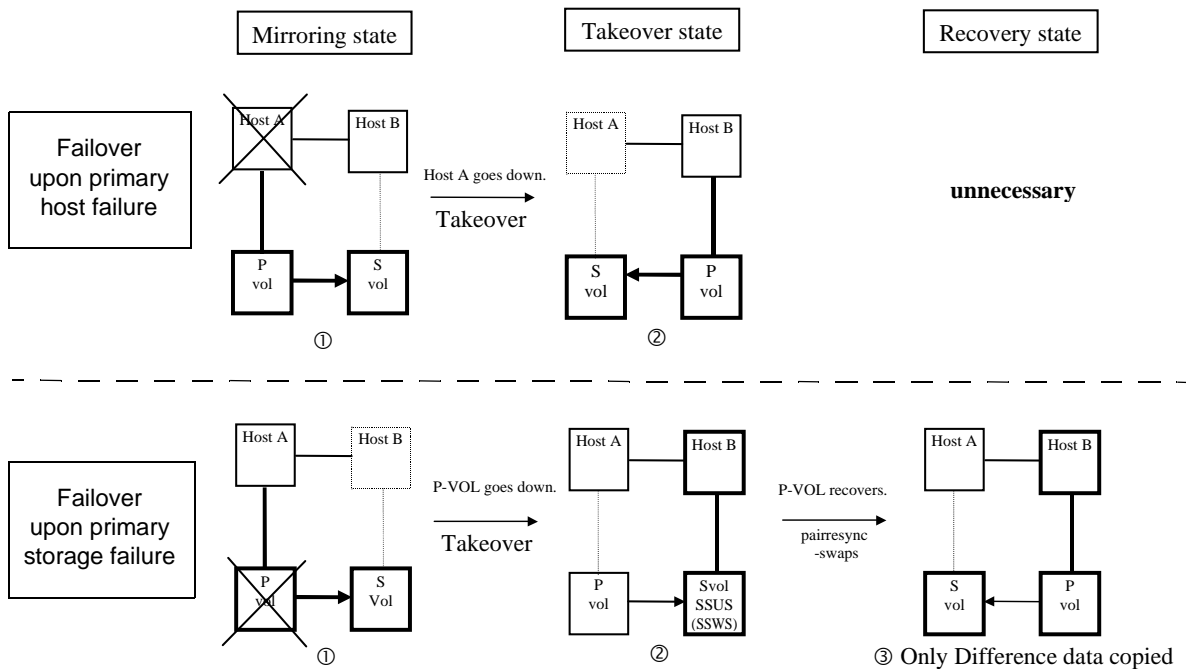


Scenario

1. A failure occurs in host A or in the P-VOL.

- Host B detects the failure in host A and issues the takeover command to make the S-VOL usable. If the S-VOL can continue processing, host B takes over from host A and continues processing.
- While host B is processing, the P-VOL and S-VOL can be swapped using full copy (**pairsplit -S, paircreate -vl**) and the data updated by host B is fed back to the new S-VOL, host A.
- When host A recovers from the failure, host A takes over processing from host B through the **horctakeover swap-takeover** command.

XP256 microcode 52-47-xx and over
XP512/48 microcode 10-00-xx and over
XP1024/128
XP12000



Scenario

1. A failure occurs in host A or in the P-VOL.
2. Host B detects the failure in host A and issues the takeover command to make the S-VOL usable. Host B takes over from host A and continues processing.

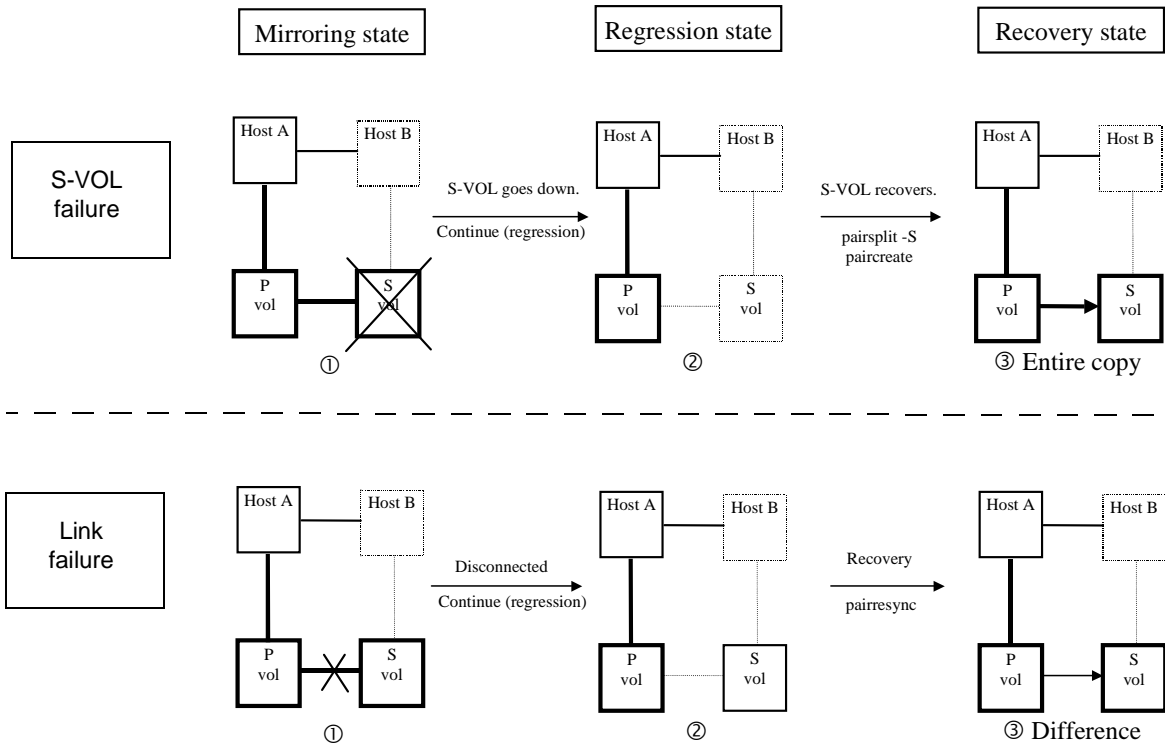
In the case of a host A failure, a takeover command executes a **swap-takeover**.

In the case of a P-VOL failure, a takeover command executes a **S-VOL-SSUS-takeover**.

3. While host B is processing, the P-VOL and S-VOL are swapped using **pairresync –swaps** and the delta data (BITMAP) updated by host B is fed back to host A.
4. When host A recovers from the failure, host A takes over processing from host B through the **horctakeover swap-takeover** command.

Regression and recovery of CA

The figure below shows a diagram of regression and recovery where **horctakeover** is not needed.



Scenario

1. The P-VOL detects a failure in the S-VOL or the link and suspends mirroring. (It depends on the fence level whether host A continues processing or host B takes over processing from host A.)
2. The P-VOL changes its paired volume status to PSUE and keeps track of data changes in a difference bitmap. The CA manager detects the status change and outputs a message to syslog. If a host A user has initiated a monitoring command, a message is displayed on the client's screen.

3. The S-VOL or the link recovers from the failure. Host A issues the **pairsplit –S**, **paircreate –vl**, or **pairresync** command to update the P-VOL data by copying all data, or copying differential data only. The updated P-VOL is fed back to the S-VOL.

CA recovery procedures

Follow these steps to recover CA operations:

1. If an error occurs in writing paired volumes (for example, pair suspension), the server software using the volumes detects the error depending on the fence level of the paired volume.
2. Issue **pairedisplay** to the paired volume or group to get status information.
3. If necessary, issue the **horctakeover** command to recover P-VOL write access if the secondary volume fails and the primary is fenced (write inhibited).
4. If the primary volume fails, split or suspend the paired volume and use the secondary volume as the substitute volume.
5. Find out the reason why the pair was split. Repair or recover the failure and resynchronize your pairs immediately.

Abnormal termination

A CA command can abnormally terminate for many reasons, for example:

- The remote server is down.
- A local server failure.
- A disk array failure.
- The disappearance of the RM instance.

Check the system log file and RM log file to identify the cause.

If a command terminates abnormally because the remote server fails, recover the remote server, and then reissue the command. If the RM instance has disappeared, reactivate the RM instance. If you find failures

for which you can take no action, check the files in the log directory and contact HP.

Failure to activate the RAID Manager instance

The failure to activate RM on a new system can be caused by an incorrect environment setting and/or configuration file definition. Check the activation log file and take any necessary actions.

Fibre Channel addressing

This appendix provides Fibre Channel conversion tables for these operating systems:

- HP-UX
- Sun Solaris
- Microsoft Windows NT
- Microsoft Windows 2000
- Microsoft Windows 2003
- OpenVMS

Fibre Channel address conversions

RM converts the Fibre Channel physical address to a target ID using conversion tables presented on the following pages.

Fibre Channel TID

Type of Port	HP-UX and Other Windows NT/2000 (HP Fibre)		Sun Solaris		Windows NT/2000 (Emulex)	
	TID	LUN	TID	LUN	TID	LUN
Fibre	0 to 63	0 to 7	0 to 125	0 to 511	0 to 31	0 to 511
SCSI	0 to 15	0 to 7	0 to 15	0 to 7	0 to 15	0 to 7

HP-UX Fibre Channel address conversion

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		

Sun Solaris Fibre Channel address conversion

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		

Windows NT/2000 Fibre Channel address conversion (QLogic or Emulex driver)

PhId1(C1)				PhId2(C2)				PhId3(C3)				PhId4(C4)				PhId5(C5)			
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
		27	15			56	15			98	15			CC	15				
3C	30	26	14	72	30	55	14	B1	30	97	14	E4	30	CB	14				
3A	29	25	13	71	29	54	13	AE	29	90	13	E2	29	CA	13				
39	28	23	12	6E	28	53	12	AD	28	8F	12	E1	28	C9	12				
36	27	1F	11	6D	27	52	11	AC	27	88	11	E0	27	C7	11				
35	26	1E	10	6C	26	51	10	AB	26	84	10	DC	26	C6	10				
34	25	1D	9	6B	25	4E	9	AA	25	82	9	DA	25	C5	9				
33	24	1B	8	6A	24	4D	8	A9	24	81	8	D9	24	C3	8				
32	23	18	7	69	23	4C	7	A7	23	80	7	D6	23	BC	7				
31	22	17	6	67	22	4B	6	A6	22	7C	6	D5	22	BA	6				
2E	21	10	5	66	21	4A	5	A5	21	7A	5	D4	21	B9	5				
2D	20	0F	4	65	20	49	4	A3	20	79	4	D3	20	B6	4				
2C	19	08	3	63	19	47	3	9F	19	76	3	D2	19	B5	3				
2B	18	04	2	5C	18	46	2	9E	18	75	2	D1	18	B4	2				
2A	17	02	1	5A	17	45	1	9D	17	74	1	CE	17	B3	1			EF	1
29	16	01	1	59	16	43	0	9B	16	73	0	CD	16	B2	0			E8	0

The table above is based on the conversion method of the QLogic or Emulex driver. If a different driver is used for the Fibre Channel adapter on the server, the target ID displayed on Windows NT/2000 may be different from the target ID shown in the table. In this case, the target ID on the RM configuration file is required to describe the target ID that is indicated by the **raidscan** command.

D

STDIN file formats

This appendix provides the format specifications for the STDIN or device special files.

The STDIN or device special files are specified in the following formats:

MPE/iX	/dev/. . .
HP-UX	/dev/rdisk/*
Solaris	/dev/rdisk/*s2 or c*s2,
Linux	/dev/sd... or /dev/rd...
MPE/iX	/dev/..., "LDEV-"
AIX	/dev/rhdisk* or /dev/hdisk* or hdisk*
Digital or Tru64	/dev/rz*c or /dev/rdisk/dsk*c
DYNIX	/dev/rdisk/sd* or sd*
	for only unpartitioned raw device
Windows NT	hdX-Y, \$LETALL, \$Phys, D:\DskX\pY, \DskX\pY
Windows2000/2003	hdX-Y,\$LETALL,\$Volume, \$Phys,D:\Vol(Dms,Dmt,Dmr)X\DskY, \Vol(Dms,Dmt,Dmr)X\DskY
OpenVMS	\$1\$* or DK* or DG* or GK*

Porting notice for MPE/iX

This appendix describes operating system requirements, restrictions, and known issues for MPE/iX.

Porting notice for MPE/iX

Introduction

MPE/iX does not fully support POSIX like UNIX. Therefore, RAID Manager has some restrictions in MPE/iX. The system calls (wait3(), gettimeofday()...) that are not supported on MPE/iX are implemented in LIB BSD; however, RM has to avoid using LIB BSD due to its availability as free software. These functions are, therefore, implemented within RM. RM has accomplished porting within standard POSIX for MPE/iX only.

Restrictions in the current release

RAID Manager has the following restrictions in porting to MPEiX.

Network function

The Bind() system call of MPE/iX POSIX cannot specify the Ip_address of its own host, so it supports only 'INADDR_ANY'. Therefore, RAID Manager needs to use 'NONE' like the following on the horcm.conf. Also a port number over 1024 must be specified in '/etc/services'.

HORCM_MON

#ip_address	service	poll (10ms)	timeout (10ms)
NONE	horcm	1000	3000

Syslog function

RAID Manager does not support the syslog function of MPE/iX due to free software availability. Also, the syslog daemon (syslogd) does not execute at normal startup on MPE/iX. As an alternative, the HORCM daemon uses a logging file.

HORCM daemon startup

HORCM can start as a daemon process from a UNIX Shell. But in the case of MPE/iX, if a parent process exits, then any child process also dies at the same time. In other words, it looks like MPE/iX POSIX cannot launch a daemon process from a POSIX Shell. Therefore, horcmstart.sh has been changed to wait until HORCM has exited after startup of the horcmgr. According to the rules for MPE/iX, horcmstart.sh is run as a MPE JOB. The following is an example of a JOB control file named JRAIDMR1 (HORCMINST=1).

```
!job jraidmr1, manager.sys;pri=cs
!setvar TZ "PST8PDT"
!xeq sh.hpbin.sys '/HORCM/usr/bin/horcmstart.sh 1'
!eoj
```

When you execute this JOB in the background by using the **STREAM** command of MPE/iX, you will have the HORCM daemon running in the background. You will be able to verify that the HORCM daemon is running as a JOB by using the **SHOWJOB** command.

```
shell/iX> callci STREAM JRAIDMR1
#J15
shell/iX> callci SHOWJOB
```

JOBNUM	STATE	IPRI	JIN	JLIST	INTRODUCED	JOB NAME
#J14	EXEC		10S	LP	WED 9:02P	JRAIDMR0,MANAGER.SYS
#J15	EXEC		10S	LP	WED 9:02P	JRAIDMR1,MANAGER.SYS
#S28	EXEC	QUIET	9	9	WED 9:10P	MANAGER.SYS

Command device

Because MPE/iX POSIX does not provide raw I/O such as UNIX, RAID manager has used the SCSI pass-thru driver to access the command device on XP256/512, and is also using the normal read/write SCSI commands for some RM control operations.

You need to confirm that MPE/iX has installed the patch MPEKXU3 before using the SCSI pass-thru driver

Installing

Since MPE/iX POSIX is unable to execute **cpio** to extract a file, the RM product is provided as a tar file.

For further information about installing RAID Manager on MPE/iX systems, see [“Installing RAID Manager on MPE/iX systems” on page 33](#).

Uninstalling

The **RMuninst** (`rm -rf /$instdir/HORCM`) command cannot remove the directory (`/HORCM/log*/curlog` only) while the HORCM is running. → For more details, see the section “Cannot remove directories using the “`rm -rf /users/HORCM`” command on page 358.

The only way to remove the log directory for the **RMuninst** (`rm -rf /$instdir/HORCM`) command is to shut down and reboot the MPE system.

Use the **RMuninst** (`rm -rf /$instdir/HORCM`) command after the MPE/iX system has been shut down and rebooted.

-zx option for RAID Manager commands

The '-zx option' for RAID Manager commands uses the `select()` function to wait for an event from STDIN, but the MPE/iX POSIX `select()` function does not support that, and `select()` for terminal (STDIN) is unable to echo back the terminal input.

Therefore the '-zx option' for RAID Manager commands will not be supported, and it will be deleted as a displayed option.

MPE socket hang

One problem is that two or more packets are queued on the MPE socket, and then the packets remain on the socket indefinitely unless a shutdown of HORCM occurs due to `select()` is not woke-up.

As a result, the command using the remote host fails with `EX_ENORMT` on multiple commands.

This problem is resolved by **RM010904(3)**, which supports the traffic control method for MPE socket.

The traffic control method is to limit sending the packets for multiple commands at the same time, and over-packets are queued (FIFO) to wait until sending the next packets. The queued packets are sent after a reply is received for the sent message. This method controls the amount of packets that are sent to the remote host at the same time. The amount of packets are controlled by the **HORCMTRFCTL** environment variable for HORCM. The default for the amount of packets for MPE is one (HORCMTRFCTL=1).

HORCMTRFCTL is effective for all (other) platforms, but the default is not controlled (HORCMTRFCTL=0).

In **RM 1.09.02** or earlier, the following is needed to avoid this problem.

- Isolates RM instances with groups performing at the same time. This will isolate the message on the socket. A maximum number of RM instances are 32 on one command device. (In the case of NO RMLIB).
- Uses '-l' option (pairsplit -l, pairresync -l, pairevtwait -l) on PVOL instance. This will not use the remote RM instance except paircreate.

Known issues and concerns

MPE panic with the “rm” command

The MPE POSIX layer is unable to execute the **rm** command when the directory you are trying to remove is pointed to by a symbolic link. In /HORCM there is a symbolic link to /users/HORCM. While trying to remove the directory under this symbolic link (two levels down), it gets a nil pointer causing a system failure. This is entirely attributed to a POSIX bug and HP realizes that it is a critical bug that needs to be fixed soon; it has been submitted as an urgent request for a fix.

Display of the “dstat” command

DSTAT can only display up to 12 characters for the product-id.

```
shell/iX> callci dstat
```

LDEV-TYPE	STATUS	VOLUME	VOLUME SET	GEN
99-OPEN-3-CVS	UNKNOWN			
100-OPEN-3-CVS	MASTER	MEMBER100	PVOL100-0	
101-OPEN-3-CVS	MASTER	MEMBER101	PVOL101-0	
102-OPEN-3-CVS	MASTER	MEMBER102	PVOL102-0	
103-OPEN-3-CVS-C	MASTER	MEMBER103	PVOL103-0	

Regarding “multiple capability” of the SCSI path-thru driver

When other commands are executed via the SCSI path-thru driver, HORCM is blocked until the other commands have completed. If RAID Manager (the HORCM daemon) is blocked while sending raw I/O to the command device, then HORCM cannot work for another RM command.

Cannot remove directories using the “rm -rf /users/HORCM” command

You can change the “horcmstart.sh” script to avoid MPE panic with the **rm** command; however, there may be a new problem in that **rm -rf** cannot remove the log directories. The following are the results of **rm -rf /users/HORCM**:

```
rm -rf /users/HORCM

rm: cannot remove directory "/users/HORCM/log0/curlog":
Permission denied
rm: cannot remove directory "/users/HORCM/log0/tmplog":
Permission denied
rm: cannot remove directory "/users/HORCM/log1/curlog":
Permission denied
rm: cannot remove directory "/users/HORCM/log1/tmplog":
Permission denied
```

Here, the **rm** command is saying “Permission denied” but if the cause is really “Permission denied”, why does it remove the directory and files under the ‘/users/HORCM/log*/curlog’ directory?

MPE/iX POSIX commands can never remove these directories.

You cannot remove the ‘/tmp/curlog’ directory even if you use the **mv /users/HORCM/log*/curlog /tmp** command.

MPE/iX startup procedures

Make a JOB control file

The following is an example of JOB control file named JRAIDMR0 (HORCMINST=0).

```
!job jraidmr0, manager.sys;pri=cs
!setvar TZ "PST8PDT"
!xeq sh.hpbin.sys '/HORCM/usr/bin/horcmstart.sh 0'
!eoj
```

Make the device special files and check the LDEV configuration

You are able to use **-inst** option that is used to make a special file as “/dev/ldev*” from “LDEV-” of the **dstat** command for the SCSI pass-thru driver.

```
shell/iX> callci dstat | ./ingraid -inst -CLI
```

DEVICE_FILE	PORT	SERIAL	LDEV	CTGC/B/12	SSID	R:Group	PRODUCT_ID
ldev100	CL1-L	35013	17	-	s/s/ss	0004 5:01-01	OPEN-3
ldev101	CL1-L	35013	18	-	s/s/ss	0004 5:01-01	OPEN-3
ldev102	CL1-L	35013	19	-	s/s/ss	0004 5:01-01	OPEN-3
ldev103	CL1-L	35013	35	-	-	-	OPEN-3-CM

Note: LDEV user here refers to the MPE/iX term.

Describe the command device on /etc/horcm*.conf

```
HORCM_MON
#ip_address      service      poll(10ms)  timeout(10ms)
NONE            horcm        1000        3000

HORCM_CMD
#dev_name        dev_name     dev_name
/dev/ldev103
```

```

HORCM_DEV
#dev_group      dev_name      port#      TargetID    LU#      MU#

HORCM_INST
#dev_group      ip_address    service

```

You will have to start HORCM without a description for HORCM_DEV and HORCM_INST **because the target ID & LUN are Unknown.**

You will be able to know about mapping a physical device with a logical device (ldev of MPE/iX term) by using **raidscan -find**.

Execute an "horcmstart.sh 0" as a JOB

```

shell/iX> callci STREAM JRAIDMR0
#J14
shell/iX> callci SHOWJOB

```

JOBNUM	STATE	IPRIJIN	JLIST	INTRODUCED	JOB NAME
#J14	EXEC	10S LP	WED 9:02P	JRAIDMR0.MANAGER.SYS	
#S28	EXEC QUIET	9 9	WED 9:10P	MANAGER.SYS	
.					
.					

Get a physical mapping of the LDEV (special device files)

```

shell/iX> export HORCMINST=0
shell/iX> callci dstat | raidscan -find

```

DEVICE_FILE	UID	S/F	PORT	TARG	LUN	SERIAL	LDEV	PRODUCT_ID
/dev/ldev100	0	S	CL1-L	0	1	35013	17	OPEN-3
/dev/ldev101	0	S	CL1-L	0	2	35013	18	OPEN-3
/dev/ldev102	0	S	CL1-L	0	3	35013	19	OPEN-3
/dev/ldev103	0	S	CL1-L	0	4	35013	35	OPEN-3-CM

Describe the known HORCM_DEV & HORCM_INST on /etc/horcm*.conf

HORCM_DEV

#dev_group	dev_name	port#	TargetID	LU#	MU#
DSG1	dsvol0	CL1-L	0	1	0
DSG1	dsvol1	CL1-L	0	2	0
DSG1	dsvol2	CL1-L	0	3	0

HORCM_INST

#dev_group	ip_address	service
DSG1	HOSTB	horcm1

Restart "horcmstart.sh 0" as a JOB

```
shell/iX> horcmshutdown.sh 0
inst 0:
HORCM Shutdown inst 0 !!!
shell/iX> callci STREAM JRAIDMR0
#J17
shell/iX> callci SHOWJOB
```

JOBNUM	STATE	IPRI	JIN	JLIST	INTRODUCED	JOB NAME
#S28	EXEC		9	9	WED 9:10P	MANAGER.SYS
#J17	EXEC		10S	LP	WED 11:34P	JRAIDMR0,MANAGER.SYS

Porting notice for OpenVMS

This appendix describes operating system requirements, restrictions, and known issues for OpenVMS.

Porting notice for OpenVMS

Introduction

RM uses the UNIX domain socket for IPC (Inter Process Communication). While OpenVMS does not support the AF_UNIX socket, RAID Manager uses the Open VMS mailbox driver for inter-process communication between RAID Manager commands and the HORCM daemon.

Requirements and restrictions

Version of OpenVMS

RM uses CRTL, and needs the following version to support the ROOT directory for POSIX:

- OpenVMS Version 7.3-1 or later
- **The CRTL version must be installed before running RM.**

SY\$POSIX_ROOT

You need to define the **POSIX_ROOT** before running RAID Manager

Example:

```
$ DEFINE/TRANSLATION=(CONCEALED,TERMINAL)
SYS$POSIX_ROOT "Device:[directory]"
```

Mailbox driver

You need to redefine **LN\$TEMPORARY_MAILBOX** in the **LN\$PROCESS_DIRECTORY** table as shown below.

Example

```
$ DEFINE/TABLE=LN$PROCESS_DIRECTORY
LN$TEMPORARY_MAILBOX LN$GROUP
```

HORCM daemon startup

In OpenVMS, horcmstart.exe is created as a detached process or batch job by using the **DCL** command.

Using the detached process:

If you want the HORCM daemon to run in background, you will need to create the detached LOGINOUT.EXE process by using the **RUN /DETACHED** command. You will also need to make a command file for LOGINOUT.EXE.

The following are examples of the **loginhorcm*.com** file given to SYS\$INPUT for LOGINOUT.EXE. They show that **VMS4\$DKB100:[SYS0.SYSMGR.]** was defined as SYS\$POSIX_ROOT.

```
$ DEFINE/TRANSLATION=(CONCEALED,TERMINAL) SYS$POSIX_ROOT
"VMS4$DKB100:[SYS0.SYSMGR.]"
```

```
$ DEFINE DCL$PATH SYS$POSIX_ROOT: [horcm.usr.bin] ,SYS$POSIX_ROOT: [horcm.etc]
```

```
$ DEFINE/TABLE=LNMS$PROCESS_DIRECTORY LNM$TEMPORARY_MAILBOX LNM$GROUP
```

```
$ horcmstart 0
```

```
$ DEFINE/TRANSLATION=(CONCEALED,TERMINAL) SYS$POSIX_ROOT
"VMS4$DKB100:[SYS0.SYSMGR.]"
```

```
$ DEFINE DCL$PATH SYS$POSIX_ROOT: [horcm.usr.bin] ,SYS$POSIX_ROOT:[horcm.etc]
```

```
$ DEFINE/TABLE=LNMS$PROCESS_DIRECTORY LNM$TEMPORARY_MAILBOX LNM$GROUP
```

```
$ horcmstart 1
```

```
$ run /DETACHED SYS$SYSTEM:LOGINOUT.EXE /PROCESS_NAME=horcm0 -
```

```
_ $ /INPUT=VMS4$DKB100:[SYS0.SYSMGR.][horcm]loginhorcm0.com -
```

```
_ $ /OUTPUT=VMS4$DKB100:[SYS0.SYSMGR.][horcm]run0.out -
```

```
_ $ /ERROR=VMS4$DKB100:[SYS0.SYSMGR.][horcm]run0.err
```

```
%RUN-S-PROC_ID, identification of created process is 00004160
```

```
$
```

```

$
$ run /DETACHED SYS$SYSTEM:LOGINOUT.EXE /PROCESS_NAME=horcm1 -
_ $ /INPUT=VMS4$DKB100:[SYS0.SYSMGR.] [horcm]loginhorcm1.com -
_ $ /OUTPUT=VMS4$DKB100:[SYS0.SYSMGR.] [horcm]run1.out -
_ $ /ERROR=VMS4$DKB100:[SYS0.SYSMGR.] [horcm]run1.err
%RUN-S-PROC_ID, identification of created process is 00004166

```

You can verify that the HORCM daemon is running as a detached process by using the **SHOW PROCESS** command.

```
$ show process horcm0
```

```

25-MAR-2003 23:27:27.72 User: SYSTEM Process ID: 00004160
                          Node: VMS4 Process name: "HORCM0"

```

```
Terminal:
```

```

User Identifier: [SYSTEM]
Base priority: 4
Default file spec: Not available
Number of Kthreads: 1

```

```
Soft CPU Affinity: off
```

```

$
$ horcmshutdown 0 1
inst 0:
HORCM Shutdown inst 0 !!!
inst 1:
HORCM Shutdown inst 1 !!!
$

```

Command device

With OpenVMS, RAID Manager uses the SCSI Class driver to access the RAID Manager command device on the XP array, and defines “DG*” or DK*” as the logical name for the device.

You will need to define the physical device as either DG* or DK* by using DEFINE/SYSTEM command.

Example: \$ show device

Device Name	Device Status	Error Count	Volume Label	Free Blocks	Trans Count	Mnt Cnt
VMS4\$DKB0:	Online	0				
VMS4\$DKB100:	Mounted	0	ALPHASYS	30782220	414	1
VMS4\$DKB200:	Online	0				
VMS4\$DKB300:	Online	0				
VMS4\$DQA0:	Online	0				
\$1\$DGA145:	(VMS4)	Online				
\$1\$DGA146:	(VMS4)	Online				
:						
:						
\$1\$DGA153:	(VMS4)	Online				


```
$ DEFINE/SYSTEM DKA145
                                $1$DGA145:
$ DEFINE/SYSTEM DKA146
                                $1$DGA146:
:
:
$ DEFINE/SYSTEM DKA153      $1$DGA153:
```

-zx option for RAID Manager commands under OpenVMS

A number of commands in this book reference a **-zx** option. RAID Manager does not support the **-zx** option for OpenVMS, and the option will not appear on displays.

Syslog function

OpenVMS does not support the syslog function. Instead, the HORCM daemon uses the HORCM logging file.

Startup log files

Under OpenVMS, RAID Manager has two startup log files, which are separated by using PID.

For example in the `SY$POSIX_ROOT:[HORCM.LOG*.CURLOG]` directory:

```
HORCMLOG_VMS4  HORCM_VMS4_10530.LOG  
HORCM_VMS4_10531.LOG
```

Option syntax and case sensitivity

RAID Manager commands are case sensitive. OpenVMS users needs to change case sensitivity in `LOGIN.COM`.

The following upper-case strings are not case sensitive.

- “DG*”, “DK*”, or logical device names
- “-CLI”, “-FCA” (“-FHORC”), or “-FBC” (“-FMRCF”) in pair management commands
- “-CLI”, “-CLIWP”, “-CLIWN”, or “-CM” options in the **inraid** command
- Environment variable names, such as `HORCMINST`, controlled by `CRTL`

Define the following logical name in your `login.com` in order to distinguish uppercase and lowercase.

```
$ DEFINE DECC$ARGV_PARSE_STYLE ENABLE  
$ SET PROCESS/PARSE_STYLE=EXTENDED
```

Using the spawn command

You can start `HORCM` using the **spawn** command. The following are examples of using the **spawn** command.


```

Example $ spawn /NOWAIT /PROCESS=horcm0 horcmstart 0
        %DCL-S-SPAWNED, process HORCM0 spawned
$
        starting HORCM inst 0
$ spawn /NOWAIT /PROCESS=horcm1 horcmstart 1
        %DCL-S-SPAWNED, process HORCM1 spawned
$
        starting HORCM inst 1
$

```

Note that the subprocess (HORCM, the RM daemon) created by **spawn** will be terminated when the terminal is logged off or the session is terminated. To run the process independently of LOGOFF, use the **RUN /DETACHED** command.

Privileges for using RAID Manager

- A user account for RAID Manager must have the same privileges as “SYSTEM” (that is, it must be able to use the SCSI class driver and Mailbox driver directly). Some OpenVMS system administrators may not allow RAID Manager to run from the system account. In this case, create another account on the system, such as “RAdmin” that has the same privileges as “SYSTEM.”
- RAID Manager uses the Mailbox driver for communication between the RAID Manager components. So, the RAID Manager command processor and RM daemon (called HORCM) must have the same privileges.

If the RAID Manager command processor and HORCM execute with different privileges, then the RAID Manager command processor will hang or be unable to attach to the daemon.

Installation

RAID Manager requires that the logical name **sys\$posix_root** exist on the system. Define **sys\$posix_root** before installing RAID Manager.

It is recommended that the three logical names shown below be defined in LOGIN.COM prior to RAID Manager installation:

```

$ DEFINE/TRANSLATION= ( CONCEALED, TERMINAL)
SYS$POSIX_ROOT "Device: [directory]"
$ DEFINE DCL$PATH SYS$POSIX_ROOT: [horcm.usr.bin],
SYS$POSIX_ROOT: [horcm.etc]
$ DEFINE/TABLE=LNM$PROCESS_DIRECTORY
LNM$TEMPORARY_MAILBOX LNM$GROUP
$ DEFINE DECC$ARGV_PARSE_STYLE ENABLE
$ SET PROCESS/PARSE_STYLE=EXTENDED

```

The Device:[directory] you choose will be defined as SYS\$POSIX_ROOT.

To install RAID Manager:

Install RAID Manager by using the file

HP-AXPVMS-RMXP-V0115-4-1.PCSI

1. Insert and mount the installation media.
2. Execute the following command.

```

$ PRODUCT INSTALL RMXP
/source=Device: [PROGRAM.RM.OVMS] /LOG -
_ $ /destination=SYS$POSIX_ROOT: [000000]

```

where Device:[PROGRAM.RMOVMS] is where file
HP-AXPVMS-RMXP-V0115-4-1.PCSI exists.

3. Confirm the installation:

```

$ raidqry -h
Model : Raid-Manager-XP/OpenVMS
Ver&Rev: 01.15.02
:
:

```

To obtain the installation history:

```

$ PRODUCT SHOW HISTORY RMXP /FULL

```

To uninstall RAID Manager:

```

$ PRODUCT REMOVE RMXP /LOG

```

Known issues and concerns

Rebooting on PAIR state (writing disabled)

OpenVMS does not show the volumes with writing disabled (e.g., SVOL_PAIR) at system startup; therefore, the S-VOLs are hidden when rebooting in PAIR state or SUSPEND-mode.

You can verify that the **show device** and **in RAID** commands do not show the S-VOLs after a reboot as shown below (that is, **DGA148** and **DGA150** devices are in the SVOL_PAIR state and do not display).

```
$ sh dev dg
```

Device Name	Device	Status	Error Count	Volume Free	Trans Mnt	Label	Blocks Count	Cnt
\$1\$DGA145:	(VMS4)	Online	0					
\$1\$DGA146:	(VMS4)	Online	0					
\$1\$DGA147:	(VMS4)	Online	0					
\$1\$DGA149:	(VMS4)	Online	0					
\$1\$DGA151:	(VMS4)	Online	0					
\$1\$DGA152:	(VMS4)	Online	0					
\$1\$DGA153:	(VMS4)	Online	0					

```
$ in RAID DKA145-153 -cli
```

DEVICE_FILE	PORT	SERIAL	LDEV	CTGC	B/12	SSID	R:Group	PRODUCT_ID
DKA145	CL1-H	30009	145	-	-	-	-	OPEN-9-CM
DKA146	CL1-H	30009	146	-	s/P/ss	0004	5:01-11	OPEN-9
DKA147	CL1-H	30009	147	-	s/S/ss	0004	5:01-11	OPEN-9
DKA148	-	-	-	-	-	-	-	-
DKA149	CL1-H	30009	149	-	P/s/ss	0004	5:01-11	OPEN-9
DKA150	-	-	-	-	-	-	--	-
DKA151	CL1-H	30009	151	-	P/s/ss	0004	5:01-11	OPEN-9

```
DKA152      CL1-H 30009 152 -   s/s/ss 0004 5:01-11 OPEN-9
DKA153      CL1-H 30009 153 -   s/s/ss 0004 5:01-11 OPEN-9
```

```
$ inqraid DKA148
sys$assign : DKA148 -> errcode = 2312
DKA148 -> OPEN: no such device or address
```

After enabling the S-VOL for writing by using either the **pairsplit** or **horctakeover** command, you will need to execute the **mcr sysman** command to use the S-VOLs for backup or disaster recovery.

```
$ pairsplit -g CAVG -rw
$ mcr sysman
SYSMAN> io auto
SYSMAN> exit
```

```
$ sh dev dg
```

Device Name	Device	Error Count	Volume Label	Free Blocks	Trans Count	Mnt Cnt
\$1\$DGA145:	(VMS4) Online	0				
\$1\$DGA146:	(VMS4) Online	0				
\$1\$DGA147:	(VMS4) Online	0				
\$1\$DGA148:	(VMS4) Online	0				
\$1\$DGA149:	(VMS4) Online	0				
\$1\$DGA150:	(VMS4) Online	0				
\$1\$DGA151:	(VMS4) Online	0				
\$1\$DGA152:	(VMS4) Online	0				
\$1\$DGA153:	(VMS4) Online	0				

Startup procedures using a detached process

Creates the shareable logical name for RAID if undefined initially. RAID Manager needs to define the physical device (**\$1\$DGA145...**) as either DG* or DK* by using the **SHOW DEVICE** command and the **DEFINE/SYSTEM** command, but then does not need to be mounted.

```
$ show device
```

Device Name	Device	Device Status	Error Count	Volume Label	Free Blocks	Trans Count	Mnt Cnt
\$1\$DGA145:	(VMS4)	Online	0				
\$1\$DGA146:	(VMS4)	Online	0				
:							
:							
\$1\$DGA153:	(VMS4)	Online	0				
\$							

```

$ DEFINE/SYSTEM DKA145 $1$DGA145:
$ DEFINE/SYSTEM DKA146 $1$DGA146:
:
:
$ DEFINE/SYSTEM DKA153 $1$DGA153:

```

Defining the environment for RAID Manager in LOGIN.COM

You need to define the path for the RAID Manager commands to **DCL\$PATH** as the foreign command.

```

$ DEFINE DCL$PATH
SYS$POSIX_ROOT:[horcm.usr.bin], SYS$POSIX_ROOT:
[horcm.etc]

```

If the RAID Manager command and the RM daemon (HORCM) will be executing in different jobs (using a different terminal), then you must redefine **LN\$TEMPORARY_MAILBOX** in the **LN\$PROCESS_DIRECTORY** table as shown below:

```

$ DEFINE/TABLE=LN$PROCESS_DIRECTORY
LN$TEMPORARY_MAILBOX LN$GROUP

```

Discovering and describing the command device on `sys$posix_root:[etc]horcm0.conf`

```

$ inqraid DKA145-151 -CLI

```

DEVICE_FILE	PORT	SERIALLDEV	CTGC/B/12	SSID	R:Group	PRODUCT_ID
--------------------	-------------	-------------------	------------------	-------------	----------------	-------------------

```

DKA145      CL1-H  30009 145  -  -      -      -      OPEN-9-CM
DKA146      CL1-H  30009 146  -  s/S/ss 0004      OPEN-9
                                         5:01-11
DKA147      CL1-H  30009 147  -  s/P/ss 0004      OPEN-9
                                         5:01-11
DKA148      CL1-H  30009 148  -  s/S/ss 0004      OPEN-9
                                         5:01-11
DKA149      CL1-H  30009 149  -  s/P/ss 0004      OPEN-9
                                         5:01-11
DKA150      CL1-H  30009 150  -  s/S/ss 0004      OPEN-9
                                         5:01-11
DKA151      CL1-H  30009 151  -  s/P/ss 0004      OPEN-9
                                         5:01-11

```

```

SYS$POSIX_ROOT:[etc]horcm0.conf
HORCM_MON

```

```

#ip_address  service  poll(10ms  timeout(10ms)
127.0.0.1    30001    1000      3000

```

```

HORCM_CMD
#dev_name    dev_name    dev_name
DKA145

```

You will have to start HORCM without a description for HORCM_DEV and HORCM_INST because **target ID and LUN are unknown**.

You can determine the mapping of a physical device with a logical name by using **raidscan -find**.

Executing a "horcmstart 0"

```

$ run /DETACHED SYS$SYSTEM:LOGINOUT.EXE
  /PROCESS_NAME=horcm0 -
_$ /INPUT=VMS4$DKB100:[SYS0.SYSMGR.][horcm]
  loginhorcm0.com -
_$ /OUTPUT=VMS4$DKB100:[SYS0.SYSMGR.][horcm]run0.out -
_$ /ERROR=VMS4$DKB100:[SYS0.SYSMGR.][horcm]run0.err
%RUN-S-PROC_ID, identification of created process is
00004160

```

Verifying physical mapping of the logical device

```
$ HORCMINST := 0
$ raidscan -pi DKA145-151 -find
```

DEVICE_FILE	UID	S/F	PORT	TARG	LUN	SERIAL	LDEV	PRODUCT_ID
DKA145	0	F	CL1-H	0	1	30009	145	OPEN-9-CM
DKA146	0	F	CL1-H	0	2	30009	146	OPEN-9
DKA147	0	F	CL1-H	0	3	30009	147	OPEN-9
DKA148	0	F	CL1-H	0	4	30009	148	OPEN-9
DKA149	0	F	CL1-H	0	5	30009	149	OPEN-9
DKA150	0	F	CL1-H	0	6	30009	150	OPEN-9
DKA151	0	F	CL1-H	0	7	30009	151	OPEN-9

```
$ horcmshutdown 0
inst 0:
HORCM Shutdown inst 0 !!!
```

Describing the known HORCM_DEV on sys\$posix_root:[etc]horcm*.conf

For horcm0.conf

```
HORCM_DEV
```

#dev_group	dev_name	port#	TargetID	LU#	MU#
VG01	oradb1	CL1-H	0	2	0
VG01	oradb2	CL1-H	0	4	0
VG01	oradb3	CL1-H	0	6	0

```
HORCM_INST
```

#dev_group	ip_address	service
VG01	HOSTB	horcm1

FOR horcm1.conf

```
HORCM_DEV
```

#dev_group	dev_name	port#	TargetID	LU#	MU#
VG01	oradb1	CL1-H	0	3	0
VG01	oradb2	CL1-H	0	5	0
VG01	oradb3	CL1-H	0	7	0

```

HORCM_INST
#dev_group          ip_address          service
VG01                HOSTA                horcm0

```

The UDP port name for HORCM communication in **"SYS\$SYSROOT:[000000.TCPIP\$ETC]SERVICES.DAT"** is defined as shown in the example below.

```

horcm0      30001/udp
horcm1      30002/udp

```

Starting "horcm 0" and "horcm 1" as detached processes

```

$ run /DETACHED SYS$SYSTEM:LOGINOUT.EXE
  /PROCESS_NAME=horcm0 -
_ $ /INPUT=VMS4$DKB100:[SYS0.SYSMGR.] [horcm]
  loginhorcm0.com -
_ $ /OUTPUT=VMS4$DKB100:[SYS0.SYSMGR.] [horcm]run0.out -
_ $ /ERROR=VMS4$DKB100:[SYS0.SYSMGR.] [horcm]run0.err
%RUN-S-PROC_ID, identification of created process is
00004160
$
$
$ run /DETACHED SYS$SYSTEM:LOGINOUT.EXE
  /PROCESS_NAME=horcm1 -
_ $ /INPUT=VMS4$DKB100:[SYS0.SYSMGR.] [horcm]
  loginhorcm1.com -
_ $ /OUTPUT=VMS4$DKB100:[SYS0.SYSMGR.] [horcm]run1.out -
_ $ /ERROR=VMS4$DKB100:[SYS0.SYSMGR.] [horcm]run1.err
%RUN-S-PROC_ID, identification of created process is
00004166

```

You will be able to verify that HORCM daemon is running as a detached process by using the **SHOW PROCESS** command.


```
$ show process horcm0
```

```
25-MAR-2003 23:27:27.72 User: SYSTEM Process ID:  
00004160
```

Node:

```
VMS4 Process name: "HORCM0"
```

Terminal:

```
User Identifier: [SYSTEM]  
Base priority: 4  
Default file spec: Not available  
Number of Kthreads: 1
```

```
Soft CPU Affinity: off
```

DCL command examples

1. Setting the environment variable by using symbol:

```
$ HORCMINST := 0
```

```
$ HORCC_MRCF := 1
```

```
$ raidqry -l
```

```
No Group Hostname HORCM_ver Uid Serial# Micro_ver Cache(MB)  
1 --- VMS4 01.12.00 0 30009  
21-04-04/00 8192
```

```
$
```

```
$ pairdisplay -g VG01 -fdc
```

Group	PairVol(L/R)	Device	FileM	Seq#	LDEV#	Status	%	P-LDEV#	M
VG01	oradb1(L)	DKA146	0	30009	146..	S-VOL PAIR,	100	147	-
VG01	oradb1(R)	DKA147	0	30009	147..	P-VOL PAIR,	100	146	-
VG01	oradb2(L)	DKA148	0	30009	148..	S-VOL PAIR,	100	149	-
VG01	oradb2(R)	DKA149	0	30009	149..	P-VOL PAIR,	100	148	-
VG01	oradb3(L)	DKA150	0	30009	150..	S-VOL PAIR,	100	151	-
VG01	oradb3(R)	DKA151	0	30009	151..	P-VOL PAIR,	100	150	-

2. Removing the environment variable:

```
$ DELETE/SYMBOL HORCC_MRCF
$ pairdisplay -g VG01 -fdc
```

Group	PairVol (L/R)	Device_File	Seq#	LDEV#	.P/S	Status	Fence	%	,P-LDEV#M	
VG01	oradb1 (L)	DKA146	30009	146..	SMPL	----	-----,	-----	----	-
VG01	oradb1 (R)	DKA147	30009	147..	SMPL	----	-----,	-----	----	-
VG01	oradb2 (L)	DKA148	30009	148..	SMPL	----	-----,	-----	----	-
VG01	oradb2 (R)	DKA149	30009	149..	SMPL	----	-----,	-----	----	-
VG01	oradb3 (L)	DKA150	30009	150..	SMPL	----	-----,	-----	----	-
VG01	oradb3 (R)	DKA151	30009	151..	SMPL	----	-----,	-----	----	-

\$

3. Changing the default log directory:

```
$ HORCC_LOG := /horcm/horcm/TEST
$ pairdisplay
PAIRDISPLAY: requires '-x xxx' as argument
PAIRDISPLAY: [EX_REQARG] Required Arg list
Refer to the command log
(SYS$POSIX_ROOT:[HORCM.HORCM.TEST]HORCC_VMS4.LOG
(/HORCM
/HORCM/TEST/horcc_vms4.log)) for details.
```

4. Resetting to the default log directory

```
$ DELETE/SYMBOL HORCC_LOG
```

5. Specifying the device described in scandev.LIS

```
$ define dev_file SYS$POSIX_ROOT:[etc]SCANDEV
$ type dev_file
DKA145-150
$
```

```
$ pipe type dev_file | inqraid -CLI
```

DEVICE_FILE	PORT	SERIALLDEV	CTG	C/B/12	SSID	R:Group	PRODUCT_ID
DKA145	CL1-H	30009 145	-	-	-	-	OPEN-9-CM
DKA146	CL1-H	30009 146	-	s/S/ss	0004	5:01-11	OPEN-9
DKA147	CL1-H	30009 147	-	s/P/ss	0004	5:01-11	OPEN-9
DKA148	CL1-H	30009 148	-	s/S/ss	0004	5:01-11	OPEN-9
DKA149	CL1-H	30009 149	-	s/P/ss	0004	5:01-11	OPEN-9
DKA150	CL1-H	30009 150	-	s/S/ss	0004	5:01-11	OPEN-9

6. Making the configuration file automatically:

You will be able to omit steps 3 to 6 on startup by using the **mkconf** command.

```
$ type dev_file
DKA145-150
$
$ pipe type dev_file | mkconf -g URA -i 9
starting HORCM inst 9
HORCM Shutdown inst 9 !!!
```

```
A CONFIG file was successfully completed.
HORCM inst 9 finished successfully.
starting HORCM inst 9
```

DEVICE_FILE	Group	PairVol	PORT	TARG	LUN	M	SERIALLDEV
DKA145	-	-	-	-	-	-	30009 145
DKA146	URA	URA_000	CL1-H	0	2	0	30009 146
DKA147	URA	URA_001	CL1-H	0	3	0	30009 147
DKA148	URA	URA_002	CL1-H	0	4	0	30009 148
DKA149	URA	URA_003	CL1-H	0	5	0	30009 149
DKA150	URA	URA_004	CL1-H	0	6	0	30009 150

```
HORCM Shutdown inst 9 !!!
Please check
'SYS$SYSROOT:[SYSMGR]HORCM9.CONF', 'SYS$SYSROOT:[SY
```

```

SMGR.LOG9.CURLOG]
HORCM_*.LOG', and modify 'ip_address & service'.
HORCM inst 9 finished successfully.
$

```

```

SYS$SYSROOT:[SYSMGR]horcm9.conf
(/sys$sysroot/sysmgr/horcm9.conf)

```

```

# Created by mkconf on Thu Mar 13 20:08:41

```

```

HORCM_MON

```

```

#ip_address      service      poll(10ms timeout(10ms)
127.0.0.1        52323        1000         3000

```

```

HORCM_CMD

```

```

#dev_name        dev_name        dev_name
#UnitID 0 (Serial# 30009)
DKA145

```

```

# ERROR [CMDDEV] DKA145 SER = 30009 LDEV =
145 [OPEN-9-CM `

```

```

HORCM_DEV

```

```

#dev_group      dev_name        port#      TargetID      LU#      MU#
# DKA146        SER = 30009 LDEV = 146 [ FIBRE FCTBL = 3 ]
URA            URA_000        CL1-H      0             2        0
# DKA147        SER = 30009 LDEV = 147 [ FIBRE FCTBL = 3 ]
URA            URA_001        CL1-H      0             3        0
# DKA148        SER = 30009 LDEV = 148 [ FIBRE FCTBL = 3 ]
URA            URA_002        CL1-H      0             4        0
# DKA149        SER = 30009 LDEV = 149 [ FIBRE FCTBL = 3 ]
URA            URA_003        CL1-H      0             5        0
# DKA150        SER = 30009 LDEV = 150 [ FIBRE FCTBL = 3 ]
URA            URA_004        CL1-H      0             6        0

```

```

HORCM_INST

```

```

#dev_group      ip_address      service
URA            127.0.0.1      52323

```

7. Using \$1\$* naming as native device name

You can use the native device without the DEFINE/SYSTEM command by specifying \$1\$* naming directly.

Examples

```
$ inqraid $1$DGA145-155 -CLI
```

DEVICE_FILE	PORT	SERIAL	LDEV	CTG	C/B/12	SSID	R:Group	PRODUCT_ID
\$1\$DGA145	CL2-H	30009	145	-	-	-	-	OPEN-9-CM
\$1\$DGA146	CL2-H	30009	146	-	s/P/ss	0004	5:01-11	OPEN-9
\$1\$DGA147	CL2-H	30009	147	-	s/S/ss	0004	5:01-11	OPEN-9
\$1\$DGA148	CL2-H	30009	148	0	P/s/ss	0004	5:01-11	OPEN-9

```
$ pipe show device | INQRAID -CLI
```

DEVICE_FILE	PORT	SERIAL	LDEV	CTG	C/B/12	SSID	R:Group	PRODUCT_ID
\$1\$DGA145	CL2-H	30009	145	-	-	-	-	OPEN-9-CM
\$1\$DGA146	CL2-H	30009	146	-	s/P/ss	0004	5:01-11	OPEN-9
\$1\$DGA147	CL2-H	30009	147	-	s/S/ss	0004	5:01-11	OPEN-9
\$1\$DGA148	CL2-H	30009	148	0	P/s/ss	0004	5:01-11	OPEN-9

```
$ pipe show device | MKCONF -g URA -i 9
```

```
starting HORCM inst 9
HORCM Shutdown inst 9 !!!
A CONFIG file was successfully completed.
HORCM inst 9 finished successfully.
```

```
starting HORCM inst 9
```

DEVICE_FILE	Group	PairVol	PORT	TARG	LUN	M	SERIAL
\$1\$DGA145	-	-	-	-	-	-	30009 145
\$1\$DGA146	URA	URA_000	CL2-H	0	2	0	30009 146
\$1\$DGA147	URA	URA_001	CL2-H	0	3	0	30009 147
\$1\$DGA148	URA	URA_002	CL2-H	0	4	0	30009 148

```
HORCM Shutdown inst 9 !!!
```

```
Please check 'SYS$SYSROOT:[SYSMGR]HORCM9.CONF',
'SYS$SYSROOT:[SYSMGR.LOG9.CURLOG]
HORCM_*.LOG', and modify 'ip_address & service'.
HORCM inst 9 finished successfully.
```

```
$
```

```
$ pipe show device | RAIDSCAN -find
```

```

DEVICE_FILE      UID  S/F  PORT   TARG  LUN   SERIAL  LDEV  PRODUCT_ID
$1$DGA145        0   F   CL2-H   0     1     30009   145  OPEN-9-CM
$1$DGA146        0   F   CL2-H   0     2     30009   146  OPEN-9
$1$DGA147        0   F   CL2-H   0     3     30009   147  OPEN-9
$1$DGA148        0   F   CL2-H   0     4     30009   148  OPEN-9

```

```
$ pairdisplay -g BCVG -fdc
```

```

Group   PairVol(L/R) Device_File   M   ,Seq#,LDEV#..P/S,Status,%   ,P-LDEV# M
BCVG    oradb1(L)     $1$DGA146   0  30009  146..P-VOL PAIR, 100     147 -
BCVG    oradb1(R)     $1$DGA147   0  30009  147..S-VOL PAIR, 100     146 -
$

```

```
$ pairdisplay -dg $1$DGA146
```

```

Group   PairVol(L/R) (Port#,TID, LU-M) ,Seq#,LDEV#..P/S,Status,Seq#,P-LDEV# M
BCVG    oradb1(L)    (CL1-H , 0,  2-0 )30009  146..P-VOL PAIR,30009  147 -
BCVG    oradb1(R)    (CL1-H , 0,  3-0 )30009  147..S-VOL PAIR,-----  146 -
$

```

Startup procedures on bash

RAID Manager is not recommended for use with **bash** because **bash** is not provided as an official release in some versions of OpenVMS.

Creating the shareable logical name for RAID if undefined initially

You need to define the physical device (**\$1\$DGA145...**) as either **DG*** or **DK*** by using the **SHOW DEVICE** command and the **DEFINE/SYSTEM** command, but then it does not need to be mounted.

```
$ show device
```

```

Device           Device           Error   Volume
Free  Trans  Mnt
Name           Status           Count   Label
Blocks Count  Cnt
$1$DGA145:      (VMS4)  Online           0
$1$DGA146:      (VMS4)  Online           0
:

```

```

:
$1$DGA153:      (VMS4)  Online          0
$
$ DEFINE/SYSTEM DKA145 $1$DGA145:
$ DEFINE/SYSTEM DKA146 $1$DGA146:
:
:
$ DEFINE/SYSTEM DKA153 $1$DGA153:

```

Defining the environment for RAID Manager in LOGIN.COM

If the RAID Manager commands and the RM daemon (HORCM_ will be executing in different jobs (on a different terminal), then you must redefine **LNMS\$TEMPORARY_MAILBOX** in the **LNMS\$PROCESS_DIRECTORY** table, as shown below:

```

$ DEFINE/TABLE=LNMS$PROCESS_DIRECTORY
LNMS$TEMPORARY_MAILBOX LNMS$GROUP

```

Discovering and describing the command device on /etc/horcm0.conf

```
bash$ inqraid DKA145-151 -CLI
```

DEVICE_FILE	PORT	SERIAL	LDEV	CTG	C/B/12	SSID	R:Group	PRODUCT_ID
DKA145	CL1-H	30009	145	-	-	-	-	OPEN-9-CM
DKA146	CL1-H	30009	146	-	s/S/ss	0004	5:01-11	OPEN-9
DKA147	CL1-H	30009	147	-	s/P/ss	0004	5:01-11	OPEN-9
DKA148	CL1-H	30009	148	-	s/S/ss	0004	5:01-11	OPEN-9
DKA149	CL1-H	30009	149	-	s/P/ss	0004	5:01-11	OPEN-9
DKA150	CL1-H	30009	150	-	s/S/ss	0004	5:01-11	OPEN-9
DKA151	CL1-H	30009	151	-	s/P/ss	0004	5:01-11	OPEN-9

```
/etc/horcm0.conf
```

```
HORCM_MON
```

```
#ip_address    service    poll(10ms timeout(10ms)
127.0.0.1      52000      1000        3000
```

```
HORCM_CMD
#dev_name      dev_name      dev_name
DKA145
```

```
HORCM_DEV
#dev_group     dev_name      port#         TargetID
LU#    MU#
```

```
HORCM_INST
#dev_group     ip_address    service
```

You will have to start HORCM without a description for HORCM_DEV and HORCM_INST **because the target ID and LUN are unknown.**

You will be able to determine the mapping of a physical device with a logical name by using the **raidscan -find** command.

Executing “horcmstart 0” as a background process

```
bash$ horcmstart 0 &
18
bash$
    starting HORCM inst 0
```

Verifying physical mapping of the logical device

```
bash$ export HORCMINST=0
bash$ raidscan -pi DKA145-151 -find
```

DEVICE	FILE	UID	S/F	PORT	TARGLUN	SERIAL	LDEV	PRODUCT_ID
DKA145		0	F	CL1-H	0 1	30009	145	OPEN-9-CM
DKA146		0	F	CL1-H	0 2	30009	146	OPEN-9
DKA147		0	F	CL1-H	0 3	30009	147	OPEN-9
DKA148		0	F	CL1-H	0 4	30009	148	OPEN-9
DKA149		0	F	CL1-H	0 5	30009	149	OPEN-9


```
DKA150    0  F  CL1-H  0   6  30009   150  OPEN-9
DKA151    0  F  CL1-H  0   7  30009   151  OPEN-9
```

Describing the known HORCM_DEV on /etc/horcm*.conf

For horcm0.conf

HORCM_DEV

#dev_group	dev_name	port#	TargetID	LU#	MU#
VG01	oradb1	CL1-H	0	2	0
VG01	oradb2	CL1-H	0	4	0
VG01	oradb3	CL1-H	0	6	0

HORCM_INST

#dev_group	ip_address	service
VG01	HOSTB	horcm1

For horcm1.conf

HORCM_DEV

#dev_group	dev_name	port#	TargetID	LU#	MU#
VG01	oradb1	CL1-H	0	3	0
VG01	oradb2	CL1-H	0	5	0
VG01	oradb3	CL1-H	0	7	0

HORCM_INST

#dev_group	ip_address	service
VG01	HOSTA	horcm0

Starting "horcmstart 0 1"

The HORCM subprocess created by bash will be terminated when bash terminates.

```
bash$ horcmstart 0 &  
19  
bash$  
starting HORCM inst 0  
bash$ horcmstart 1 &  
20  
bash$  
starting HORCM inst 1
```

Glossary

ACA	HP StorageWorks Asynchronous Continuous Access XP.
ACP	Array Control Processor. The ACP handles passing data between cache and the physical drives. ACPs work in pairs. In the event of an ACP failure, the redundant ACP takes control. Both ACPs work together sharing the load.
AL	Arbitrated loop.
AL-PA	Arbitrated loop physical address
allocation	The ratio of allocated storage capacity versus total capacity as a percentage. “Allocated storage” refers to those LDEVs that have paths assigned to them. Allocated storage capacity is the sum of the storage of these LDEVs. Total capacity is the sum of the capacity of all LDEVs on the array.
array group	The number of physical disk drives contained in a RAID group. This number depends on the RAID configuration. For example, there may be two or four physical disks in a RAID1 group. There are four or eight physical disks in a RAID5 group.
BC	HP StorageWorks Business Copy XP. BC lets you maintain up to nine internal copies of logical volumes on the disk array.
CA	HP StorageWorks Continuous Access XP. CA lets you create and maintain duplicate copies of logical volumes on a remote disk array.
cache	Very high speed memory used to speed I/O transaction time. All reads and writes to the array are sent to the cache. The data is buffered there until the transfer to/from physical disks (with slower data throughput) is complete. Cache memory speeds I/O throughput to the application.
CE	Customer engineer.

CFW	Cache fast write.
CH	Channel.
CHA (channel adapter)	The channel adapter (CHA) provides the interface between the disk array and the external host system. Occasionally this term is used synonymously with the term channel host interface processor (CHIP)
CHIP (channel host interface processor)	Synonymous with the term channel adapter (CHA).
CHP (channel processor)	The processor(s) located on the channel adapter (CHA).
CHPID	Channel path identifier.
CKD	Count key data.
CLI	Command line interface.
command device	A volume on the disk array that accepts CA or BC control operations, which are then executed by the disk array.
Command View	HP StorageWorks Command View XP. A software product used for managing XP arrays.
configuration file	A file that defines the pair configurations.
CTGID (consistency group ID)	The group identifier for which the disk array guarantees the sequence of asynchronous data transfer.
control unit (CU)	To organize the storage space attached, you can group similarly configured logical devices (LDEVs) with unique control unit images (CUs). CUs are numbered sequentially. An LDEV requires both a CU number and an LDEV number to identify it.
CU	control unit.
CVS	CVS devices (OPEN-x CVS) are custom volumes that are smaller than normal fixed-sized logical disk devices (volumes). They are created using the Volume Size Configuration feature of Command View XP or LUN Configuration Manager XP.

disk group	The physical disk locations associated with a parity group.
disk type	The manufacturer's label in the physical disk controller firmware. In most cases, the disk type is identical to the disk model number.
DKA (disk adapter)	Synonymous with the ACP.
DKC (disk controller unit)	The array cabinet that houses the channel adapters and service processor (SVP).
DKU (disk cabinet unit)	The array cabinets that house the physical disks.
DRR (disk recovery and restore unit)	The unit responsible for data recovery and restoration in the event of a cache failure. Located on the ACP.
daemon	A process in the UNIX system that waits for events and that does not disappear after an event is carried out.
DW	Duplex write.
DWL	Duplex write line.
emulation modes	Emulation modes can be assigned to LDEVs to make them operate like standard OPEN system disk drives. The emulation mode of an LDEV determines its capacity.
EPO	Emergency power-off.
ESCON	Enterprise Systems Connection (an IBM trademark). A set of IBM and vendor products that interconnect S/390 computers with each other and with attached storage, locally attached workstations, and other devices using optical fiber technology and switches called ESCON Directors.
expanded LUN	A LUN is normally associated with only a single LDEV. The LUN Size Expansion (LUSE) feature allows a LUN to be associated with 2-36 LDEVs. Essentially, LUSE makes it possible for applications to access a single large pool of storage. The LUSE feature is available when the HP StorageWorks LUN Configuration Manager product is installed.
ExSA	Extended serial adapter.

FC	Fibre Channel.
FC-AL	Fibre Channel arbitrated loop.
failover	Automatic change from using a failed portion and to an alternate component in order to continue operations.
fence level	A method for setting rejection of a write I/O request from the host according to the condition of mirroring consistency.
FICON	An I/O technology developed by IBM to connect mainframes to storage devices at higher speeds and greater distances than the earlier Enterprise Systems Connection (ESCON)
GB	Gigabytes.
HA	High availability.
HBA	Host bus adapter.
HORCM_CMD	A section of the RM instance configuration file that defines the disk devices used by RM to communicate with the disk array.
HORCM_DEV	A section of the RM instance configuration file that defines the volumes of the instance.
HORCM_INST	A section of the RM instance configuration file that defines how RM groups link to remote RM instances.
HORCM_MON	A section of the RM instance configuration file that defines the instance you are configuring.
host mode	Each port can be configured with various host modes. The host mode determines the array's behavior toward a specific type of host.
hot standby	Using one or more servers as a standby in case of a primary server failure.
HP	Hewlett-Packard Company.
instance	An independent copy of RM. Instances are local or remote and can run on the same host.

instance configuration file	A file that defines the link between a volume and an RM instance. This file consists of four sections: HORCM_MON , HORCM_CMD , HORCM_DEV and HORCM_INST .
LCP	Local control port.
LDEV	Logical device. An LDEV is created when a RAID group is divided into sections using a host emulation mode (for example, OPEN-9 or OPEN-M). The number of resulting LDEVs depends on the emulation mode. The term LDEV is often used synonymously with the term volume.
local disk	A disk in the local array. Sometimes refers to a disk in a local host.
local instance	The RM instance currently being configured or the instance to which commands are issued.
LUN	Logical unit number. A LUN results from mapping a SCSI logical unit number, port ID, and LDEV ID to a RAID group. The size of the LUN is determined by the emulation mode of the LDEV and the number of LDEVs associated with the LUN. For example, a LUN associated with two OPEN-3 LDEVs will have a size of 4,693 MB.
LUSE	Logical Unit Size Expansion, a feature which logically combines LDEVs so they appear as a larger LDEV. This allows a LUN to be associated with 2 to 36 LDEVs. Essentially, LUSE makes it possible for applications to access data requiring a large amount of disk space.
MB	Megabytes.
MCU	Main control unit.
OPEN-x	A general term describing any one of the supported OPEN emulation modes (for example, OPEN-L).
OS	Operating system.
PA	Physical address.
parity group	Synonymous with the term RAID group.
partition	To divide a disk according to the UNIX kernel or device driver layer into two or more areas, which will be treated as if they were two or more physical disks.

path	“Path” and “LUN” are synonymous. Paths are created by associating a port, a target, and a LUN ID with one or more LDEVs.
PB	Petabytes.
port	A connector on a channel adapter card in the disk array. A port passes data between the disk array and external devices, such as a host. Ports are named using a port group and port letter, for example, CL1-A.
P-VOL	The primary or main volume that contains data to be copied.
RAID	Redundant array of independent disks.
remote console PC	The PC running HP StorageWorks Remote Control XP.
Remote Control (RC)	HP StorageWorks Remote Control XP. A software product used for managing XP arrays.
remote instance	The instance with which the local instance communicates, as configured in the HORCM_INST section of the RM instance configuration file.
R-SIM	Remote service information message.
RAID group	Synonymous with RAID group. A group of disks configured to provide enhanced redundancy, performance, or both.
RCP	Remote control port.
RCU	Remote control unit.
shell script	A command sequence executed by a UNIX shell.
script file	A file containing a shell script.
SCSI	Small computer system interface.
SIM	Service information message.
SNMP	Simple Network Management Protocol.
special files	Files that indicate physical devices and are different from regular files in the UNIX system. The functions of the device drivers (that is, access to system peripherals) become available through these special files.

SSID	Storage subsystem identification.
S-VOL	Secondary or remote volume. The volume that receives the data from the primary volume.
SVP	Service processor. A notebook computer built into the disk array. The SVP provides a direct interface to the disk array. It is used only by the HP service representative.
takeover	The actions of a standby server that takes over processing from the previously active server.
TB	Terabytes.
TID	Target ID.
Volume	Synonymous with LDEV.
VSC	Volume Size Configuration. A feature that defines custom volumes (CVS volumes) that are smaller than normal fixed-sized logical disk devices (volumes).
XDF	Extended distance feature (for ExSA channels).

A

addresses

 Fibre Channel conversion in RM 363

C

command devices, switching 25

commands

 using RAID Manager 60

configuration

 setting up 29

configuration file examples 265

configuration file parameters 43

conventions, documentation 10

D

DCL command examples 377

disk array(s)

 supported 9

disk arrays 9

documentation conventions 10

drivescan command option 208

E

env command option 210

environment variables 74

error codes 254

error reporting 250

F

features 18

Fibre Channel

 addressing in RM 363

Fibre Channel addressing 345

findcmddev command 211

G

general commands 104

glossary 387

H

HA control script state transitions 312

HA system configuration

 recovery procedures 339

help, getting 10

horcctl command 107

horcmshutdown command 110

horcmstart command 111

host log file output 72

HP storage website 11

HP technical support 11

I

inqraid command 116

installing
 MPE/iX 33
 OpenVMS 39
installing RAID Manager
 UNIX systems 30
instances 22
 RAID Manager 22

L
log directories 72

M
mkconf command 126
mount command option 212
MPE socket hang 356
MPE/iX
 installing 33, 356
 known issues 357
 porting notice 354
 restrictions 354
 start-up procedures 359
 uninstalling 356

N
norctakeover command 112

O
OpenVMS
 installing 39, 369
 known issues 371

P
paircreate command 129
paircurchk command 136
pairdisplay command 139
pairevolchk command 180
pairevtwait command 149
pairmon command 155

pairresync command 158
pairsplit command 166
pairsyncwait command 173
parameters, configuration file 43
porting notice, MPE/iX 354
portscan command option 215
PVOL-takeover function 338

R
RAID Manager
 command devices 24
 features 18
 general commands 104
 instances 22
 product description 17
 system requirements 28
 topologies 22
 using 59
 Windows NT/2000/2003 command
 options 207
RAID Manager commands 60
raidar command 189
raidqry command 192
raidscan command 195
raidvchdsp command 233
raidvchkdsp command 233
raidvchkscant command 239
recovery procedures 343
 HA system configuration 339
related documentation 9

S
scripts with RAID Manager commands 61
SCSI pass-thru driver 355
setenv command option 217
setting up RAID Manager 29
sleep command option 218

Start-up procedures using detached process on

DCL 372

state transistions 312

StorageWorks, supported arrays 9

supported disk arrays 9

Surestore, supported arrays 9

S-VOL data consistency function 329

SVOL-takeover function 336

swap-takeover function 334

switching command devices 25

sync command option 219

system requirements

RAID Manager 28

T

takeover-switch function 332

technical support, HP 11

topologies 22

troubleshooting 103

U

umount command option 222

UNIX systems

installing 30

user files

creating 73

usetenv command 224

using RAID Manager 59

using RAID Manager commands 60

V

variables, environment 74

W

website, HP storage 11

Windows NT/2000/2003 command

options 207

