

Hitachi AMS 2000 Family ShadowImage In-system Replication User Guide

FASTFIND LINKS

[Document organization](#)

[Release notes and readme](#)

[Getting help](#)

[Table of Contents](#)

Copyright © 2010 Hitachi Ltd., Hitachi Data Systems Corporation, ALL RIGHTS RESERVED

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

Hitachi, Ltd. and Hitachi Data Systems reserve the right to make changes to this document at any time without notice and assume no responsibility for its use. Hitachi, Ltd. and Hitachi Data Systems products and services can only be ordered under the terms and conditions of Hitachi Data Systems' applicable agreements.

All of the features described in this document may not be currently available. Refer to the most recent product announcement or contact your local Hitachi Data Systems sales office for information on feature and product availability.

Notice: Hitachi Data Systems products and services can be ordered only under the terms and conditions of Hitachi Data Systems' applicable agreement(s). The use of Hitachi Data Systems products is governed by the terms of your agreement(s) with Hitachi Data Systems.

Hitachi is a registered trademark of Hitachi, Ltd. in the United States and other countries. Hitachi Data Systems is a registered trademark and service mark of Hitachi in the United States and other countries.

All other trademarks, service marks, and company names are properties of their respective owners.



Table of Contents

Preface	vii
Intended audience	viii
Product version	viii
Release notes and readme	viii
Document revision level	viii
Changes in this revision	viii
Document organization	ix
Document conventions	ix
Convention for storage capacity values	x
Accessing product documentation	xi
Getting help	xvi
Comments	xvi
1 Overview	1-1
ShadowImage in-system replication software	1-2
Hardware and software configuration	1-2
How ShadowImage works	1-3
Volume pairs (P-VOLs and S-VOLs)	1-4
Creating pairs	1-6
Initial copy operation	1-6
Automatically split the pair following pair creation	1-7
Splitting pairs	1-7
Re-synchronizing pairs operations	1-8
Normal	1-8
Reverse	1-9
Re-synchronizing for split or split pending pair	1-10
Re-synchronizing for suspended pair	1-10
Suspending pairs operation	1-10
Releasing pairs operation	1-11
Differential Management LUs (DMLU)	1-11
Consistency group (CTG)	1-11

Quick mode	1-12
ShadowImage pair status	1-13
Interfaces for performing ShadowImage operations	1-15
Cascade connection of ShadowImage with TrueCopy	1-16
Cascade restrictions with ShadowImage P-VOL	1-21
LU shared with P-VOL on ShadowImage and P-VOL on TrueCopy	1-21
One LU used for P-VOL on ShadowImage and S-VOL on TrueCopy	1-22
Cascade restrictions with S-VOL of ShadowImage	1-23
Simultaneous cascading restrictions with ShadowImage P-VOL and S-VOL	1-24
Cascade restrictions with data pool of SnapShot	1-24
2 Planning and design.	2-1
The plan and design workflow	2-2
Establishing how often a copy is made (copy frequency)	2-2
Establishing how long a copy is held (copy lifespan)	2-3
Lifespan based on backup requirements	2-3
Lifespan based on business uses	2-3
Establishing the number of copies	2-4
Best practice	2-5
Requirements, recommendations for Logical Units	2-5
RAID configuration for ShadowImage LUs.	2-6
Operating system considerations	2-7
Identifying P-VOL and S-VOL in Windows	2-7
LU mapping with CCI	2-7
AIX	2-7
Microsoft Cluster Server (MSCS).	2-8
Veritas Volume Manager (VxVM)	2-8
Windows 2000	2-8
Windows Server 2003/2008 and ShadowImage configuration	2-8
Linux and LVM configuration	2-9
Concurrent use with Volume Migration	2-9
Concurrent use with Cache Partition Manager	2-9
Concurrent use with Dynamic Provisioning	2-10
Windows Server 2003 and 2008/Windows 2000 and Dynamic Disk	2-15
VMWare and ShadowImage configuration	2-17
Cascading ShadowImage with TrueCopy	2-18
Calculating maximum capacity	2-19
3 Requirements	3-1
System requirements	3-2
Displaying the hardware revision number	3-3
Supported platforms	3-4

4	Install and enable	4-1
	Installing ShadowImage	4-2
	Enabling/disabling ShadowImage	4-3
	Uninstalling ShadowImage	4-4
5	Configuration	5-1
	Configuration workflow	5-2
	Setting Up primary, secondary volumes	5-2
	Setting the DMLU	5-2
	Setting the ShadowImage I/O switching mode	5-3
6	Using ShadowImage	6-1
	ShadowImage workflow	6-2
	Check pair status	6-2
	Create the pair	6-2
	Copy pace	6-3
	Create a ShadowImage pair	6-4
	Split the pair	6-6
	Resync the pair	6-7
	Delete a pair	6-9
	Edit a pair	6-9
	Restore the P-VOL	6-9
	Use the S-VOL for tape backup, testing, reports, etc.	6-10
7	Monitoring and troubleshooting	7-1
	Monitor pair status	7-2
	Troubleshooting	7-3
	Pair failure	7-3
	Path failure	7-3
A	ShadowImage specifications	A-1
	General specifications	A-2
B	Operations using CLI	B-1
	Installing and uninstalling ShadowImage	B-2
	Installing ShadowImage	B-2
	Uninstalling ShadowImage	B-3
	Enabling or disabling ShadowImage	B-3
	Setting the DMLU	B-4
	Setting the ShadowImage I/O switching mode	B-5
	ShadowImage operations	B-5

Confirming pairs status	B-5
Creating ShadowImage pairs	B-6
Splitting ShadowImage pairs	B-7
Re-synchronizing ShadowImage pairs	B-8
Restoring the P-VOL	B-8
Releasing ShadowImage pairs	B-9
Editing pair information	B-9
Creating ShadowImage pairs that belong to a group	B-10
Splitting ShadowImage pairs that belong to a group	B-11
Sample back up script for Windows	B-12
C Operations using CCI	C-1
Setting up CCI	C-2
Setting the command device	C-2
Setting LU mapping	C-3
Defining the configuration definition file	C-3
Setting the environment variable	C-6
ShadowImage operations using CCI	C-8
Confirming pair status	C-8
Creating pairs (paircreate)	C-9
Pair creation using a consistency group	C-10
Splitting pairs (pairsplit)	C-11
Resynchronizing pairs (pairresync)	C-11
Releasing pairs (pairsplit -S)	C-11
Pair, group name differences in CCI and Navigator 2	C-12
D I/O switching mode feature	D-1
I/O switching mode feature	D-2
Specifications	D-3
Recommendations	D-3
Enabling I/O switching mode	D-4
Recovery from a drive failure	D-5

Glossary

Index



Preface

This document provides instructions on assessing your ShadowImage requirements, designing an implementation to meet those requirements, and implementing and operating ShadowImage software using the Storage Navigator 2 graphical user interface.

This preface includes the following information:

- ❑ [Intended audience](#)
- ❑ [Product version](#)
- ❑ [Release notes and readme](#)
- ❑ [Document revision level](#)
- ❑ [Changes in this revision](#)
- ❑ [Document organization](#)
- ❑ [Document conventions](#)
- ❑ [Convention for storage capacity values](#)
- ❑ [Accessing product documentation](#)
- ❑ [Getting help](#)

Intended audience

This document is intended for system administrators, Hitachi Data Systems representatives, and Authorized Service Providers who install, configure, and operate Hitachi Adaptable Modular System (AMS) 2000 family storage systems.

Product version

This document applies to Hitachi AMS 2000 Family firmware version 0890/A or later.

Release notes and readme

Read the release notes and readme file before installing and using this product. They may contain requirements or restrictions that are not fully described in this document and/or updates or corrections to this document.

Document revision level

Revision	Date	Description
MK-97DF8129-01	October 2008	Initial Release
MK-97DF8129-02	December 2008	This revision supersedes revision 01.
MK-97DF8129-03	March 2009	This revision supersedes revision 02.
MK-97DF8129-04	June 2009	This revision supersedes revision 03.
MK-97DF8129-05	August 2009	This revision supersedes revision 04.
MK-97DF8129-06	November 2009	This revision supersedes revision 05.
MK-97DF8129-07	January 2010	This revision supersedes revision 06.
MK-97DF8129-08	April 2010	This revision supersedes revision 07.
MK-97DF8129-09	August 2010	This revision supersedes revision 08.

Changes in this revision

This release includes the following:

- Updated information about the cascade connection in [Concurrent use with Dynamic Provisioning on page 2-10](#).
- Added information about [VMWare and ShadowImage configuration on page 2-17](#).
- Added a procedure about [on page 5-3](#).

Document organization

Thumbnail descriptions of the chapters are provided in the following table. Click the [chapter title](#) in the first column to go to that chapter. The first page of every chapter or appendix contains links to the contents.

Chapter/Appendix Title	Description
Chapter 1, Overview	Provides descriptions of ShadowImage components and how they work together.
Chapter 2, Planning and design	Provides detailed planning and design information.
Chapter 3, Requirements	Provides ShadowImage requirements.
Chapter 4, Install and enable	Provides instructions for enabling ShadowImage.
Chapter 5, Configuration	Provides configuration information.
Chapter 6, Using ShadowImage	Provides information and procedures for using ShadowImage.
Chapter 7, Monitoring and troubleshooting	Provides monitoring and maintenance information.
Appendix A, ShadowImage specifications	Provides ShadowImage specifications.
Appendix B, Operations using CLI	Provides detailed Navigator 2 Command Line Interface instructions for configuring and using ShadowImage.
Appendix C, Operations using CCI	Provides detailed Command Control Interface instructions for configuring and using ShadowImage.
Appendix D, I/O switching mode feature	Host access remains continuous despite drive failure with this feature.
Glossary	Provides definitions for terms and acronyms found in this document.
Index	Provides locations to specific information in this document.

Document conventions

This document uses the following symbols to draw attention to important safety and operational information.

Symbol	Meaning	Description
	Tip	Tips provide helpful information, guidelines, or suggestions for performing tasks more effectively.
	Note	Notes emphasize or supplement important points of the main text.
	Caution	Cautions indicate that failure to take a specified action could result in damage to the software or hardware.

The following typographic conventions are used in this document.

Convention	Description
Bold	Indicates text on a window, other than the window title, including menus, menu options, buttons, fields, and labels. Example: Click OK .
<i>Italic</i>	Indicates a variable, which is a placeholder for actual text provided by the user or system. Example: copy <i>source-file target-file</i> Angled brackets (< >) are also used to indicate variables.
screen/code	Indicates text that is displayed on screen or entered by the user. Example: # <code>pairdisplay -g oradb</code>
< > angled brackets	Indicates a variable, which is a placeholder for actual text provided by the user or system. Example: # <code>pairdisplay -g <group></code> Italic font is also used to indicate variables.
[] square brackets	Indicates optional values. Example: [a b] indicates that you can choose a, b, or nothing.
{ } braces	Indicates required or expected values. Example: { a b } indicates that you must choose either a or b.
vertical bar	Indicates that you have a choice between two or more options or arguments. Examples: [a b] indicates that you can choose a, b, or nothing. { a b } indicates that you must choose either a or b.
underline	Indicates the default value. Example: [<u>a</u> b]

Convention for storage capacity values

Physical storage capacity values (e.g., disk drive capacity) are calculated based on the following values:

Physical capacity unit	Value
1 KB	1,000 bytes
1 MB	1,000 KB or 1,000 ² bytes
1 GB	1,000 MB or 1,000 ³ bytes
1 TB	1,000 GB or 1,000 ⁴ bytes
1 PB	1,000 TB or 1,000 ⁵ bytes
1 EB	1,000 PB or 1,000 ⁶ bytes

Logical storage capacity values (e.g., logical device capacity) are calculated based on the following values:

Logical capacity unit	Value
1 block	512 bytes
1 KB	1,024 (2 ¹⁰) bytes
1 MB	1,024 KB or 1024 ² bytes
1 GB	1,024 MB or 1024 ³ bytes
1 TB	1,024 GB or 1024 ⁴ bytes
1 PB	1,024 TB or 1024 ⁵ bytes
1 EB	1,024 PB or 1024 ⁶ bytes

Accessing product documentation

The AMS 2000 Family user documentation is available on the Hitachi Data Systems Portal: <https://portal.hds.com>. Please check this site for the most current documentation, including important updates that may have been made after the release of the product.

This documentation set consists of the following documents.

Release notes

- Adaptable Modular Storage System Release Notes
- Storage Navigator Modular 2 Release Notes



Please read the release notes before installing and/or using this product. They may contain requirements and/or restrictions not fully described in this document, along with updates and/or corrections to this document.

Installation and getting started

The following documents provide instructions for installing an AMS 2000 Family storage system. They include rack information, safety information, site-preparation instructions, getting-started guides for experienced users, and host connectivity information. The symbol  identifies documents that contain initial configuration information about Hitachi AMS 2000 Family storage systems.

 AMS2100/2300 Getting Started Guide, MK-98DF8152

Provides quick-start instructions for getting an AMS 2100 or AMS 2300 storage system up and running as quickly as possible.

 AMS2500 Getting Started Guide, MK-97DF8032

Provides quick-start instructions for getting an AMS 2500 storage system up and running as quickly as possible.

AMS 2000 Family Site Preparation Guide, MK-98DF8149

Contains site planning and pre-installation information for AMS 2000 Family storage systems, expansion units, and high-density expansion units. This document also covers safety precautions, rack information, and product specifications.

AMS 2000 Family Fibre Channel Host Installation Guide,
MK-08DF8189

Describes how to prepare Hitachi AMS 2000 Family Fibre Channel storage systems for use with host servers running supported operating systems.

AMS 2000 Family iSCSI Host Installation Guide, MK-08DF8188

Describes how to prepare Hitachi AMS 2000 Family iSCSI storage systems for use with host servers running supported operating systems.

Storage and replication features

The following documents describe how to use Storage Navigator Modular 2 (Navigator 2) to perform storage and replication activities.

Storage Navigator 2 Advanced Settings User's Guide, MK-97DF8039

Contains advanced information about launching and using Navigator 2 in various operating systems, IP addresses and port numbers, server certificates and private keys, boot and restore options, outputting configuration information to a file, and collecting diagnostic information.

Storage Navigator Modular 2 User's Guide, MK-99DF8208

Describes how to use Navigator 2 to configure and manage storage on an AMS 2000 Family storage system.

AMS 2000 Family Dynamic Provisioning Configuration Guide, MK-09DF8201

Describes how to use virtual storage capabilities to simplify storage additions and administration.

Storage Navigator 2 Storage Features Reference Guide for AMS, MK-97DF8148

Contains concepts, preparation, and specifications for Account Authentication, Audit Logging, Cache Partition Manager, Cache Residency Manager, Data Retention Utility, LUN Manager, Performance Monitor, SNMP Agent, and Modular Volume Migration.

AMS 2000 Family Copy-on-write SnapShot User Guide, MK-97DF8124

Describes how to create point-in-time copies of data volumes in AMS 2100, AMS 2300, and AMS 2500 storage systems, without impacting host service and performance levels. Snapshot copies are fully read/write compatible with other hosts and can be used for rapid data restores, application testing and development, data mining and warehousing, and nondisruptive backup and maintenance procedures.

AMS 2000 Family ShadowImage In-system Replication User Guide, MK-97DF8129 — this document

Describes how to perform high-speed nondisruptive local mirroring to create a copy of mission-critical data in AMS 2100, AMS 2300, and AMS 2500 storage systems. ShadowImage keeps data RAID-protected and fully recoverable, without affecting service or performance levels. Replicated data volumes can be split from host applications and used for

system backups, application testing, and data mining applications while business continues to operate at full capacity.

AMS 2000 Family TrueCopy Remote Replication User Guide,
MK-97DF8052

Describes how to create and maintain multiple duplicate copies of user data across multiple AMS 2000 Family storage systems to enhance your disaster recovery strategy.

AMS 2000 Family TrueCopy Extended Distance User Guide,
MK-97DF8054

Describes how to perform bi-directional remote data protection that copies data over any distance without interrupting applications, and provides failover and recovery capabilities.

AMS 2000 Data Retention Utility User's Guide, MK-97DF8019

Describes how to lock disk volumes as read-only for a certain period of time to ensure authorized-only access and facilitate immutable, tamper-proof record retention for storage-compliant environments. After data is written, it can be retrieved and read only by authorized applications or users, and cannot be changed or deleted during the specified retention period.

Storage Navigator Modular 2 online help

Provides topic and context-sensitive help information accessed through the Navigator 2 software.

Hardware maintenance and operation

The following documents describe how to operate, maintain, and administer an AMS 2000 Family storage system. They also provide a wide range of technical information and specifications for the AMS 2000 Family storage systems. The symbol  identifies documents that contain initial configuration information about Hitachi AMS 2000 Family storage systems.

 AMS 2100/2300 Storage System Hardware Guide, MK-97DF8010

Provides detailed information about installing, configuring, and maintaining an AMS 2100/2300 storage system.

 AMS 2500 Storage System Hardware Guide, MK-97DF8007

Provides detailed information about installing, configuring, and maintaining an AMS 2500 storage system.

 AMS 2000 Family Storage System Reference Guide,
MK-97DF8008

Contains specifications and technical information about power cables, system parameters, interfaces, logical blocks, RAID levels and configurations, and regulatory information about AMS 2100, AMS 2300, and AMS 2500 storage systems. This document also contains remote adapter specifications and regulatory information.

AMS 2000 Family Storage System Service and Upgrade Guide,
MK-97DF8009

Provides information about servicing and upgrading AMS 2100, AMS 2300, and AMS 2500 storage systems.

AMS 2000 Family Power Savings User Guide, MK-97DF8045

Describes how to spin down volumes in selected RAID groups when they are not being accessed by business applications to decrease energy consumption and significantly reduce the cost of storing and delivering information.

Command and Control (CCI)

The following documents describe how to install the Hitachi AMS 2000 Family Command Control Interface (CCI) and use it to perform TrueCopy and ShadowImage operations.

AMS 2000 Family Command Control Interface (CCI) Installation Guide, MK-97DF8122

Describes how to install CCI software on open-system hosts.

AMS 2000 Family Command Control Interface (CCI) Reference Guide, MK-97DF8121

Contains reference, troubleshooting, and maintenance information related to CCI operations on AMS 2100, AMS 2300, and AMS 2500 storage systems.

AMS 2000 Family Command Control Interface (CCI) User's Guide,
MK-97DF8123

Describes how to use CCI to perform TrueCopy and ShadowImage operations on AMS 2100, AMS 2300, and AMS 2500 storage systems.

Command Line Interface (CLI)

The following documents describe how to use Hitachi Storage Navigator Modular 2 to perform management and replication activities from a command line.

Storage Navigator Modular 2 Command Line Interface (CLI) Unified Reference Guide, MK-97DF8089

Describes how to interact with all Navigator 2 bundled and optional software modules by typing commands at a command line.

Storage Navigator 2 Command Line Interface Replication Reference Guide for AMS, MK-97DF8153

Describes how to interact with Navigator 2 to perform replication activities by typing commands at a command line.

Dynamic Replicator documentation

The following documents describe how to install, configure, and use Hitachi Dynamic Replicator to provide AMS Family storage systems with continuous data protection, remote replication, and application failover in a single, easy-to-deploy and manage platform.

Dynamic Replicator - Scout Release Notes, RN-99DF8211

Dynamic Replicator - Scout Host Administration Guide, MK-98DF8212

Dynamic Replicator - Scout Installation and Configuration Guide, MK-98DF8213

Dynamic Replicator - Scout Quick Start Guide, MK-98DF8214

Dynamic Replicator - Scout Host Troubleshooting Guide, MK-98DF8215

Dynamic Replicator DR-Scout ICAT Utility Guide, MK-98DF8216

Dynamic Replicator - Scout RX Server Deployment Guide, MK-98DF8217

Dynamic Replicator VX Solution for Oracle (Solaris), MK-98DF8218

Dynamic Replicator - Scout Solution for SharePoint 2007, MK-98DF8219

Dynamic Replicator - Scout Solution for MySQL (Windows), MK-98DF8220

Protecting Citrix XenServer Using Hitachi Dynamic Replicator - Scout, MK-98DF8221

Dynamic Replicator Quick Install/Upgrade Guide, MK-98DF8222

Dynamic Replicator - Scout Protecting MS SQL Server, MK-98DF8223

Dynamic Replicator - Scout - Protecting Microsoft Exchange Server, MK-98DF8224

Getting help

If you need to contact the Hitachi Data Systems support center, please provide as much information about the problem as possible, including:

- The circumstances surrounding the error or failure.
- The exact content of any messages displayed on the host system(s).
- The exact content of any messages displayed on Storage Navigator Modular 2.
- The Storage Navigator Modular 2 configuration information. This information is used by service personnel for troubleshooting purposes.

The Hitachi Data Systems customer support staff is available 24 hours a day, seven days a week. If you need technical support, please log on to the Hitachi Data Systems Portal for contact information: <https://portal.hds.com>

Comments

Please send us your comments on this document: doc.comments@hds.com. Include the document title, number, and revision, and refer to specific section(s) and paragraph(s) whenever possible.

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

Overview

ShadowImage™ In-System Replication Software uses local mirroring technology to create a copy of any volume in the AMS array. During copying, host applications can continue to read/write to and from the primary production volume. Replicated data volumes can be split as soon as they are created for use with other applications.

This guide provides instructions for planning and design, configuring, using, and monitoring ShadowImage. In this chapter, the following are discussed:

- ❑ [ShadowImage in-system replication software](#)
- ❑ [Hardware and software configuration](#)
- ❑ [How ShadowImage works](#)
- ❑ [Interfaces for performing ShadowImage operations](#)
- ❑ [Cascade connection of ShadowImage with TrueCopy](#)

ShadowImage in-system replication software

Hitachi's ShadowImage uses local mirroring technology to create full-volume copies within the AMS array. In a ShadowImage pair operation, all data blocks in the original data volume are sequentially copied onto the secondary volume.

The original and secondary data volumes remain synchronized until they are split. While synchronized, updates to the original data volume are continually mirrored to the secondary volume. When the secondary volume is split from the original volume, it contains a mirror image of the original volume at that point in time.

After the pair is split, the secondary volume can be used for offline testing or analytical purposes since there is no common data sharing with the original volume. Since there are no dependencies between the original and secondary volumes, each can be written to by separate hosts. Changes to both volumes are tracked so they can be re-synchronized.

Hardware and software configuration

The typical configuration includes a Hitachi AMS array, a host connected to the array, and software to configure and operate ShadowImage (management software). The host is connected to the array using fibre channel or iSCSI connections. The management software is connected to arrays via a management LAN.

ShadowImage employs a primary volume, a secondary volume or volumes, and the Hitachi Storage Navigator Modular2 (Navigator 2) graphical user interface (GUI). Additional user functionality is made available through Navigator 2 Command-Line Interface (CLI) or Hitachi Command Control Interface (CCI).

[Figure 1-1](#) shows a typical ShadowImage environment.




Figure 1-1: ShadowImage Environment

The following sections describe how these components work together.

How ShadowImage works

ShadowImage creates a duplicate volume of another volume. This volume “pair” is created when you:

- Select a volume that you want to replicate
- Identify another volume that will contain the copy
- Associate the primary and secondary volumes
- Copy all primary volume data to the secondary volume

When the initial copy is made, all data on the P-VOL is copied to the S-VOL. The P-VOL remains available for read/write I/O during the operation. Write operations performed on the P-VOL are always duplicated to the S-VOL.

When the pair is split, the primary volume continues being updated, but data in the secondary volume remains as it was at the time of the split. At this time:

- The secondary volume becomes available for read/write access by secondary host applications.
- Changes to primary and secondary volumes are tracked by differential bitmaps.
- The pair can be made identical again by re-synchronizing changes from primary-to-secondary or secondary-to-primary.

Volume pairs (P-VOLs and S-VOLs)

Each ShadowImage pair can consist of up to 8 secondary volumes (S-VOL). The primary and secondary volumes are located in the same array.

The S-VOL only becomes accessible to a host after the pair is split. When a pair is split, the pair status becomes Split. While a pair is split, the array keeps track of changes to the P-VOL and S-VOL in differential bitmaps.

When the pair is re-synchronized, differential (changed) data in the P-VOL is copied to the S-VOL; then the S-VOL is again identical to the P-VOL. A reverse-resync can also be performed when you want to update the P-VOL with the S-VOL data. In a reverse resync, the differential data in the S-VOL is copied to the P-VOL.

[Figure 1-2](#) illustrates basic operations. [Figure 1-3](#) illustrates pair operations using Storage Navigator Modular 2 GUI.




Figure 1-2: Basic ShadowImage Operations




Figure 1-3: Pair Operations

Creating pairs

The ShadowImage creating pairs operation establishes the newly specified ShadowImage pair(s). The volumes, which will become the P-VOL and S-VOL, must both be in the Simplex state before becoming a ShadowImage pair. Also, when the P-VOL to be an operation target configures a ShadowImage pair in the Paired status, the Paired Internally Synchronizing status, the Synchronizing status, the Reverse Synchronizing status, the Split Pending status, or the Failure status with another S-VOL, a pair cannot be newly created using the Create button.

Initial copy operation

The ShadowImage initial copy operation takes place when you create a new ShadowImage pair, as shown in [Figure 1-4](#). The ShadowImage initial copy operation copies all data on the P-VOL to the associated S-VOL. The P-VOL

remains available to all hosts for read and write I/Os throughout the initial copy operation. Write operations performed on the P-VOL during the initial copy operation will always be duplicated to the S-VOL. The status of the pair is Synchronizing while the initial copy operation is in progress. The pair status changes to Paired when the initial copy is complete.

When creating pairs, you can select the pace for the initial copy operation(s): Slow, Medium, and Fast. The default setting is Medium.




Figure 1-4: Adding a ShadowImage Pair

Automatically split the pair following pair creation

When creating a new ShadowImage pair, you can execute an initial copy in the background and perform pair creation and pair split continuously by specifying Quick Mode. If you execute the command, you will be able to make Read/Write access for the S-VOL immediately and the S-VOL data accessed by the host becomes the same as the P-VOL data at the time of command execution. The pair status under the initial copy is Split Pending. The pair status changes to Split when the initial copy is completed.

Splitting pairs

The ShadowImage splitting pairs operation splits the paired P-VOL and S-VOL, and changes the pair state of the P-VOL and S-VOL to Split.

When splitting pairs is performed the S-VOL becomes identical to the P-VOL and then provides full Read/Write access to the S-VOL.

You can split the pair whose status is Synchronizing or Paired Internally Synchronizing by executing the split operation in Quick Mode. To perform the split operation in Quick Mode for the pair whose status is Synchronizing, specify the option at the time of command execution. Moreover, the split operation for the pair whose status is Paired Internally Synchronizing is executed in Quick Mode regardless of the Quick Mode specification option.

In the split operation in Quick Mode, if you execute the command, you will be able to make Read/Write access for the S-VOL immediately, and the S-VOL data accessed by the host becomes the same as the P-VOL data at the time of command execution. The data to make the S-VOL data same as the P-VOL data at the time of command execution is copied in the background, and the status becomes Split Pending until the copy is completed. The status changes to Split when the copy is completed.

This feature provides point-in-time backup of your data, and also facilitates real data testing by making the ShadowImage copies (S-VOLs) available for host access.

When the split operation is complete, the pair status changes to Split or Split Pending, and you have full Read/Write access to the split S-VOL. While the pair is split, the array establishes a track map for the split P-VOL and S-VOL and records all updates to both volumes. The P-VOL remains fully accessible during splitting pairs operation. Splitting pairs operations cannot be performed on suspended (Failure) pairs.

Re-synchronizing pairs operations

The ShadowImage re-synchronizing pairs operation resynchronizes the split pairs (Split or Split Pending) or the suspended on error pairs (Failure). When re-synchronizing pairs operation starts, the pair status changes to Synchronizing, Reverse Synchronizing, or Paired Internally Synchronizing. The pair status changes to Paired when re-synchronizing pairs copy operations are completed.

ShadowImage allows you to perform two types of re-synchronizing pairs operations:

Normal

The normal re-synchronizing pairs operation (see [Figure 1-5](#)) resynchronizes the S-VOL with the P-VOL. The copy direction for a normal re-synchronizing pairs operation is P-VOL to S-VOL. The pair status during a normal re-synchronizing copy operation is Synchronizing. The S-VOL becomes inaccessible to all hosts for write operation, and the P-VOL remains accessible to all hosts for both read and write operations during a normal re-synchronizing pairs. The normal re-synchronizing pairs operation can be executed for pairs with a status of Split, Split Pending, and Failure. However, when the P-VOL that will be an operation target configures a ShadowImage pair in the Paired status, the Paired Internally Synchronizing status, the Synchronizing status, and the Split Pending status or the Failure status with another S-VOL, re-synchronizing pairs operation cannot be executed.

Moreover, if you specify Quick Mode when executing the re-synchronizing operation, the pair status becomes Paired Internally Synchronizing. Although the P-VOL data is copied to the S-VOL just like the pair in the Synchronizing status, you can execute the split operation without specifying Quick Mode option for the pair in the Paired Internally Synchronizing status.

When re-synchronizing pairs operation is executed for a pair in the Failure (S-VOL Switch) status, the restore (reverse resynchronization) operation is performed. When operating restore, you cannot specify Quick Mode.

Reverse

Restore pairs operation (see [Figure 1-6](#)) synchronizes the P-VOL with the S-VOL. However, when the P-VOL that will be an operation target configures a ShadowImage pair in the Paired status, the Paired Internally Synchronizing status, the Synchronizing status, the Reverse Synchronizing status, the Split Pending status or the Failure status with another S-VOL, the restore operation cannot be executed. The copy direction for a restore pairs operation is S-VOL to P-VOL. The pair status during a restore operation is Reverse Synchronizing, and the S-VOL becomes inaccessible to all hosts for write operations during a restore pairs operation. The P-VOL remains accessible for both read and write operations, and the write operation on P-VOL will always be reflected to S-VOL (see [Figure 1-7](#)). When operating restore, you cannot specify Quick Mode.




Figure 1-5: Normal Re-synchronizing Pairs Operation




Figure 1-6: Reverse Re-synchronizing Pairs Operation




Figure 1-7: Reflecting Write Data to S-VOL during Reverse Re-synchronizing Pairs Operation (Restore)

ShadowImage allows you to perform re-synchronizing operations on split, Split Pending, and Failure pairs.

Re-synchronizing for split or split pending pair

When a re-synchronizing pairs operation is performed on a split pair or split pending (status = Split or Split Pending), the array merges the S-VOL differential track map into the P-VOL track differential map and copies all flagged data from the P-VOL to the S-VOL. When a reverse re-synchronizing pairs operation is performed on a split pair, the array merges the P-VOL differential track map into the S-VOL differential track map and then copies all flagged tracks from the S-VOL to the P-VOL. This ensures that the P-VOL and S-VOL are properly resynchronized in the desired direction, and also greatly reduces the time needed to resynchronize the pair.

Re-synchronizing for suspended pair

When a re-synchronizing pairs operation is performed on a suspended pair (status = Failure), the array copies all data on the P-VOL to the S-VOL, since all P-VOL tracks were flagged as difference data when the pair was suspended. It takes the same time as the initial copy of ShadowImage until the copy for resynchronization is completed and the status changes to Paired.

Suspending pairs operation

The ShadowImage suspending pairs operation (split pair with **Suspend operation in progress and force the pair in failure state**) immediately suspends the ShadowImage copy operations to the S-VOL of the pair. The user can suspend a ShadowImage pair at any time. When a ShadowImage pair is suspended on error (status = Failure), the array stops performing ShadowImage copy operations to the S-VOL, continues accepting write I/O operations to the P-VOL, and marks the entire P-VOL track map as

difference data. When a re-synchronizing pairs operation is performed on a suspended pair, the entire P-VOL is copied to the S-VOL (when a restore operation is performed, the entire S-VOL is copied to the P-VOL). While re-synchronizing pairs operation for a split or split pending ShadowImage pair greatly reduces the time needed to resynchronize the pair, re-synchronizing pairs operation for a suspended on error pair will take as long as the initial copy operation.

The array will automatically suspend a ShadowImage pair when copy operation cannot be continued or cannot keep the pair mirrored for any reason. When the array suspends a pair, a file is output to the system log or event log to notify the host (CCI only). The array will automatically suspend a pair under the following conditions:

- When the ShadowImage volume pair has been suspended or deleted from the UNIX[®]/PC host using the CCI.
- When the array detects an error condition related to an initial copy operation. When a volume pair with Synchronizing status is suspended on error, the array aborts the initial copy operation, changes the status of the P-VOL and S-VOL to Failure and accepts all subsequent write I/Os to the P-VOL.

Releasing pairs operation

The ShadowImage releasing pairs operation stops the ShadowImage copy operations to the S-VOL of the pair and releases the volume in paired status. The user can release a ShadowImage pair at any time except when the volumes are already in Simplex or Split Pending status. In both ShadowImage volumes, the status will change to Simplex.

Differential Management LUs (DMLU)

The DMLU is an exclusive volume used for storing ShadowImage information when the array is powered down. The DMLU is treated the same as other volumes in the storage array, but is hidden from a host. See [Setting the DMLU on page 5-2](#) to configure.

Consistency group (CTG)

Application data often spans more than one volume. With ShadowImage, it is possible to manage operations spanning multiple volumes as a single group. In a “consistency group” (CTG), all primary logical volumes are treated as a single entity.

Managing ShadowImage primary volumes as a consistency group allows multiple operations to be performed on grouped volumes concurrently. Write order is guaranteed across application logical volumes, since pairs can be split at the same time.

Quick mode

If you use the Quick Mode for creating or updating the copying volume (S-VOL) by ShadowImage, the Read/Write access for the S-VOL becomes available immediately. Since the S-VOL data accessed from the host becomes the same as the P-VOL data at the time when the command was executed, you can start the backup from the S-VOL without waiting for the completion of the data copy.

Table 1-1: Quick Mode Characteristics

	Advantages	Disadvantages
With Quick Mode	Access from the host to the S-VOL becomes possible without waiting for data copy completion while creating or updating the copying volume.	Since the P-VOL data is used when accessing the S-VOL, the load for the S-VOL affects the I/O performance of the P-VOL. The I/O performance of the S-VOL is affected by the load of the P-VOL and is more deteriorated than the case where Quick Mode is not used.
Without Quick Mode	Since access from the host to the S-VOL is independent of the P-VOL, the I/O performance is less affected.	Access to the S-VOL cannot begin unless the data copy is completed while creating or updating the copying volume. Therefore, wait for data copy completion before accessing the S-VOL.
HP	HP-UX 11i V1.0 (PA-RISC)	
	HP-UX 11i V2.0 (PA-RISC)	
	HP-UX 11i V3.0 (PA-RISC)	
	HP-UX 11i V2.0 (IPF)	
	HP-UX 11i V3.0 (IPF)	
	Tru64 UNIX 5.1	
IBM®	AIX 5.1	
	AIX 5.2	
	AIX 5.3	
SGI	IRIX 6.5.x	

ShadowImage pair status

The ShadowImage displays the pair status of all ShadowImage volumes (LUs). [Figure 1-8](#) shows the ShadowImage pair status transitions and the relationship between the pair status and the ShadowImage operations.




Figure 1-8: ShadowImage Pair Status Transitions

Table 1-2 lists and describes the ShadowImage pair status conditions. If a volume is not assigned to a ShadowImage pair, its status is Simplex.

When you create a ShadowImage pair not specifying Quick Mode, the status of the P-VOL and S-VOL changes to Synchronizing. When the initial copy operation is complete, the pair status becomes Paired. When specifying Quick Mode, if the ShadowImage pair creation starts in the pair creation operation, the pair status becomes Split Pending. In this status, the initial copy is in progress in the background. When the initial copy operation is completed, the pair status changes to Split.

If the array cannot maintain the data copy for any reason or if you suspend the pair, the pair status changes to Failure. When you split a pair, the pair status changes to Synchronizing.

When splitting pairs operation is complete, the pair status changes to Split or Split Pending to enable you to access the split S-VOL. When you start a re-synchronizing pairs operation, the pair status changes to Synchronizing or Paired Internally Synchronizing. When you specify reverse mode for a re-synchronizing pairs operation (restore), the pair status changes to Reverse Synchronizing (data is copied in the reverse direction from the S-VOL to the P-VOL). When re-synchronizing pairs operation is complete, the pair status changes to Paired. When you release a pair, the pair status changes to Simplex.

Table 1-2: ShadowImage Pair Status

Pair Status	Description	S-VOL Access
Simplex	The volume is not assigned to a ShadowImage pair. The array accepts read and write I/Os for all Simplex volumes.	N/A (There is no S-VOL yet).
Synchronizing	Creating a pair or re-synchronizing a pair, the copy operation is in progress. The array continues to accept read and write operations for the P-VOL but does not accept write operations for the S-VOL. When a split pair is resynchronized in normal mode, the array copies only the P-VOL differential data to the S-VOL. When creating a pair or a Failure pair is resynchronized, the array copies the entire P-VOL to the S-VOL.	Read only.
Paired Internally Synchronizing	The copy operation is in progress as same as Synchronizing. The P-VOL and the S-VOL are not yet the same. The pair split in the Paired Internally Synchronizing status operates in the Quick Mode even without specifying the option and changes to Split Pending.	Read only.
Paired	The copy operation is complete, and the array starts copying the write operation taken to the P-VOL data onto the S-VOL. The P-VOL and S-VOL of a duplex pair (Paired status) is identical. The array rejects all write I/Os for S-VOLs with the status Paired.	Read only.
Split	The array starts accepting write I/Os for Split S-VOLs. The array keeps track of all updates to the split P-VOL and S-VOL so that the pair can be resynchronized quickly.	Read and write. the S-VOL can be mounted.
Split Pending	Although the array starts accepting the I/O operation of Write for the S-VOL in the Split Pending status, the data copy from the P-VOL to the S-VOL is in progress in the background. The array records the positions of all updates to the split P-VOL and S-VOL. You cannot delete the pair in the Split Pending status.	Read and write. the S-VOL can be mounted.
Reverse Synchronizing	The array does not accept write I/Os for Reverse Synchronizing S-VOLs. When a split pair is resynchronized in reverse mode, the array copies only the S-VOL differential data to the P-VOL.	Read only.
Failure	The array continues accepting read and write I/Os for a Failure (suspended under error) P-VOL (however, if the status transits from Reverse Synchronizing, all access to P-VOL is disabled). The array marks the entire P-VOL track map as difference data, so that the entire P-VOL is copied to the S-VOL when the Failure pair is resumed. Use re-synchronizing pairs operation to resume a Failure pair.	Read only. (If the status transits from Reverse Synchronizing, all access to S-VOL is disabled.)

Interfaces for performing ShadowImage operations

ShadowImage can be operated using of the following interfaces:

- The GUI (Hitachi Storage Navigator Modular 2 Graphical User Interface) is a browser-based interface from which ShadowImage can be setup, operated, and monitored. The GUI provides the simplest method for performing operations, requiring no previous experience. Scripting is not available.
- CLI (Hitachi Storage Navigator Modular 2 Command Line Interface), from which ShadowImage can be setup and all basic pair operations can be performed—create, split, resynchronize, restore, swap, and delete. The GUI also provides these functionalities. CLI also has scripting capability.
- CCI (Hitachi Command Control Interface, used to display volume information and perform all copying and pair-managing operations. CCI provides a full scripting capability which can be used to automate replication operations. CCI requires more experience than the GUI or CLI. CCI is required on Windows 2000 Server for performing mount/unmount operations.

HDS recommends using the GUI to begin operations for new users with no experience with CLI or CCI. Users who are new to replication software but have CLI experience in managing arrays may want to continue using CLI, though the GUI is an option. The same recommendation applies to CCI users.



NOTE: Hitachi Replication Manager can be used to manage and integrate ShadowImage. It provides a GUI representation of the ShadowImage system, with monitoring, scheduling, and alert functions. For more information, visit the Hitachi Data Systems website.




CAUTION! Storage Navigator 2 CLI is provided for users with significant storage management expertise. Improper use of this CLI could void your Hitachi warranty. Please consult with your reseller before using CLI.

Cascade connection of ShadowImage with TrueCopy

ShadowImage volumes can be cascaded with those of TrueCopy, as shown in the Figure 1-9 through Figure 1-13. ShadowImage P-VOL volumes can be cascaded with those of TrueCopy in the Split Pending or Paired Internally Split status. ShadowImage S-VOL volumes cannot be cascaded with those of TrueCopy in the Split Pending or Paired Internally Split status.

ShadowImage cannot be cascaded with ShadowImage or SnapShot.

■ Cascade with a P-VOL of ShadowImage



■ Cascade with an S-VOL of ShadowImage




Figure 1-9: Cascade Connection of ShadowImage with TrueCopy (P-VOL: S-VOL=1: 1)

- Cascade with a P-VOL and an S-VOL of ShadowImage




Figure 1-10: Simultaneous Cascade Connection of ShadowImage with TrueCopy (P-VOL: S-VOL=1: 1)

■ Cascade with a P-VOL of ShadowImage




Figure 1-11: Cascade Connection of ShadowImage with TrueCopy (P-VOL: S-VOL=1: 3)

■ Cascade with an S-VOL of ShadowImage




Figure 1-12: Cascade Connection of ShadowImage with S-VOL

- Cascade with a P-VOL and an S-VOL of ShadowImage




Figure 1-13: Simultaneous Cascade Connection of ShadowImage with TrueCopy (P-VOL: S-VOL=1: 3)

Cascade restrictions with ShadowImage P-VOL

If by using TrueCopy, the P-VOL on ShadowImage is cascaded, there is a possibility that the data of P-VOL on the local side does not equal the data of P-VOL on the remote side due to restore being executed using ShadowImage. To counteract this, the restriction that enables restores using ShadowImage to be executed in the *Split* status only is provided on TrueCopy. When restores using ShadowImage are executed, TrueCopy must be in the *Split* status. If restore using ShadowImage is executed in the *Synchronizing* status or *Paired* status of TrueCopy, the data in the LUs for P-VOL that are cascaded using TrueCopy on the local side and the remote side cannot be assured of being equal.

LU shared with P-VOL on ShadowImage and P-VOL on TrueCopy

Table 1-3 shows whether or not a read/write from/to a P-VOL of ShadowImage on the local side is possible in the case where a P-VOL of ShadowImage and a P-VOL of TrueCopy are the same LU.

Table 1-3: A Read/Write Instruction to a ShadowImage P-VOL on the Local Side

TrueCopy P-VOL		ShadowImage P-VOL						
		Paired (Including Paired Internally Synchronizing)	Synchronizing	Synchronizing (Restore)		Split (Including Split Pending)	Failure	Failure (Restore)
				R/W	W			
Paired		+, R/W	+, R/W	–	–	+, R/W	+, R/W	–
Synchronizing		+, R/W	+, R/W	–	–	+, R/W	+, R/W	–
Split	R/W	+, R/W	+, R/W	+, R/W	+, W	+, R/W	+, R/W	Δ , R/W
	R	+, R	+, R	–	–	+, R	+, R	–
Failure	R/W	+, R/W	+, R/W	Δ , R/W	Δ , W	+, R/W	+, R/W	Δ , R/W
	R	+, R	+, R	–	–	+, R	Δ , R	–
	R/W	+, R/W	+, R/W	–	–	+, R/W	Δ , R/W	–

+ indicates a possible case, – indicates an impossible case

Δ indicates a case where a pair operation causes an error (a case that can occur as a result of a change of the pair status to Failure)

R/W: Read/Write by a host is possible.

R: Read by a host is possible but write is impossible.

W: Write by a host is possible but read is impossible.

R/W: Read/Write by a host is impossible.



NOTE: *Failure* in this table excludes any condition where LU access is not possible (for example, LU blockage).

When one P-VOL configures a pair with one or more S-VOLs, decide which item is applied as the pair status of the P-VOL of the above-mentioned ShadowImage with the following procedure.

1. If all the pairs that the P-VOL concerned configures are in the *Split* status, the item of *Split* is applied.
2. If all the pairs that the P-VOL concerned configures are in the *Split* status or the *Failure* status, the item of *Split* is applied. However, when including the pair that became *Failure* during restore, the items of *Failure* (Restore) are applied.
3. If a pair in the *Paired* status, the *Synchronizing* status, or the *Reverse Synchronizing* status is included in the pair that the P-VOL concerned configures, the item of *Paired*, *Synchronizing*, and *Reverse Synchronizing* is applied, respectively. (Two or more pairs in the *Paired* status, the *Synchronizing* status, and the *Reverse Synchronizing* status are never included in the pair that the P-VOL concerned configures.)

One LU used for P-VOL on ShadowImage and S-VOL on TrueCopy

Table 1-4 shows whether a read/write from/to a P-VOL of ShadowImage on the remote side is possible or not in the case where a P-VOL of ShadowImage and an S-VOL of TrueCopy are the same LU.

Table 1-4: A Read/Write Instruction to a ShadowImage P-VOL on the Remote Side

TrueCopy P-VOL		ShadowImage P-VOL						
		Paired (Including Paired Internally Synchronizing)	Synchronizing	Synchronizing (Restore)		Split (Including Split Pending)	Failure	Failure (Restore)
				R/W	W			
Paired		+, R	+, R	–	–	+, R	+, R	–
Synchronizing		+, R	+, R	–	–	+, R	+, R	–
Split	R/W	+, R/W	+, R/W	+, R/W	+, W	+, R/W	+, R/W	Δ , R/W
	R	+, R	+, R	–	–	+, R	+, R	–
Failure		+, R	+, R	–	–	+, R	Δ , R	–

+ indicates a possible case, – indicates an impossible case

Δ indicates a case where a pair operation causes an error (a case that can occur as a result of a change of the pair status to Failure)

R/W: Read/Write by a host is possible.

R: Read by a host is possible but write is impossible.

W: Write by a host is possible but read is impossible.

R/W: Read/Write by a host is impossible.



NOTE: *Failure* in this table excludes any condition where LU access is not possible (for example, LU blockage).

When one P-VOL configures a pair with one or more S-VOLs, decide which item is applied as the pair status of the P-VOL of the above-mentioned ShadowImage with the following procedure.

1. If all the pairs that the P-VOL concerned configures are in the *Split* status, the item of *Split* is applied.
2. If all the pairs that the P-VOL concerned configures are in the *Split* status or the *Failure* status, the item of *Split* is applied. However, when including the pair that became *Failure* during restore, the items of *Failure* (Restore) are applied.
3. If a pair is in the *Paired* status, the *Synchronizing* status, or the *Reverse Synchronizing* status is included in the pair that the P-VOL concerned configures, the item of *Paired*, *Synchronizing*, and *Reverse Synchronizing* is applied, respectively. (Two or more pairs in the *Paired* status, the *Synchronizing* status, and the *Reverse Synchronizing* status are never included in the pair that the P-VOL concerned configures.)

Cascade restrictions with S-VOL of ShadowImage

Cascade of a TrueCopy LU with a ShadowImage S-VOL is supported only when the ShadowImage S-VOL and TrueCopy P-VOL are the same LU. Also, the operations of the ShadowImage and TrueCopy pairs are restricted depending on the status of the pairs.

Create a ShadowImage pair first when cascading TrueCopy volumes with an S-VOL of the ShadowImage. When a TrueCopy pair is created earlier, split the TrueCopy pair once and create a pair using ShadowImage.

When changing the status of a ShadowImage pair, the status of a TrueCopy pair must be *Split* or *Failure*. When changing the status of a TrueCopy pair, the status of a ShadowImage pair must be *Split*.

Table 1-5 shows whether a read/write to/from a ShadowImage S-VOL on the local side is possible in the case where a ShadowImage S-VOL and a TrueCopy P-VOL are the same LU.

Table 1-5: A Read/Write Instruction to a ShadowImage S-VOL on the Local Side

TrueCopy P-VOL		ShadowImage S-VOL						
		Paired (Including Paired Internally Synchronizing)	Synchron-izing	Synchronizing (Restore)	Split	Split Pending	Failure	Failure (Restore)
Paired		–	–	–	+, R/W	–	–	–
Synchronizing		–	–	–	+, R/W	–	–	–
Split	R/W	+, R	+, R	+, R	+, R/W	+, R/W	Δ, R	Δ, R/W
	R	+, R	+, R	+, R	+, R	+, R	Δ, R	Δ, R/W

Table 1-5: A Read/Write Instruction to a ShadowImage S-VOL on the Local Side

TrueCopy P-VOL		ShadowImage S-VOL						
		Paired (Including Paired Internally Synchronizing)	Synchron- izing	Synchronizing (Restore)	Split	Split Pending	Failure	Failure (Restore)
Failure	R/W	+, R	+, R	+, R	+, R/W	+, R/W	Δ, R	Δ, R/W
	R	+, R	+, R	+, R	+, R	+, R	Δ, R	Δ, R/W
	R/W	+, R/W	+, R/W	+, R/W	+, R/W	+, R/W	Δ, R/W	Δ, R/W

+ indicates a possible case, – indicates an impossible case

Δ indicates a case where a pair operation causes an error (a case that can occur as a result of a change of the pair status to Failure)

R/W: Read/Write by a host is possible.

R: Read by a host is possible but write is impossible.

W: Write by a host is possible but read is impossible.

R/W: Read/Write by a host is impossible.



NOTE: *Failure* in this table excludes any condition where LU access is not possible (for example, LU blockage).

Simultaneous cascading restrictions with ShadowImage P-VOL and S-VOL

When operating a ShadowImage pair where both ShadowImage P-VOL and S-VOL are cascaded with TrueCopy volumes, restrictions are placed on the operation if any restrictions are applied to cascade either ShadowImage P-VOL or S-VOL. To operate the TrueCopy pair, restrictions on the respective cascade configuration are applied depending on which ShadowImage P-VOL and S-VOL is the cascaded volume.

When creating a TrueCopy pair that is to be cascaded with the ShadowImage S-VOL, create the ShadowImage pair first. When changing the status of the TrueCopy pair that is cascaded with the ShadowImage S-VOL, it is required that the ShadowImage pair status is *Split*.

When changing the status of the ShadowImage pair it is required that the status of the TrueCopy pair, which is cascaded with the ShadowImage S-VOL, is *Split* or *Failure*. Furthermore, when executing the restore of ShadowImage, the TrueCopy pair cascaded with the ShadowImage P-VOL in the *Split* status must change.

Cascade restrictions with data pool of SnapShot

A ShadowImage pair cannot be created using a data pool.

Planning and design

With ShadowImage you create copies of your production data so that it can be used to restore the P-VOL or used for tape backup, development, data warehousing, and so on.

This chapter guides you in planning a design that meets your organizational business requirements.

- ❑ [The plan and design workflow](#)
- ❑ [Establishing how often a copy is made \(copy frequency\)](#)
- ❑ [Establishing how long a copy is held \(copy lifespan\)](#)
- ❑ [Establishing the number of copies](#)
- ❑ [Requirements, recommendations for Logical Units](#)
- ❑ [Cascading ShadowImage with TrueCopy](#)
- ❑ [Calculating maximum capacity](#)

The plan and design workflow

A ShadowImage system can only be successful when your business needs are assessed. Business needs determine your ShadowImage copy frequency, lifespan, and number of copies.

- Copy frequency means how often a P-VOL is copied.
- Copy lifespan means how long the copy is held before it is updated.
- Knowing the frequency and lifespan help you determine the number of copies that are required.

These objectives are addressed in detail in this chapter. Three additional tasks are required before your design can be implemented, which are also addressed in this chapter.

- The primary and secondary logical volumes must be set up. Recommendations and supported configurations are provided.
- The ShadowImage maximum capacity must be calculated and compared to the array maximum supported capacity. This has to do with how the array manages storage segments.
- Equally important in the planning process are the ways that various host operating systems interact with ShadowImage. Make sure to review the information at the end of the chapter.

Establishing how often a copy is made (copy frequency)

How often copies are made is determined by how much data could be lost in a disaster before business is significantly impacted.

Ideally, a business desires no data loss. In the real world, disasters occur and data is lost. You or your organization's decision makers must decide the number of business transactions that could be lost, the number of hours required to key in lost data, and so on to decide how often copies must be made.

For example, if losing 4 hours of business transaction could be tolerated, but not more, then copies should be planned for every 4 hours. If 24 hours of business transaction could be lost, copies should be planned every 24 hours.

Figure 2-1 illustrates copy frequency.




Figure 2-1: Copy Frequency

Establishing how long a copy is held (copy lifespan)

Copy lifespan is the length of time a copy (S-VOL) is held, before a new backup is made to the volume. Lifespan is determined by two factors:

- Your organization's data retention policy for holding onto backup copies
- Secondary business uses of the backup data

Lifespan based on backup requirements

- If the copy is to be used for tape backups, the minimum lifespan must be => the copy time from S-VOL to tape. For example:

Hours to copy an S-VOL to tape = 4

Therefore, S-VOL lifespan = 4 hours

- If the copy is to be used as a disk-based backup available for online recovery, you can determine the lifespan by multiplying the number of copies you will require to keep online times the copy frequency. For example:

Copies held = 4

Copy frequency = 4 hours

$4 \times 4 = 16$ hours

S-VOL lifespan = 16 hours

Lifespan based on business uses

- If copy data is used for testing an application, the testing requirements determine the amount of time a copy is held.
- If copy data is used for development purposes, development requirements determine the time the copy is held.
- If copy data is used for business reports, the reporting requirements determine the backup's lifespan.

Figure 2-2 illustrates copy lifespan.




Figure 2-2: Copy Lifespan

Establishing the number of copies

Data retention and business-use requirements for the secondary volume determine a copy's lifespan. They also determine the number of S-VOLs needed per P-VOL.

For example: If your data must be backed up every 12 hours, and business-use of secondary volume data requires holding it 36 hours, then your ShadowImage system requires 3 S-VOLs. This is illustrated in [Figure 2-3](#).




Figure 2-3: Number of S-VOLs Required

Best practice

Hitachi recommends setting up at least two S-VOLs per P-VOL. When an S-VOL is re-synchronizing with the P-VOL, it is in an inconsistent state and therefore not usable. Thus, if at least two S-VOLs exist, one is always available for restoring the P-VOL in an emergency.

A workaround when employing one S-VOL only is to backup the S-VOL to tape. However, this operation can be lengthy, and recovery time from tape is more time-consuming than from an S-VOL. Also, if a failure occurs during the updating of the copy, both the P-VOL and the single S-VOL are invalid.

Requirements, recommendations for Logical Units

This section relates mostly to primary and secondary volumes. However, recommendations for the DMLU and command device are also included.

Please review the following key requirements and recommendations.

Also, review the information on setting up LUs for ShadowImage volumes, DMLUs, and command devices in:

- [Chapter 3, Requirements](#)
- [Appendix A, ShadowImage specifications](#)

When preparing for ShadowImage, please observe the following regarding the P-VOL and S-VOL:

- They must be the same in size, with identical block counts. You can verify block size. In the Navigator 2 GUI, navigate to Groups/RAID Groups/Logical Units tab. Click the desired LUN. On the popup window that appears, review the Capacity field. This shows block size.
- They must be assigned to the same controller.
- Use SAS or SSD drives to increase performance. Performance is lower when SATA drives are used.
- Assign four or more disks to the data disks.
- Volumes used for other purposes should not be assigned as a primary volume. If such a volume must be assigned, move as much of the existing write workload to non-ShadowImage volumes as possible.
- When locating multiple P-VOLs in the same parity group, performance is best when the status of their pairs are the same (Split, Paired, Resync, and so on).

RAID configuration for ShadowImage LUs

Please observe the following regarding RAID levels when setting up ShadowImage pairs and Differential Management LUs.

- Volumes should be assigned to different RAID groups on the array to reduce I/O impact.
- If assigned to the same RAID group, limit the number of pairs in the group to reduce impact on performance.
- Avoid locating P-VOLs and S-VOLs within the same ECC group of the same RAID group for the following reasons:
 - A single drive failure causes status degeneration in the P-VOL and S-VOL.
 - Initial copy, coupling, and resync processes incurs a drive bottleneck which decreases performance.
- A RAID level with redundancy is recommended for both P-VOLs and S-VOLs.
- Redundancy for the P-VOL should be the same as the redundancy for the S-VOL.
- The recommended RAID configuration for P-VOLs and S-VOLs is RAID 5 (4D+1)/RAID 5 (4D+1).
- When two or more DMLUs or command devices (when using CCI) are set within the one AMS array, assign them to the respective RAID groups for redundancy.

Operating system considerations

This section describes the restrictions that apply to ShadowImage volumes.

Identifying P-VOL and S-VOL in Windows

In Navigator 2, the P-VOL and S-VOL are identified by their LU number. In Windows, LUs are identified by HLUN. To map LUN to HLUN on Windows, proceed as follows. These instructions provide procedures for the fibre channel and iSCSI interfaces.

1. Identify the HLUN of your Windows disk.
 - a. From the Windows Server 2003 Control Panel, select Computer Management/Disk Administrator.
 - b. Right-click the disk whose HLUN you want to know, then select **Properties**. The number displayed to the right of "LUN" in the dialog window is the HLUN.
2. Identify HLUN-to-LUN Mapping for the iSCSI interface as follows. (If using Fibre Channel, skip to Step 3.)
 - a. In the Navigator 2 GUI, select the desired array.
 - b. In the array tree that displays, click the **Group** icon, then click the **iSCSI Targets** icon in the Groups tree.
 - c. On the iSCSI Target screen, select an iSCSI target.
 - d. On the target screen, select the **Logical Units** tab. Find the identified HLUN. The LUN displays in the next column.
 - e. If the HLUN is not present on a target screen, on the iSCSI Target screen, select another iSCSI target and repeat Step d.
3. Identify HLUN-to-LUN Mapping for the Fibre Channel interface, as follows:
 - a. In the Navigator 2 GUI, select the desired array.
 - b. In the array tree, click the **Group** icon, then click the **Host Groups**.
 - c. Click the host group that the volume is mapped to.
 - d. On the screen for the host group, click the **Logical Units** tab. The volumes mapped to the Host Group display. You can confirm the LUN is mapped to the H-LUN.

LU mapping with CCI

The P-VOL and S-VOL must be mapped to a port that is specified in the configuration definition file in order to be paired. Use LUN Manager when you want to hide P-VOLs and S-VOLs from a host.

AIX

To ensure that the same host recognizes both a P-VOL and an S-VOL, version 04-00-/B or later of HDLM (JP1/HiCommand Dynamic Link Manager) is required.

Microsoft Cluster Server (MSCS)

To make an S-VOL recognized by a host, observe the following:

- Use the CCI `mount` command. Do not use Disk Administrator.
- Do not place the MSCS Quorum Disk in CCI.
- Shutdown MSCS before executing the CCI `sync` command.
- Map the S-VOL to different host than the P-VOL's host. A host cannot recognize a P-VOL and its S-VOL at the same time.
- The command device cannot be shared between the different hosts in the cluster.
- Assign the exclusive command device to each host.

Veritas Volume Manager (VxVM)

A host cannot recognize both a P-VOL and its S-VOL at the same time. Map the P-VOL and S-VOL to separate hosts.

Windows 2000

- A host cannot recognize both a P-VOL and its S-VOL at the same time. Map the P-VOL and S-VOL to separate hosts.
- When mounting a volume, you must use the CCI `mount` command, even if you are operating the pairs using Navigator 2 GUI or CLI. Do not use the Windows `mountvol` command because the data residing in server memory is not flushed. The CCI `mount` command flushes data in server memory, which is necessary for ShadowImage operations. For more information, see the *Hitachi Adaptable Modular Storage Command Control Interface (CCI) Reference Guide*.

Windows Server 2003/2008 and ShadowImage configuration

- A host cannot recognize both a P-VOL and its S-VOL at the same time. Map the P-VOL and S-VOL to separate hosts. When mounting volumes, use either the CCI `mount` command or Windows `mountvol /P` command. If using Windows' `mountvol`, the command specification should be understood and sufficient testing performed before including it in your operation. When using the CCI `mount` command, you must use the CCI `umount` to unmount. For more information on CCI commands, see the *Hitachi Adaptable Modular Storage Command Control Interface (CCI) Reference Guide*.
- In Windows Server™ 2008, use the CCI `umount` command to flush the data on the server memory at the time of the unmount. Do not use the standard Windows `mountvol` command. Moreover, do not use the directory mount at the time of the mount and only use the mount by the drive letters. Refer to the *Hitachi Adaptable Modular Storage Command Control Interface (CCI) Reference Guide* for detail of the restrictions of Windows Server™ 2008 when using the mount/umount command.
- In Windows Server™ 2008, set only the P-VOL to be recognized by the host and let another host recognize the S-VOL. When you have created

two or more S-VOLs for one P-VOL, do not make the same host recognize two or more S-VOLs (shared P-VOLs) at the same time.

- When CCI is used to operate pairs, specify a command device in the configuration definition file as Volume{GUID}.
- When a path detachment, which is caused by a controller detachment or interface failure, continues for longer than one minute, the command device may be unable to be recognized at the time when recovery from the path detachment is made. To make the recovery, execute the “re-scanning of the disks” of Windows. When Windows cannot access the command device although CCI is able to recognize the command device, restart CCI. Windows may write for the un-mounted volume. If a pair is resynchronized while retaining S-VOL data on the server memory, the compatible backup cannot be collected. Therefore, execute the CCI sync command immediately before re-synchronizing the pair for the un-mounted S-VOL.

Linux and LVM configuration

A host cannot recognize both a P-VOL and its S-VOL at the same time. Map the P-VOL and S-VOL to separate hosts.

Concurrent use with Volume Migration

The maximum number of pairs and copying operations are limited when ShadowImage is used with Volume Migration, as follows:

- Number of ShadowImage and Volume Migration pairs:
 - AMS2500/AMS2300: 2,047
 - AMS2100: 1,023
- The number of copying operations that can be performed at the same time in the background is called *copying multiplicity*. Volume Migration and ShadowImage share copying multiplicity, which is limited to four per controller for AMS2300/2100, and eight per controller for AMS2500.

AMS executes copying operations for ShadowImage and Volume Migration in sequential order. There may be times when copying is not started immediately when previously issued copy operations are being performed.

Concurrent use with Cache Partition Manager

ShadowImage can be used with Cache Partition Manager. See the section on restrictions in the *Hitachi Storage Navigator Modular 2 Storage Features Reference Guide* for more information.

Concurrent use with Dynamic Provisioning

This section describes things to keep in mind when using ShadowImage and Dynamic Provisioning together. Refer to the Hitachi Adaptable Modular Storage Dynamic Provisioning User's Guide for detailed information regarding Dynamic Provisioning. The LU created in the RAID group is called a normal LU and the LU created in the DP pool created by Dynamic Provisioning is called a DP-VOL.

- LU type that can be set for a P-VOL or an S-VOL of ShadowImage
The DP-VOL created by Dynamic Provisioning can be used for a P-VOL or an S-VOL of ShadowImage. [Table 2-1](#) shows a combination of a DP-VOL and a normal LU that can be used for a P-VOL or an S-VOL of ShadowImage. When using the DP-VOL, which is already used as an S-VOL at the time of the ShadowImage pair creation, a pair can be created by using the used DP-VOL. In that case, however, the initial copy time may be long. Therefore, create a pair after initializing the DP-VOL.

Table 2-1: Combination of a DP-VOL and a Normal LU

ShadowImage P-VOL	ShadowImage S-VOL	Comments
DP-VOL	DP-VOL	Available.
DP-VOL	Normal LU	Available. When executing the restore, the DP pool of the same capacity as the normal LU (S-VOL) is used.
Normal LU	DP-VOL	Available. In this combination, the DP pool of the same capacity as the normal LU (P-VOL) is used. Therefore, this combination is not recommended.

- LU type that can be set for a DMLU
The DP-VOL created by Dynamic Provisioning cannot be set for a DMLU. Set the normal LU for the DMLU.
- LU type that can be set for a command device
The DP-VOL created by Dynamic Provisioning cannot be set for a command device. Set the normal LU for a command device.
- Assigning the controlled processor core of a P-VOL or an S-VOL which uses the DP-VOL
When the controlled processor core of the DP-VOL used for a ShadowImage P-VOL or S-VOL differs from the normal LU, switch the S-VOL controlled processor core assignment to the P-VOL controlled processor core automatically, and create a pair. This applies to AMS2500.
- DP pool designation of a P-VOL or S-VOL which uses the DP-VOL
When using the DP-VOL created by Dynamic Provisioning for a ShadowImage P-VOL or S-VOL, using the DP-VOL designated in a separate DP pool of a P-VOL or S-VOL is recommended considering the performance implications.
- Pair status at the time of DP pool capacity depletion

When the DP pool is depleted after operating the ShadowImage pair which uses the DP-VOL created by Dynamic Provisioning, the pair status of the pair concerned may be a Failure. [Table 2-2](#) shows the pair statuses before and after the DP pool capacity depletion. When the pair status becomes a Failure caused by the DP pool capacity depletion, add the DP pool capacity whose capacity is depleted, and execute the pair operation again.

Table 2-2: Pair Statuses before and after DP Pool Capacity Depletion

Pair Statuses before the DP Pool Capacity Depletion	Pair Statuses after the DP Pool Capacity Depletion belonging to P-VOL	Pair Statuses after the DP Pool Capacity Depletion belonging to S-VOL
Simplex	Simplex	Simplex
Synchronizing	Synchronizing Failure*	Failure
Reverse Synchronizing	Failure	Reverse Synchronizing Failure*
Paired	Paired Failure*	Failure
Paired Internally Synchronizing	Paired Internally Synchronizing Failure*	Failure
Split	Split	Split
Split Pending	Split Pending Failure*	Failure
Failure	Failure	Failure

* When write is performed to the P-VOL to which the capacity depletion DP pool belongs, the copy cannot be continued and the pair status changes to Failure.

- DP pool status and availability of pair operation

When using the DP-VOL created by Dynamic Provisioning for a P-VOL or S-VOL of the ShadowImage pair, the pair operation may not be executed depending on the status of the DP pool to which the DP-VOL belongs.

[Table 2-3](#) shows the DP pool status and availability of the ShadowImage pair operation. When the pair operation fails due to the DP pool status, correct the DP pool status and execute the pair operation again.

Table 2-3: DP Pool Statuses and availability of ShadowImage Pair Operation

Pair Operation	DP Pool, DP Pool Capacity, and DP Pool Optimization Statuses					
	Normal	Capacity in Growth	Capacity Depletion	Regressed	Blocked	DP in Optimization
Create pair	+*	–	+	+	–	+ (Pairs cannot be created if the array firmware version is less than 0890/A.)
Create pair (split option)	+	–	+	+	–	+
Split pair	+	–	+	+	–	+
Resync pair	+*	–	+*	+	–	+
Restore pair	+**	–	+**	+	–	+
Delete pair	+	+	+	+	+	+

+ indicates a possible case, – indicates an impossible case

* Refer to the status of the DP pool to which the DP-VOL of the S-VOL belongs. If the status exceeds the DP pool capacity belonging to the S-VOL by the pair operation, the pair operation cannot be executed.

** Refer to the status of the DP pool to which the DP-VOL of the P-VOL belongs. If the status exceeds the DP pool capacity belonging to the P-VOL by the pair operation, the pair operation cannot be executed.



NOTE: When the DP pool was created or the capacity was increased, the DP pool underwent formatting. If pair creation, pair resynchronization, or restoration is performed during formatting, depletion of usable capacity may occur. Since the formatting progress is displayed when checking the DP pool status, check if sufficient usable capacity is secured according to the formatting progress, and then start the operation.

- Operation of the DP-VOL while using ShadowImage
When using the DP-VOL created by Dynamic Provisioning for a ShadowImage P-VOL or S-VOL, any of the operations among the capacity growing, capacity shrinking, and LU deletion of the DP-VOL in use cannot be executed. To execute the operation, delete the ShadowImage pair of which the DP-VOL to be operated is in use, and then execute it again. The attribute edit and capacity addition of the DP pool can be executed regardless of the ShadowImage pair.
- ShadowImage I/O Switching function
When the DP-VOLs created by Dynamic Provisioning are used for a ShadowImage P-VOL or S-VOL, ShadowImage I/O Switching function cannot operate.

- Volume write during Split Pending

When using the DP-VOL created by Dynamic Provisioning for a ShadowImage P-VOL or S-VOL, if writing to a P-VOL or an S-VOL when the pair status is Split Pending, the capacity of the DP pool to which both volumes belong may be consumed.


- Cascade connection

When the firmware array version is 0893/A or more, using the DP-VOL created by Dynamic Provisioning for a ShadowImage P-VOL or an S-VOL, for the pair including the normal LU, the cascade connection with TrueCopy can be executed.


However, when the firmware version of the array is less than 0893/A, see the following restrictions:

When using the DP-VOL for a ShadowImage P-VOL or S-VOL, for the pair including the normal LU, the cascade connection with TrueCopy, TCE, and SnapShot, cannot be executed. The cascade connection between ShadowImage also cannot be executed. [Figure 2-4](#) shows an example of the configuration that cannot execute the cascade connection. For example, in each cascade configuration, the cascade connection is executed with the TrueCopy volume for the normal LU. However, when the normal LU and DP-VOL executing the cascade connection are in the ShadowImage pair relation, the cascade connection also cannot be executed for the normal LU.

■ Cascade with a P-VOL (normal LU) of ShadowImage



■ Cascade with an S-VOL (normal LU) of ShadowImage



■ Cascade with a P-VOL (normal LU) of ShadowImage




Figure 2-4: Example of a non-supported Cascade Connection of ShadowImage with TrueCopy (P-VOL: S-VOL=1: 1)

Windows Server 2003 and 2008/Windows 2000 and Dynamic Disk

Observe the following when using Windows Server 2003, Windows Server 2008 or Windows 2000 and dynamic disk:

- A P-VOL and S-VOL can be made into a dynamic disk on Windows Server 2003; they cannot be made into a dynamic disk on Windows Server 2000 and Windows Server 2008.
- When using an S-VOL with a secondary host, ensure that the pair status is Split or Split Pending after the pair is created.
- A host cannot recognize both a P-VOL and its S-VOL at the same time. Map the P-VOL and S-VOL to separate hosts.
- An LU in which two or more dynamic disk volumes co-exist cannot be copied.
- Do not use a dynamic disk function for volumes other than a S-VOL on the secondary host side.

When copying, hide all the dynamic disks that exist on the primary side using the CCI `raidvchkset -vg idb` command. No restriction is placed on the primary side. Hide all the dynamic disk volumes to be restored on the primary side at the time of restoration.

If any one of the dynamic disks is left un-hidden, a **Missing** drive occurs. When this occurs, delete it manually using the `diskpart delete` command.




Figure 2-5: Supported, un-Supported Dynamic Disk Configurations

- Copy dynamic disk volumes that consist of two or more LUs only after hiding all LUs from a host. When the copy is completed, you can have them recognized by the host.




Figure 2-6: Supported Copying for Dynamic Disk Volumes

- A dynamic disk cannot be used with a cluster (MSCS, VCS, etc.) or VxVM and HDLM.

VMWare and ShadowImage configuration


When creating a backup of the virtual disk in the vmfs format using ShadowImage, shutdown the virtual machine that accesses the virtual disk, and then split the pair.

If one LU is shared by multiple virtual machines, it is required to shutdown all the virtual machines that share the LU when creating a backup. Therefore, it is not recommended to share one LU by multiple virtual machines in the configuration that creates a backup using ShadowImage.

Cascading ShadowImage with TrueCopy

The ShadowImage volumes can be cascaded (shared) with TrueCopy volumes, as shown in Figure 2-8. Several other cascade configurations are supported using TrueCopy and ShadowImage P-VOLs and S-VOLs on both the local and remote arrays.

- Cascade with a P-VOL of ShadowImage



- Cascade with an S-VOL of ShadowImage




Figure 2-7: Basic ShadowImage Cascade Examples with TCE

Read/write and pair operations are limited when a ShadowImage pair is cascaded with TrueCopy. Please see the chapter on cascading in the *Hitachi AMS 2000 Family TrueCopy Remote Replication User Guide* (MK-DF978052) for complete information.

Calculating maximum capacity

AMS manages capacity in segments of 15.75 GB for P-VOLs. Because of this it is necessary to calculate your ShadowImage system's managed capacity and compare it to the maximum supported capacity.

To calculate ShadowImage capacity

1. Add the sizes of all ShadowImage S-VOLs.

The value must be less than the maximum supported capacity for your array/memory capacity, as shown in [Table 2-4](#).

Table 2-4: Maximum Supported Capacity per Controller

Array	Mounted Memory Capacity (GB/CTL)	Maximum Supported Capacity (TB)
AMS 2100	1	261
	2	748
	4	897
AMS 2300	1	Not supported
	2	673
	4	1,870
	8	3,740
AMS 2500	2	486
	4	1,870
	6	2,618
	8	3,366
	10	4,488
	12	5,610
	16	6,733



NOTE: Maximum supported capacity is also limited by the size of the array.

If SnapShot, TrueCopy, or TCE are used on same array as ShadowImage

1. If SnapShot, TrueCopy, and/or TCE are used on the same array as ShadowImage, calculate the maximum supported capacity using the following formula:

Maximum ShadowImage Capacity

=> Max. Supported Capacity

- (17 x Total SnapShot P-VOL size)
 - (Total TrueCopy P-VOL and S-VOL size)
 - (51 x Total TCE P-VOL and S-VOL size)
- / 2

In the following example, the array/memory capacity is AMS2100/2 GB; therefore, maximum supported capacity is 748 TB.

- a. List ShadowImage Capacity: 5 TB
- b. List total SnapShot P-VOL size: 10 TB (if present)
- c. List total TrueCopy P-VOL and S-VOL size: 20 TB (if present)
- d. List total TCE P-VOL and S-VOL size: 15 TB (if present)
- e. Then:

5 TB => 748 TB -

(17 x 10 TB = 170 TB) = 578 TB -

20 TB = 558 TB -

(51 X 15 TB = 765 TB) = -187 TB

In this example, maximum supported capacity is exceeded.

= 748 TB - 170 TB = 578 TB

578 TB is available for the ShadowImage system. If:

ShadowImage S-VOL 1 = 15000 GB = 15 TB

ShadowImage S-VOL 2 = 2500 GB = 2.5 TB

Total ShadowImage S-VOL size = 17.5 TB

Then in this example, the total ShadowImage S-VOL size of 17.5 TB is well below the available capacity. of 578 TB.

If your calculated capacity value exceeds the maximum allowable capacity, you can do one or more of the following:

- Reduce the number of P-VOLs
- Change P-VOL sizes
- Reduce the lifespan of the S-VOL
- Reduce ShadowImage P-VOL size

Requirements

This chapter describes minimum operational requirements.

- ❑ [System requirements](#)
- ❑ [Supported platforms](#)

System requirements

Table 3-1 shows the minimum requirements for ShadowImage. See [Appendix A, ShadowImage specifications](#) for additional information.

Table 3-1: ShadowImage Requirements

Item	Minimum Requirements
Firmware	Version 0832/B or later is required for an AMS 2100 or 2300 array when the hardware Rev. is 0100. Version 0840/A or later is required for an AMS2500 array when the hardware Rev. is 0100. When executing the pair operation in Quick Mode, version 0861/A or more is required for an AMS array when the hardware Rev. is 0100. Version 0890/A or later is required for an AMS 2100, 2300, and 2500 when the hardware Rev. is 0200.
Storage Navigator 2	Version 3.21 or later is required for the management PC for an AMS 2100 or 2300 array when the hardware Rev. is 0100. Version 4.00 or later is required for the management PC for an AMS 2500 array when the hardware Rev. is 0100. When executing the pair operation in Quick Mode, version 6.10 or more is required for the management PC. Version 9.00 or later is required for the management PC for an AMS 2100, 2300, or 2500 when the hardware Rev. is 0200.
CCI	Version 01-21-03/06 or later is required for the host when CCI is used for the ShadowImage operation.
License key	Required for ShadowImage
Number of controllers	2 (dual configuration). Primary and secondary volumes must be on the same controller.
Command devices	Maximum of 128. The command device is required only when CCI is used for ShadowImage operations. CCI is provided for advanced users only. The command device volume size must be greater than or equal to 33 MB.
DMLUs	Maximum of 2. The Differential Management LU size must be greater than or equal to 10 GB. Two per array recommended in different RAID groups.
LU size	The S-VOL block count must be equal to the P-VOL block count.

Displaying the hardware revision number

The hardware revision (Rev.) number can be displayed when an individual array is selected from the Arrays list using Navigator 2, version 9.00 or later.

The screenshot shows the Hitachi Navigator 2 interface. On the left is a navigation pane with a tree view containing 'AMS2300_85000026', 'Components', 'Groups', 'Replication', 'Settings', 'Security', 'Performance', and 'Alerts & Events'. The 'AMS2300_85000026' item is selected. The main area displays the array's details. At the top, the array name 'AMS2300_85000026' is shown in blue, with a 'Refresh Information' button and a 'Help' button to its right. Below this is a 'Summary' table. The 'H/W Rev.' row is circled in red, showing the value '0100'. Other rows include 'Status' (Ready), 'Type' (AMS2300), 'Serial No.' (85000026), 'Array ID' (85000026), and 'Firmware' (0890/A-M). To the right of the 'H/W Rev.' row, there are columns for 'Capacity of All LU' (1.0TB), 'Raw Capacity of All Drives' (103.5TB), and IP addresses for both controllers. Below the summary table is a 'Common Array Tasks' section with several task cards: 'Initial Setup', 'Install License', 'Create Logical Unit and Mapping', 'Update Firmware', 'Backup Volume', 'Check for Errors', and 'Look at All Arrays'. Each card has an icon and a brief description of the task.

Summary			
Status	Ready	Capacity of All LU	1.0TB
Type	AMS2300	Raw Capacity of All Drives	103.5TB
H/W Rev.	0100	IPv4 Address	Controller 0 172.16.11.230
Serial No.	85000026	Controller 1	172.16.11.231
Array ID	85000026	IPv6 Address	Controller 0 2001::200:87ff:fec6:46e7
Firmware	0890/A-M	Controller 1	2001::200:87ff:fec6:46e9

Common Array Tasks

The following menu helps you perform typical tasks. For additional settings, please use the array tree menu.

- Initial Setup**
Configure several items on the array to make it ready to use.
- Install License**
Install Licenses for optional storage features.
- Create Logical Unit and Mapping**
Create logical unit and mapping easily.
- Update Firmware**
Update the firmware in the array from a local file or from the support website.
- Backup Volume**
Copy the selected volume to prevent data loss.
- Check for Errors**
View the Alerts & Events screen and show the latest status of the array.
- Look at All Arrays**
Log out of this array and return to the main arrays list. You can then choose another array to manage.

Supported platforms

Table 3-2 shows the supported platforms and operating system versions required for ShadowImage.

Table 3-2: Supported Platforms

Platforms	Operating System Version
SUN	Solaris 8 (SPARC)
	Solaris 9 (SPARC)
	Solaris 10 (SPARC)
	Solaris 10 (x86)
	Solaris 10 (x64)
PC Server (Microsoft)	Windows 2000
	Windows Server 2003 (IA32)
	Windows Server 2008 (IA32)
	Windows Server 2003 (x64)
	Windows Server 2008 (x64)
	Windows Server 2003 (IA64)
	Windows Server 2008 (IA64)
Red Hat	Red Hat Linux AS2.1 (IA32)
	Red Hat Linux AS/ES 3.0 (IA32)
	Red Hat Linux AS/ES 4.0 (IA32)
	Red Hat Linux AS/ES 3.0 (AMD64/EM64T)
	Red Hat Linux AS/ES 4.0 (AMD64/EM64T)
	Red Hat Linux AS/ES 3.0 (IA64)
	Red Hat Linux AS/ES 4.0 (IA64)
HP	HP-UX 11i V1.0 (PA-RISC)
	HP-UX 11i V2.0 (PA-RISC)
	HP-UX 11i V3.0 (PA-RISC)
	HP-UX 11i V2.0 (IPF)
	HP-UX 11i V3.0 (IPF)
	Tru64 UNIX 5.1
IBM®	AIX 5.1
	AIX 5.2
	AIX 5.3
SGI	IRIX 6.5.x

Install and enable

This chapter provides instructions for installing and enabling ShadowImage.

- ❑ [Installing ShadowImage](#)
- ❑ [Enabling/disabling ShadowImage](#)
- ❑ [Uninstalling ShadowImage](#)

Installing ShadowImage

If ShadowImage was purchased at the same time as the order for AMS was placed, then ShadowImage is bundled with the array and no installation is necessary. Proceed to [Enabling/disabling ShadowImage on page 4-3](#).

If ShadowImage was purchased on an order separate from Adaptable, it must be installed before enabling.



NOTE: A key code or key file is required to install or uninstall. If you do not have the key file or code, you can obtain it from the download page on the HDS Support Portal, <http://support.hds.com>.

- Installation instructions are provided here for Navigator 2 GUI.
- For CLI instructions, see [Installing and uninstalling ShadowImage on page B-2](#) (advanced users only).

Before installing or uninstalling ShadowImage, verify that the array is operating in a normal state. Installation/un-installation cannot be performed if a failure has occurred.

To install ShadowImage

1. In the Navigator 2 GUI, click the check box for the array where you want to install ShadowImage.
2. Click the **Show & Configure Array** button. The tree view appears as shown in [Figure 4-1](#).




Figure 4-1: Tree View

3. Click **Settings**, then click **Licenses**.
4. On the Licenses page, click the ShadowImage check box, then click the **Install License** button at the bottom of the page.

The Install License screen appears, as shown in [Figure 4-2](#).

Install License

Help

License Property

Enter the information for the license to be installed.

*Install with:

☐ Key File:

Input the Key File Name.

☒ Key Code:

Input the Key Code.

* Required field

Figure 4-2: Install License Screen

5. Select the **Key File** or **Key Code** radio button, then enter the file name or key code. You may browse for the Key File.
6. Click **OK**.
7. Click **Close** on the confirmation screen.

Enabling/disabling ShadowImage

Enable or disable ShadowImage using the following procedure.



NOTE: Before disabling, ShadowImage pairs must be deleted.

To enable or disable ShadowImage

1. In the Navigator 2 GUI, click the check box for the array where you want to enable or disable ShadowImage
2. Click the **Show & Configure Array** button.
3. Click **Settings**, then click **Licenses**.
4. Select **SHADOWIMAGE** in the Licenses list.
5. Click **Change Status**. The Change License screen appears, as shown in [Figure 4-3](#).

Change License - SHADOWIMAGE Help




Figure 4-3: Change License Screen

6. To enable, click the **Enable: Yes** check box.
To disable, clear the **Enable: Yes** check box.
7. Click **OK**.
8. A message appears confirming that ShadowImage is enabled or disabled. Click **Close**.

Uninstalling ShadowImage

ShadowImage pairs must be released and their status returned to Simplex before uninstalling.

To uninstall ShadowImage

1. In the Navigator 2 GUI, click the check box for the array where you want to uninstall ShadowImage
2. Click the **Show & Configure Array** button.
3. In the tree view, click **Settings**, then click **Licenses**.
4. On the Licenses screen, select **SHADOWIMAGE** in the Licenses list and click the **De-install License** button.
5. On the De-Install License screen, enter the code in the **Key Code** box, and then click **OK**.
6. On the confirmation screen, click **Close** to confirm.

Configuration

This chapter provides required information for setting up your system for ShadowImage.

- ❑ [Configuration workflow](#)
- ❑ [Setting Up primary, secondary volumes](#)
- ❑ [Setting the DMLU](#)

Configuration workflow

Setup for ShadowImage consists of making certain that primary and secondary volumes are set up correctly.

- Setting up the DMLU.

Setting Up primary, secondary volumes

The primary and secondary volumes must be set up prior to making ShadowImage copies. When doing so, adhere to the following:

- The P-VOL and S-VOL must have identical block counts. Verify block size in the Navigator 2 GUI by navigating to Groups/Groups/Logical Units tab. Click the LUN whose block size you want to check. On the popup window that appears, review the Capacity field. This shows block size.
- The P-VOL and S-VOL must be assigned to the same controller.

Please refer to the following section for all key requirements and recommendations:

- [Appendix A, ShadowImage specifications](#)

Setting the DMLU

The DMLU is an exclusive volume used for storing ShadowImage information when the array is powered down. You must set up the DMLU before using ShadowImage.

Prerequisites

- LUs for the DMLU must be set up and formatted. Minimum size: 10 GB.
- Though only one DMLU is needed, two are recommended, with the second used as a backup.
- For RAID considerations, see the bullet on DMLU in [RAID configuration for ShadowImage LUs on page 2-6](#).
- Also see DMLU items in [Appendix A, ShadowImage specifications](#).

To set up a DMLU

1. In the navigation tree, select **Settings**, then **DMLU**. The Differential Management Logical Units screen displays.
2. Click **Add DMLU**. The Add DMLU screen displays.

Add DMLU

Help

Differential Management Logical Unit Property

Select one or two logical units as the differential management logical unit(DMLU).

* Logical units to add as DMLU :

Assignable Logical Units

Rows/Page: 25 | Page 1 of 1

☐	LUN	Capacity	RAID Group	DP Pool	RAID Level
☐	0000	50.0GB	000	N/A	RAID6(8D)
☐	0001	50.0GB	000	N/A	RAID6(8D)
☐	1002	10.0GB	000	N/A	RAID6(8D)
☒	1003	10.0GB	000	N/A	RAID6(8D)

Filter Filter Off

* Required field



3. Select the **LUNs** that you want to assign as DMLUs, and then click **OK**. A confirmation message displays.
4. Select the **Yes, I have read ...** check box, then click **Confirm**. When the success message displays, click **Close**.

Add DMLU - 1003

 **Confirm the volume that you have selected.**
Be aware that the data in the selected volume will be overwritten.

YOU CANNOT UNDO THIS OPERATION.

☒ Yes, I have read the above warning and want to set these logical units as DMLUs.

Confirm Cancel

Confirm

Setting the ShadowImage I/O switching mode

The following procedure is for setting the ShadowImage I/O switching mode to ON.

To display the applet screen of Navigator 2, it is necessary to set the **Java Runtime Parameters** by a computer that starts Navigator 2.


When you use the JRE less than 1.6.0_10, setting the **Java Runtime Parameters** are necessary in a client to start Navigator 2. When you use the JRE 1.6.0_10 or more, setting the **Java Runtime Parameters** are not necessary in a client to start Navigator 2. However, starting the **Open Advanced Settings**, when "DMEG0002F0: Since memories required for the Advanced Settings are insufficient, a screen cannot be displayed. Change a setup of Java Plug-in installed in the client and increase the usable memories." appears, set the following **Java Runtime Parameters**.

For Windows:

1. In the Windows **Start** menu, choose **Settings, Control Panel**.
2. From the **Control Panel**, select **Java**.
3. Click **View** of the upper position in the **Java** tab.
4. Enter "-Xmx216m" to the Java Runtime Parameters field.
5. Click **OK**.
6. Click **OK** in the Java tab.
7. Close the **Control Panel**.

To set the ShadowImage I/O Switching Mode:

1. Start Navigator 2.
2. Log in as registered user to Navigator 2.
3. Select the array in which you will set ShadowImage.
4. Click **Show & Configure Array**.
5. Select the **Advanced Settings** icon in the **Settings** tree view.
6. Click **Open Advanced Settings**.




After a few minutes, an Array Unit (Applet) screen displays.




NOTE: The Applet screen is displayed connected to the **SNM2 Server**. If 20 minutes elapse while displaying the Applet screen, you will not be able to operate it due to the automatic logoff function. If the operation is completed, close the screen. If the Applet screen cannot be displayed, the login to the **SNM2 Server** may have failed. In this case, the Applet screen cannot be displayed again. The code: 0x000000000000b045 or "DMEG800003: The error occurred in connecting RMI server." is displayed on the Applet screen.

Perform the following:

- a. Close the Web browser, stop the **SNM2 Server** once, restart it, and display the screen of the array you want to operate.
 - b. Close the Web browser; confirm the **SNM2 Server** is started. If it has stopped, start it and display the screen of the array that you want to operate.
 - c. Return to the array screen after 20 minutes elapse and display the screen of the array you want to operate.
7. Select the **Configuration Settings** icon in the Array Unit tree view.





8. Click **Set**.
- The **Configuration Settings** tabbed page appears.
9. Click the **System Parameter** tab.



10. Select the **ShadowImage I/O Switch Mode** in the **Options** and click **Apply**.

11. A message displays. Click **OK**.



12. Click **Close** in the **Configuration Settings** tabbed page.

13. Click **Close** in an Array Unit screen.



NOTE: When turning off the ShadowImage I/O Switching mode, it is required to make pair statuses of all ShadowImage pairs to those other than Split (S-VOL Switch) and Synchronizing (S-VOL Switch).

Using ShadowImage

This chapter describes ShadowImage operations.

- ❑ [ShadowImage workflow](#)
- ❑ [Check pair status](#)
- ❑ [Create the pair](#)
- ❑ [Split the pair](#)
- ❑ [Resync the pair](#)
- ❑ [Delete a pair](#)
- ❑ [Edit a pair](#)
- ❑ [Restore the P-VOL](#)
- ❑ [Use the S-VOL for tape backup, testing, reports, etc.](#)

ShadowImage workflow

A typical ShadowImage workflow consists of the following:

- Check pair status. Each operation requires a pair to have a specific status.
- Create the pair, in which the S-VOL becomes a duplicate of the P-VOL.
- Split the pair, which separates the primary and secondary volumes and allows use of the data in the S-VOL by secondary applications.
- Re-synchronize the pair, in which the S-VOL again mirrors the on-going, current data in the P-VOL.
- Restore the P-VOL from the S-VOL.
- Delete a pair.
- Edit pair information.

(For an illustration of basic ShadowImage operations, see [Figure 1-2 on page 1-5](#).)

Check pair status

When you want to perform a specific ShadowImage operation the pair must be in a state that allows the operation. For instructions on checking pair status plus status definitions, see [Monitor pair status on page 7-2](#).

Create the pair

Please review the following before creating a pair.

- When the primary volume *is not* part of another ShadowImage pair, both primary and secondary volumes must be in the SMPL (simplex) state.
- When the primary volume *is* part of another ShadowImage pair or pairs, the pair status of those pairs must be Split.
- The primary and secondary volumes must have identical block counts and must be assigned to the same controller.
- Because pair creation affects the performance on the host, observe the following:
 - Create a pair when I/O load is light.
 - Limit the number of pairs that you create simultaneously.
- During the Create Pair operation, the following takes place:
 - All data in the P-VOL is copied to the S-VOL.
 - The P-VOL remains available to the host for read/write throughout the copy operation.
 - Writes to the P-VOL during pair creation are copied to the S-VOL.
 - Pair status is Synchronizing while the initial copy operation is in progress.

- Status changes to Paired when the initial copy is complete.
- New writes to the P-VOL continue to be copied to the S-VOL in the Paired status.

Copy pace

Copy pace is the speed at which a pair is created or re-synchronized. You select the copy pace when you create or resynchronize a pair (if using CLI, you enter a copy pace parameter).

Copy pace impacts host I/O performance. A slow copy pace has less impact than a medium or fast pace. The pace is divided on a scale of 1 to 15, as follows:

- Slow — between 1-5. The process takes longer when host I/O activity is heavy. The amount of time to complete an initial copy or resync cannot be guaranteed.
- Medium — between 6-10. (Recommended) The process is performed continuously, but the amount of time to complete the initial copy or resync cannot be guaranteed. Actual pace varies according on host I/O activity.
- Fast — between 11-15. The copy/resync process is performed continuously and takes priority. Host I/O performance is restricted. The amount of time to complete an initial copy or resync is guaranteed.

Create a ShadowImage pair

To use CLI, see [Creating ShadowImage pairs on page B-6.](#))

1. In Navigator 2 GUI, select the desired array, then click the **Show & Configure Array** button.
2. From the Replication tree, select the **Local Replication** icon. The **Pairs** screen displays.
3. Click the **Create Pair** button at the bottom of the screen. View further instructions by clicking the Help button.

Create Pair Help

Local Pair Property

Select a copy type of local replication and enter property information as desired.

* Copy Type : ☒ ShadowImage ☐ SnapShot

Basic **Advanced**

Select the primary volume and the secondary volume to pair and give the pair a unique name.

Pair Name :

31 characters or less (alphanumeric characters, special symbols "%", "*", "+", "-", ".", "/", "=", "@", "_", ":", "[", or "]")

* Select a Primary Volume :

Primary Volume					
Rows/Page: 25 Page 1 of 1					
	LUN	Capacity	RAID Group	DP Pool	RAID Level
☐	0004	20.0GB	000	N/A	RAID5(4D+1P)
☒	0006	10.0GB	000	N/A	RAID5(4D+1P)
☐	0007	10.0GB	000	N/A	RAID5(4D+1P)
☐	0008	20.0GB	000	N/A	RAID5(4D+1P)

* Select a Secondary Volume :

Secondary Volume					
Rows/Page: 25 Page 1 of 1					
	LUN	Capacity	RAID Group	DP Pool	RAID Level
☐	0004	20.0GB	000	N/A	RAID5(4D+1P)
☐	0006	10.0GB	000	N/A	RAID5(4D+1P)
☒	0007	10.0GB	000	N/A	RAID5(4D+1P)
☐	0008	20.0GB	000	N/A	RAID5(4D+1P)

* Required field

OK Cancel

4. After making selections on the **Basic** tab, further customize the pair by clicking the **Advanced** tab.

Create Pair

Help

Local Pair Property

Select a copy type of local replication and enter property information as desired.

* Copy Type : ☒ ShadowImage ☐ SnapShot

Basic **Advanced**

Enter any additional properties for the pair.

Copy Pace : Medium

After pair creation, add the pair to a group.- Group Assignment:

☒ {Ungrouped}

☐ New or existing Group Number :

☐ Existing Group Name :

Do initial copy from the primary volume to the secondary volume : ☒ Yes

Allow read access to the secondary volume after the pair is created : ☒ Yes

Automatically split the pair immediately after they are created: ☐ Yes

* Required field

OK Cancel
OK

5. From the **Copy Pace** dropdown list, select the speed at which copies will be made. Select Slow, Medium, or Fast. (See [Copy pace](#), above, for more information.)
6. In the **Group Assignment** area, you have the option of assigning the new pair to a consistency group. (For a description, see [Consistency group \(CTG\) on page 1-11.](#)) Do one of the following:
 - If you do not want to assign the pair to a consistency group, leave the **Ungrouped** button selected.
 - To create a group and assign the new pair to it, click the **New or existing Group Number** button and enter a new number for the group in the box.
 - To assign the pair to an existing group, enter its number in the **Group Number** box, or enter the group name in the **Existing Group Name** box.



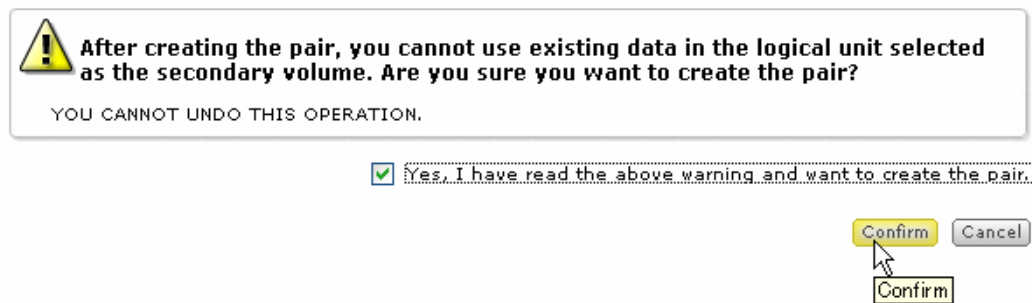
NOTE: You can also add a Group Name for a consistency group as follows:

- a. After completing the create pair procedure, on the **Pairs** screen, check the box for the pair belonging to the group.
 - b. Click the **Edit Pair** button.
 - c. On the Edit Pair screen, enter the **Group Name** then click **OK**.
7. In the **Do initial copy from the primary volume...** field, leave **Yes** checked to copy the primary to the secondary volume.

Clear the check box to create a pair without copying the P-VOL at this time, and thus reduce the time it takes to create the pair. The system treats the two volumes as a pair.

8. In the **Allow read access to the secondary volume after the pair is created** field, leave **Yes** checked to allow access to the secondary volume after the pair is created. Clear the check box to prevent read/write access to the S-VOL from a host after the pair is created. This option (un-checking) insures that the S-VOL is protected and can be used as a backup.
9. Add a check mark to the box **Automatically split the pair immediately after they are created** when you want to automatically split the pair after creation.
10. Click **OK**.
11. A confirmation message displays. Check the **Yes, I have read the above warning and want to create the pair** check box, and click **Confirm**.

Create Pair



12. A confirmation message displays. Click **Close**.

Create Pair



Split the pair

When a primary and secondary volume are in Pair status, all data that is written to the primary volume is copied to the secondary volume. This continues until the pair is split.

When it is split, updates continue to be written to the primary volume, but not to the secondary volume. Data in the S-VOL is frozen at the time of the split. After the Split Pair operation:

- The secondary volume becomes available for read/write access by secondary host applications.
- Separate track tables record updates to the P-VOL and to the S-VOL.
- The pair can be made identical again by re-synchronizing from primary-to-secondary or secondary-to-primary.

To split the pair

To use CLI, see [Splitting ShadowImage pairs on page B-7.](#))

1. In Navigator 2 GUI, select the desired array, then click the **Show & Configure Array** button.
2. From the Replication tree, select the **Local Replication** icon. The **Pairs** screen displays.
3. Select the pair you want to split.
4. Click the **Split Pair** button at the bottom of the screen. View further instructions by clicking the Help button, as needed.

SplitPair - SI_LU1000_LU1001 Help

Local Pair Split Option

Select the option of the split operation.

Option :

Suspend operation in progress and force the pair into a failure state: ☐ Yes

Attach description to identify the pair upon split :

31 characters or less (alphanumeric characters, special symbols "%", "*", "+", "-", ":", "/", "=", "@", "_", ";", "[", or "]")

Quick Mode : ☐ Yes

OK Cancel

5. Mark the check box of the **Suspend operation in progress and force the pair in failure state** if necessary.
6. Enter a character strings to the **Attach description to identify the pair upon split** if necessary.
7. When you want to split Quick Mode, add a check mark to **Quick Mode**.
8. Click **OK**.
9. A confirmation message displays. Click **Close**.

Resync the pair

Re-synchronizing a pair that has been split updates the S-VOL so that it is again identical with the P-VOL. A reverse resync updates the P-VOL so that it is identical with the S-VOL.

- The pair must be in Split status.
- Pair status during a normal re-synchronizing is Synchronizing.
- Status changes to Paired when the resync is complete.
- When the pair is re-synchronized, it can then be split for tape backup or other uses of the updated S-VOL.



NOTE: Because updating the S-VOL affects performance in the RAID group to which the pair belongs, best results are realized by performing the operation when I/O load is light. Priority should be given to the Resync process.

To resync the pair

To use CLI, see [Re-synchronizing ShadowImage pairs on page B-8.](#))

1. In Navigator 2 GUI, select the desired array, then click the **Show & Configure Array** button.
2. From the Replication tree, select the **Local Replication** icon. The **Pairs** screen displays.
3. Select the pair you want to resync.
4. Click the **Resync Pair** button. The Resync Pair screen appears as shown below. View further instructions by clicking the Help button, as needed.

Resync Pair - SI_LU1000_LU1001

Help

Option of Re-synchronize Local Pair

Select the option of the re-sync operation.

Quick Mode : ☐ Yes

OK Cancel
OK

5. When you want to re-synchronize Quick Mode, place a check mark in the Yes box for **Quick Mode**.
6. Click OK.

A confirmation message displays.

Resync Pair - SI_LU1000_LU1001



After re-synchronizing the pair, you cannot use existing data in the secondary volume. Are you sure you want to re-synchronize the pair?

YOU CANNOT UNDO THIS OPERATION.

☒ Yes, I have read the above warning and want to re-synchronize selected pairs.

Confirm Cancel
Confirm

7. For **Yes, I have read the above warning and want to re-synchronize selected pairs**. Place a check in the box, and click **Confirm**.
8. A confirmation message displays. Click **Close**.

Delete a pair

You can delete a pair when you no longer need it. When you delete a pair, the primary and secondary volumes return to their SIMPLEX state. Both are available for use in another pair.

To delete a ShadowImage pair

To use CLI, see [Releasing ShadowImage pairs on page B-9.](#))

1. In Navigator 2 GUI, select the desired array, then click the **Show & Configure Array** button.
2. From the Replication tree, select the **Local Replication** icon. The **Pairs** screen displays.
3. Select the pair you want to delete.
4. Click **Delete Pair**.

Edit a pair

You can edit the name, group name, and copy pace for a pair.

To edit pairs

To use CLI, refer to *Storage Navigator Modular 2 Command Line Interface (CLI) Reference Guide (for Replication)*.

1. In Navigator 2 GUI, select the desired array, then click the **Show & Configure Array** button.
2. From the Replication tree, select the **Local Replication** icon. The **Pairs** screen displays.
3. Select the pair that you want to edit.
4. Click the **Edit Pair** button. View instructions by clicking the Help button.

Restore the P-VOL

ShadowImage enables you to restore your P-VOL to a previous point in time. You can restore from any S-VOL paired with the P-VOL.

The amount of time it takes to restore your data is dependent on the size of the P-VOL and the amount of data that has changed.

To restore the P-VOL from the S-VOL

1. Shut down the host application.
2. Un-mount the P-VOL from the production server.
3. In the Storage Navigator 2 GUI, select the **Local Replication** icon in the Replication tree view.

Advanced users using the Navigator 2 CLI, please refer to [Restoring the P-VOL on page B-8.](#)

4. In the GUI, select the pair to be restored in the **Pairs** list.

5. Click **Restore Pair**. View further instructions by clicking the Help button as needed.
6. Mount the P-VOL.
7. Re-start the application.

Use the S-VOL for tape backup, testing, reports, etc.

Your ShadowImage copies can be used on a secondary server to fulfill a number of data management tasks. These might include backing up production data to tape, using the data to develop or test an application, generating reports, populating a data warehouse, and so on.

Whatever the task, the process for preparing and making your data available is the same. The following process can be performed using the Navigator 2 GUI or CLI, in combination with an operating system scheduler. The process should be performed during non-peak hours for the host application.

To use the S-VOL for secondary functions

1. Un-mount the S-VOL if it is being used by a host.
2. Resync the pair before stopping or quiescing the host application. This is done to minimize the down time of the production application.
 - Navigator 2 GUI users, please see [Resync the pair on page 6-7](#).
 - Advanced users using CLI, please see [Re-synchronizing ShadowImage pairs on page B-8](#).



NOTE: Some applications can continue to run during a backup operation, while others must be shut down. For those that continue running (placed in backup mode or quiesced rather than shut down), there may be a host performance slowdown.

3. When pair status becomes Paired, shut down or quiece (quiet) the production application, if possible.
4. Split the pair. Doing this insures that the backup will contain the latest mirror image of the P-VOL.
 - GUI users please see [Split the pair on page 6-6](#).
 - Advanced users using CLI, please see [Splitting ShadowImage pairs on page B-7](#).
5. Un-quiesce or start up the production application so that it is back in normal operation mode.
6. Mount the S-VOL on the server, if needed.
7. Run the backup program using the S-VOL.

Monitoring and troubleshooting

This chapter provides information and instructions for monitoring and troubleshooting the ShadowImage system.

- ❑ [Monitor pair status](#)
- ❑ [Troubleshooting](#)

Monitor pair status

Monitoring pair status insures the following:

- A pair is in the correct status for the ShadowImage operation you wish to perform.
- Pairs are operating correctly and status is changing to the appropriate state during and after an operation.
- Data is being updated from P-VOL to S-VOL in a pair resync, and from S-VOL to P-VOL in a pair reverse resync.
- Differential data management is being performed in the Split status. The Status column on the **Pairs** screen shows the percentage of synchronization. This can be used to estimate the amount of time a resync will take.

To check pair status

To use CLI, see [Confirming pairs status on page B-5.](#))

1. In Navigator 2 GUI, select the desired array, then click the **Show & Configure Array** button.
2. From the Replication tree, select the **Local Replication** icon. The **Pairs** screen displays.
3. Locate the pair and review the **Status** field.



NOTE: The pair status and pair identical rate are displayed.

[Table 7-1](#) shows Navigator2 GUI status and descriptions. For CCI statuses see [Confirming pair status on page C-8.](#)

Table 7-1: Pair Statuses

GUI Pair Status	Description
Simplex	Status when the volume has not have been paired or when the pair has been deleted. The volume does not have a ShadowImage association with another volume.
Paired	S-VOL is a duplicate of the P-VOL. Updates to the P-VOL are copied to the S-VOL.
Paired Internally Synchronizing	Re-synchronizing pairs specifying the Quick Mode.
Synchronizing	Initial or re-synchronization copy is in progress. The array continues to accept read and write operations for the P-VOL but does not accept write operations for the S-VOL. When a split pair is resynchronized in normal mode, the array copies only the P-VOL differential data to the S-VOL. When creating a pair or when a Failure pair is resynchronized, the array copies the entire P-VOL to the S-VOL.
Split	Updates stop from P-VOL to S-VOL. The S-VOL remains a copy of the P-VOL at the time of the split. P-VOL continues being updated by the host application.
Split Pending	A pair split specifying the Quick Mode.
Resynchronizing	The S-VOL is updated from the P-VOL. When this operation is completed, the status changes to Paired.

Table 7-1: Pair Statuses

GUI Pair Status	Description
Reverse Synchronizing	P-VOL restoration from S-VOL is in progress.
Failure	Copying is suspended due to a failure occurrence. The array marks the entire P-VOL as differential data; thus, it must be copied in its entirety to the S-VOL when a Resync is performed.



NOTE: The identical rate displayed with the pair status shows the identical ratio of the P-VOL and S-VOL data that can be accessed from the host. When the pair status is Split Pending, even though the background copy is performed, if the P-VOL and S-VOL data viewed from the host is matched, the identical rate becomes 100%. The ratio that the background copy is completed is indicated by the Progress. You can check the Progress by the detail information of each pair.

Troubleshooting

Pair failure

A pair failure occurs when one of the following takes place:

- A hardware failure occurs.
- Forcible release is performed by the user. This occurs when you halt a Pair Split operation. The array places the pair in Failure status.

If the pair was not forcibly suspended, the cause is hardware failure.

To restore pairs after a hardware failure

1. If the volumes were re-created after the failure, the pairs must be re-created.
2. If the volumes were recovered and it is possible to resync the pair, then do so. If resync is not possible, delete then re-create the pairs.
3. If a P-VOL restore was in progress during a hardware failure, delete the pair, restore the P-VOL if possible, and create a new pair.

To restore pairs after forcible release operation

Create or re-synchronize the pair. When an existing pair is re-synchronized, the entire P-VOL is re-copied to the S-VOL.

Path failure

When using CCI, and a path fails for more than one minute, the command device may not be recognized when the path is recovered. Execute Windows' "re-scan the disks" to recovery. Restart CCI if Windows cannot access the command device even, if CCI is able to recognize Windows.

ShadowImage specifications

This appendix provides specifications for ShadowImage.

- [General specifications](#)

General specifications

Table A-1 lists and describes the external specifications for ShadowImage.

Table A-1: Specifications

Item	Specification
Adaptable Modular Storage model	AMS2500 AMS2300 AMS2100 (For dual configuration only.)
Host interface	Fibre channel or iSCSI.
Number of pairs	- AMS2500, AMS2300: maximum of 2047 pairs - AMS2100: maximum of 1023 pairs Note: When a P-VOL is paired with eight S-VOLs, the number of pairs is eight.
Command devices	Required for CCI. - Maximum: 128 per array - Volume size: 33MB or greater
Unit of pair management	Volumes are the target of ShadowImage pairs, and are managed per logical unit.
Pair structure (number of S-VOLs per P-VOL)	1 P-VOL:8 S-VOLs
Differential Management LU (DMLU)	- Minimum size: 10 GB - One DMLU is required, two (maximum) are recommended. When using two, each must be: - In different RAID groups than the other - Under different controller than the other - The DP-VOL created by Dynamic Provisioning cannot be set for DMLU.
RAID level	- P-VOL: RAID 0 (2D to 16D), RAID 1+0 (2D+2D to 8D+8D), RAID 5 (2D+1P to 15D+1P), RAID 6 (2D+2P to 28D+2P), RAID 1 (1D+1D) ("with redundancy" recommended) - S-VOL: RAID 0 (2D to 16D), RAID 1+0 (2D+2D to 8D+8D), RAID 5 (2D+1P to 15D+1P), RAID 6 (2D+2P to 28D+2P), RAID 1 (1D+1D) ("with redundancy" recommended) - RAID 0 cannot be set for the SATA disk drive.
Combination of RAID groups	P-VOL and S-VOL should be paired on different RAID groups. The number of data disks does not have to be the same.
Size of P-VOL and S-VOL	P-VOL = S-VOL (block count must be equal)
Types of drive for the P-VOL and S-VOL	If the drive types are supported by the array, they can be set for the P-VOL and S-VOL. Both SAS/SSD and SATA-drive LUs can be assigned to a P-VOL or S-VOL. However, SAS or SSD drives are recommended for the P-VOL.

Table A-1: Specifications

Item	Specification
Consistency Group (CTG) number	CTG per array: Max 256 Pairs per group: - AMS2500/2300: Max 2047 - AMS2100: Max 1023
Concurrent use with TrueCopy	Yes. The cascade connection cannot be executed with the ShadowImage pair including the DP-VOL created by Dynamic Provisioning. See Cascading ShadowImage with TrueCopy on page 2-18 for more information.
Concurrent use of TCE	ShadowImage and TCE can be used on the same array; however, a cascaded system between ShadowImage and TCE is not supported.
Concurrent use of Dynamic Provisioning	The DP-VOL created by Dynamic Provisioning can be used as a ShadowImage P-VOL or S-VOL. For more details, see Concurrent use with Dynamic Provisioning on page 2-10 .
Concurrent use of SnapShot	SnapShot and ShadowImage can be used together at the same time, but cascading between SnapShot and ShadowImage is not supported. When using SnapShot and ShadowImage together, the maximum number of CTGs is 256.
Formatting, growing/shrinking, deleting during Coupling (RAID group, P-VOL, S-VOL).	Not available. However, when pair status is failure (S-VOL Switch), a P-VOL can be formatted. When pair status is Simplex, you can grow or shrink LUs.
Concurrent use of Volume Migration	Yes, however a P-VOL, an S-VOL, and a reserved LU of Volume Migration cannot be specified as a ShadowImage P-VOL. The maximum number of the pairs and the number of pairs whose data can be copied in the background is limited when ShadowImage is used together with Volume Migration.
Concurrent use of Cache Residency Manager	Yes, however the LU specified for Cache Residency (LU cache residence) cannot be used as a P-VOL, S-VOL.
Concurrent use of Cache Partition Manager	Yes
Concurrent use of SNMP Agent	Yes. Traps are sent following when occurs. Pair status changes to Failure.
Concurrent use of Data Retention Utility	Yes. However, when S-VOL Disable is set for an LU, the LU cannot be used in a ShadowImage pair. When S-VOL Disable is set for an LU that is already a S-VOL, no suppression of the pair takes place, unless the pair status is split.

Table A-1: Specifications

Item	Specification
Concurrent use of Power Saving	Yes. However, when a P-VOL or S-VOL is included in a RAID group in which Power Saving is enabled, the only ShadowImage pair operation that can be performed are pair split and the pair release.
Concurrent use of unified LU	Yes.
Concurrent use of LUN Manager	Yes.
Concurrent use of Password Protection	Yes.
ShadowImage I/O switching function	Yes. DP-VOLs created by Dynamic Provisioning can be used for a P-VOL or an S-VOL of ShadowImage. For details, see Appendix D, I/O switching mode feature .
Maximum supported capacity value of S-VOL (TB)	See Calculating maximum capacity on page 2-19 for details.
License	ShadowImage must be installed using the key code.
Management of LUs while using ShadowImage	Formatting and deleting LUs are not available. When formatting and deleting LUs, split ShadowImage pair(s) using the pairsplit command.
Restriction for formatting the volumes	Do not execute ShadowImage operations while formatting the volume. Formatting takes priority and the ShadowImage operations will be suspended.
Restriction during RAID group expansion	A RAID group with a ShadowImage P-VOL or S-VOL can be expanded only when the pair status is Simplex or Split.
Failures	When a failure of the copy operation from P-VOL to S-VOL occurs, ShadowImage will suspend the pair and the status changes to failure. If a volume failure occurs, ShadowImage suspends the pair. If a drive failure occurs, the ShadowImage pair status is not affected because of the RAID architecture.
Reduction of memory	The memory cannot be reduced when ShadowImage, SnapShot, or TrueCopy are enabled. Reduce memory after disabling the functions.

Operations using CLI

This chapter describes basic Navigator 2 CLI procedures for performing ShadowImage operations.

- ❑ [Installing and uninstalling ShadowImage](#)
- ❑ [ShadowImage operations](#)
- ❑ [Creating ShadowImage pairs that belong to a group](#)
- ❑ [Sample back up script for Windows](#)



NOTE: For additional information on the commands and options used in this appendix, see the *Navigator 2 Command Line Interface (CLI) Reference Guide for Replication*.

Installing and uninstalling ShadowImage

If ShadowImage was purchased when the order for AMS was placed, then ShadowImage came bundled with the system and no installation is necessary. Proceed to [Enabling or disabling ShadowImage on page B-3](#).

If you purchased ShadowImage on an order separate from your AMS, it must be installed before enabling. A key code or key file is required.



NOTE: Before installing/uninstalling ShadowImage, verify that the array is operating in a normal state. If a failure such as a controller blockade has occurred, installation/un-installation cannot be performed.

Installing ShadowImage

To install ShadowImage, the key code or key file provided with the optional feature is required. You can obtain it from the download page on the HDS Support Portal, <http://support.hds.com>.

To install ShadowImage

1. From the command prompt, register the array in which the ShadowImage is to be installed, then connect to the array.
2. Execute the `auopt` command to install ShadowImage. For example:

```
% auopt -unit subsystem-name -lock off -licensefile license-
      file-path\license-file-name
No.   Option Name
    1  ShadowImage in-system replication
Please specify the number of the option to unlock.
When you unlock two or more options, partition the numbers
given in the list with space(s). When you unlock all options,
input 'all'. Input 'q', then break.
The number of the option to unlock. (number/all/q [all]): 1
Are you sure you want to unlock the option?
(y/n [n]): y

Option Name                                     Result
ShadowImage in-system replication              Unlock

The process was completed.
%
```

ShadowImage is installed and the status is "Enable". Installation of ShadowImage is now complete.

Uninstalling ShadowImage

To uninstall ShadowImage, the key code provided with the optional feature is required. Once uninstalled, ShadowImage cannot be used again until it is installed using the key code or key file.

To uninstall ShadowImage:

1. All ShadowImage pairs must be released (the status of all LUs are Simplex) before uninstalling ShadowImage.
2. From the command prompt, register the array in which the ShadowImage is to be uninstalled, then connect to the array.
3. Execute the `auopt` command to uninstall ShadowImage. For example:

```
% auopt -unit subsystem-name -lock on -keycode downloaded-48
characters-key code
Are you sure you want to lock the option?
(y/n [n]): y
The option is locked.
%
```

4. Execute the `auopt` command to confirm whether ShadowImage has been uninstalled. For example:

```
% auopt -unit subsystem-name -refer
DMEC002015: No information displayed.
%
```

Uninstalling ShadowImage is now complete.

Enabling or disabling ShadowImage

Once ShadowImage is installed, it can be enabled or disabled.

The following describes the enabling/disabling procedure.

1. If you are disabling ShadowImage, all pairs must be released (the status of all LUs are Simplex).
2. From the command prompt, register the array in which the status of the feature is to be changed, then connect to the array.
3. Execute the `auopt` command to change the status (enable or disable). The following is an example of changing the status from enable to disable. If you want to change the status from disable to enable, enter enable after the `-st` option, as shown in the example below.

```
% auopt -unit subsystem-name -option SHADOWIMAGE -st disable
Are you sure you want to disable the option? (y/n [n]): y
The option has been set successfully.
%
```

4. Execute the `auopt` command to confirm whether the status has been changed. For example:

```
% auopt -unit subsystem-name -refer
Option NameType      Term      Status
SHADOWIMAGEPermanent ---        Disable
%
```

Enabling or disabling ShadowImage is now complete.

Setting the DMLU

The DMLU is an exclusive volume used for storing ShadowImage information when the array is powered down. You must set up the DMLU before using ShadowImage.

Prerequisites

- LUs for the DMLUs must be set up and formatted.
- DMLU size must be at least 10 GB.
- One DMLU is needed but two are recommended, with the second used as a backup.
- For RAID considerations, see the bullet on DMLU in [Cascading ShadowImage with TrueCopy on page 2-18](#).
- Also see [Appendix B, Operations using CLI](#).

To set up DMLU

1. From the command prompt, register the array on which you want to create the DMLU and connect to that array.
2. Execute the `auDM-LU` command to create a DMLU. This command first displays LUs that can be assigned as DMLUs and later creates a DMLU. For example:

```
% auDM-LU -unit array-name -availablelist
Available Logical Units
  LUN Capacity    RAID Group DP Pool   RAID Level  Type  Status
    0    10.0 GB         0    N/A    5( 4D+1P)  SAS   Normal
%
% auDM-LU -unit array-name -set -lu 0
Are you sure you want to set the DM-LU? (y/n [n]): y
The DM-LU has been set successfully.
%
```

3. To release an already set DMLU, specify the `-rm` and `-lu` options in the `auDM-LU` command. For example:

```
% auDM-LU -unit array-name -rm -lu 0
Are you sure you want to release the DM-LU? (y/n [n]): y
The DM-LU has been released successfully.
%
```


Setting the ShadowImage I/O switching mode

The following procedure explains how to set the ShadowImage I/O switching mode to ON. For more information, see [Appendix D, I/O switching mode feature](#).

To set the ShadowImage I/O Switching Mode

1. From the command prompt, register the array on which you want to set the ShadowImage I/O Switching Mode. Connect to the array.
2. Execute the `ausystemparam` command.

When you want to reset the ShadowImage I/O Switching Mode, enter `disable` following the `-set -ShadowImageIOSwitch` option. For example:

:

```
% ausystemparam -unit array-name -set -ShadowImageIOSwitch
enable
Are you sure you want to set the system parameter? (y/n [n]): y
The system parameter has been set successfully.
%
```

3. Execute the `ausystemparam` command to verify that the ShadowImage I/O Switching Mode has been set. For example:

:

```
% ausystemparam -unit array-name -refer
Options
Turbo LU Warning = OFF
:
:
ShadowImage I/O Switch Mode = ON
:
:
Operation if the Processor failures Occurs = Reset a Fault
:
:
%
```



NOTE: When turning off the I/O Switching Mode, pair status must be other than Split (S-VOL Switch) and Synchronizing (S-VOL Switch).

ShadowImage operations

The `aureplicationlocal` command operates ShadowImage pair. To refer the `aureplicationlocal` command and its options, type in `aureplicationlocal -help` at the command prompt.

Confirming pairs status

To confirm the ShadowImage pairs, use the `aureplicationlocal -refer` command.

1. From the command prompt, register the array on which you want to confirm the ShadowImage pairs. Connect to the array.

2. Execute the `aureplicationlocal` command to confirm the ShadowImage pairs, as shown in the example below.

```
% aureplicationlocal -unit subsystem-name -refer -si
Pair Name                LUN Pair LUN Status
    Copy Type      Group
SI_LU1020_LU1021                1020      1021 Paired(
    0%)                      S
ShadowImage ---:Ungrouped
%

```

Creating ShadowImage pairs

The following procedure explains how to create on pair. To create pairs in a group, refer to [Creating ShadowImage pairs that belong to a group on page B-10](#).

To create the ShadowImage pairs, use the `aureplicationlocal -create` command.

1. From the command prompt, register the array on which you want to create the ShadowImage pairs. Connect to the array.
2. Execute the `aureplicationlocal` command to create the ShadowImage pairs.

When you want to automatically split the pair immediately after creation is completed, create the pair specifying the `-compsplit` option. In this case, the pair status immediately after pair creation becomes Split Pending.

In the following example, the P-VOL LUN is 1020 and the S-VOL LUN is 1021.

```
% aureplicationlocal -unit subsystem-name -create -si -pvol
    1020 -svol 1021
Are you sure you want to create pairs "SI_LU1020_LU1021"?
(y/n [n]): y
The pair has been created successfully.
%

```

3. Verify the pair status, as shown in the example below.

```
% aureplicationlocal -unit subsystem-name -refer -si
Pair Name                LUN Pair LUN Status
    Copy Type      Group
SI_LU1020_LU1021                1020      1021 Synchronizing(
    40%)                      ShadowImage ---:Ungrouped
%

```

The ShadowImage pair is created.

Splitting ShadowImage pairs

To split the ShadowImage pairs, use the `aureplicationlocal -split` command.

1. From the command prompt, register the array on which you want to split the ShadowImage pairs. Connect to the array.
2. Execute the `aureplicationlocal` command to split the ShadowImage pairs.

When you want to split Quick Mode, split the pair specifying the `-quick` option. In this case, the pair status immediately after pair splitting becomes Split Pending.

In the following example, the P-VOL LUN is 1020 and the S-VOL LUN is 1021.

```
% aureplicationlocal -unit subsystem-name -split -si -pvol
    1020 -svol 1021
Are you sure you want to split pairs?
(y/n [n]): y
The pair has been split successfully.
%
```

3. Verify the pair status as shown in the example below.

```
% aureplicationlocal -unit subsystem-name -refer -si
Pair Name          LUN Pair LUN Status
    Copy Type      Group
SI_LU1020_LU1021    1020   1021 Split(100%)
    ShadowImage    ---:Ungrouped
%
```

The ShadowImage pair is split.

Re-synchronizing ShadowImage pairs

To re-synchronize the ShadowImage pairs, use the `aureplicationlocal -resync` command.

1. From the command prompt, register the array on which you want to re-synchronize the ShadowImage pairs. Connect to the array.
2. Execute the `aureplicationlocal` command to re-synchronize the ShadowImage pairs.

When you want to re-synchronize Quick Mode, re-synchronize the pair specifying the `-quick` option. In this case, the pair status immediately after pair re-synchronizing becomes Paired Internally Synchronizing.

In the following example, the P-VOL LUN is 1020 and the S-VOL LUN is 1021.

```
% aureplicationlocal -unit subsystem-name -resync -si -pvol
    1020 -svol 1021
Are you sure you want to re-synchronize pairs?
(y/n [n]): y
The pair has been re-synchronized successfully.
%
```

3. Verify the pair status as shown in the example below.

```
% aureplicationlocal -unit subsystem-name -refer -si
Pair Name          LUN Pair LUN Status
    Copy Type      Group
SI_LU1020_LU1021    1020    1021 Synchronizing(
    40%)           ShadowImage ---:Ungrouped
%
```

The ShadowImage pair is resynchronized.

Restoring the P-VOL

To restore the ShadowImage pairs, use the `aureplicationlocal -restore` command.

1. From the command prompt, register the array on which you want to restore the ShadowImage pairs. Connect to the array.
2. Execute the `aureplicationlocal` command to restore the ShadowImage pairs.

In the following example, the P-VOL LUN is 1020 and the S-VOL LUN is 1021.

```
% aureplicationlocal -unit subsystem-name -restore -si -pvol
    1020 -svol 1021
Are you sure you want to restore pairs?
(y/n [n]): y
The pair has been restored successfully.
%
```

3. Verify the pair status as shown in the example below.

```
% aureplicationlocal -unit subsystem-name -refer -si
Pair Name                LUN Pair LUN Status
Copy Type    Group
SI_LU1020_LU1021                1020    1021 Reverse
synchronizing( 40%) ShadowImage ---:Ungrouped
%
```

The ShadowImage pair is restored.

Releasing ShadowImage pairs

To release the ShadowImage pairs, use the `aureplicationlocal -simplex` command.

1. From the command prompt, register the array on which you want to release the ShadowImage pairs. Connect to the array.
2. Execute the `aureplicationlocal` command to release the ShadowImage pairs.

In the following example, the P-VOL LUN is 1020 and the S-VOL LUN is 1021.

```
% aureplicationlocal -unit subsystem-name -simplex -si -pvol
1020 -svol 1021
Are you sure you want to release pairs?
(y/n [n]): y
The pair has been released successfully.
%
```

3. Verify the pair status as shown in the example below.

```
% aureplicationlocal -unit subsystem-name -refer -si
DMEC002015: No information displayed.
%
```

The ShadowImage pair is released.

Editing pair information

You can change the pair name, group name, and/or copy pace.

To change the pair information:

1. From the command prompt, register the array on which you want to change the ShadowImage pair information. Connect to the array.
2. Execute the `aureplicationlocal` command to change the ShadowImage pair information.

In the following example, the P-VOL LUN is 1020 and the S-VOL LUN is 1021.

```
% aureplicationlocal -unit subsystem-name -chg -si -pace slow
    -pvol 1020 -svol 1021
Are you sure you want to change pair information?
(y/n [n]): y
The pair information has been changed successfully.
%
```

The ShadowImage pair information is changed.

Creating ShadowImage pairs that belong to a group

To create multiple ShadowImage pairs that belong to a group:

1. Create the first pair that belongs to a group specifying an unused group number for the new group with the `-gno` option. The new group has been created and in this group, the new pair has been created too.

```
% aureplicationlocal -unit array-name -create -si -pvol 1020
    -svol 1021 -gno 20
Are you sure you want to create pairs "SI_LU1020_LU1021"?
(y/n [n]): y
The pair has been created successfully.
%
```

2. Add the name to the group if necessary using command to change the pair information.

```
% aureplicationlocal -unit array-name -chg -si -gno 20 -
    newgname group-name
Are you sure you want to change pair information?
(y/n [n]): y
The pair information has been changed successfully.
%
```

3. Create the next pair that belongs to the created group specifying the number of the created group with `-gno` option.
4. By repeating the step 3, the multiple pairs that belong to the same group can be created.



NOTE: You cannot use the options of the group number specification and automatic split after pair creation at the same time. To create two or more pairs that utilize the group by using Quick Mode, create all pairs belonging to the group, specify the quick option, and execute the split by group unit.

Splitting ShadowImage pairs that belong to a group

To split two or more ShadowImage pairs that belong to a group:

1. Execute the **aureplicationlocal** command to split the ShadowImage pairs. Display the status of the pairs belonging to the group to be the target, and split the pairs after checking that all pairs are in the split-able status.

```
% aureplicationlocal -unit array-name -refer -si
Pair Name          LUN Pair LUN Status
Copy Type          Group
SI_LU1000_LU1003   1000      1003 Paired(100%)
ShadowImage 0:
SI_LU1001_LU1004   1001      1004 Paired(100%)
ShadowImage 0:
SI_LU1002_LU1005   1002      1005 Paired(100%)
ShadowImage 0:
%
% aureplicationlocal -unit array-name -si -split -gno 0
Are you sure you want to split pair?
(y/n [n]): y
The pair has been split successfully.
%
```



NOTE: If a pair that cannot be split is mixed in the specified group, the pair split in the group unit does not operate. When this occurs, an error as a response to the pair split operation may or may not be displayed. Also, the pair split-able status differs depending on whether Quick Mode is used or not. Therefore, check that all the pairs belonging to the group to be the pair split target are in the following statuses according to each case.

When using Quick Mode:

- Paired
- Paired Internally Synchronizing
- Synchronizing

When not using the Quick Mode:

- Paired
- Paired Internally Synchronizing

2. Verify the pair status executing the **aureplicationlocal** command.

```
% aureplicationlocal -unit array-name -refer -si
Pair Name          LUN Pair LUN Status
Copy Type          Group
SI_LU1000_LU1003   1000      1003 Split(100%)
ShadowImage 0:
SI_LU1001_LU1004   1001      1004 Split(100%)
ShadowImage 0:
SI_LU1002_LU1005   1002      1005 Split(100%)
ShadowImage 0:
%
```

The ShadowImage pair is split.

Sample back up script for Windows

This section provides sample script for backing a volume on Windows.

```
echo off
REM Specify the registered name of the arrays
set UNITNAME=Array1
REM Specify the group name (Specify "Ungroup" if the pair
    doesn't belong to any group)
set G_NAME=Ungrouped
REM Specify the pair name
set P_NAME=SI_LU0001_LU0002
REM Specify the directory path that is mount point of P-VOL
    and S-VOL
set MAINDIR=C:\main
set BACKUPDIR=C:\backup
REM Specify GUID of P-VOL and S-VOL
PVOL_GUID=xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxxxx
SVOL_GUID=yyyyyyyyy-yyyy-yyyy-yyyy-yyyyyyyyyyyyy

REM Unmounting the S-VOL
pairedisplay -x umount %BACKUPDIR%
REM Re-synchronizing pair (Updating the backup data)
aureplicationlocal -unit %UNITNAME% -si -resync -pairname
    %P_NAME% -gname %G_NAME%
aureplicationmon -unit %UNITNAME% -evwait -si -pairname
    %P_NAME% -gname %G_NAME% -st paired -pvol

REM Unmounting the P-VOL
pairedisplay -x umount %MAINDIR%
REM Splitting pair (Determine the backup data)
aureplicationlocal -unit %UNITNAME% -si -split -pairname
    %P_NAME% -gname %G_NAME%
aureplicationmon -unit %UNITNAME% -evwait -si -pairname
    %P_NAME% -gname %G_NAME% -st split -pvol
REM Mounting the P-VOL
pairedisplay -x mount %MAINDIR% Volume{%PVOL_GUID%}

REM Mounting the S-VOL
pairedisplay -x mount %BACKUPDIR% Volume{%SVOL_GUID%}
< The procedure of data copy from C:\backup to backup
    appliance>
```



NOTE: For Windows 2000 or Windows Server 2003/2008 environments, The CCI mount/unmount commands must be used when mounting/unmounting a volume.



Operations using CCI

This chapter describes basic CCI procedures for setting up and performing ShadowImage operations.

- ❑ [Setting up CCI](#)
- ❑ [ShadowImage operations using CCI](#)
- ❑ [Pair, group name differences in CCI and Navigator 2](#)

Setting up CCI

The following sub-sections describe necessary set up procedures for CCI for ShadowImage.

Setting the command device

The command device(s) and LU mapping setting is used Navigator 2.

To designate command device(s)

1. From the command prompt, register the array to which you want to set the command device. Connect to the array.
2. Execute the `aucmddev` command to set a command device. First, display LUs to be assignable command device, and later set a command device. When you want to use the protection function of CCI, enter `enable` following the `-dev` option.

The following example specifies LU 2 for command device 1.

```
% aucmddev -unit array-name -availablelist
Available Logical Units
  LUN  Capacity  RAID Group  DP Pool  RAID Level  Type Status
    2   35.0 MB        0      N/A      6( 9D+2P)  SAS  Normal
    3   35.0 MB        0      N/A      6( 9D+2P)  SAS  Normal
%
% aucmddev -unit array-name -set -dev 1 2
Are you sure you want to set the command devices?
(y/n [n]): y
The command devices have been set successfully.
%
```

3. Execute the `aucmddev` command to verify that the command device has been set. For example:

```
% aucmddev -unit array-name -refer
Command Device LUN  RAID Manager Protect
1              2   Disable
%
```



NOTE: To set the alternate command device function or to avoid data loss and array downtime, designate two or more command devices. For details on alternate Command Device function, refer to the *Hitachi Adaptable Modular Storage Command Control Interface (CCI) User's Guide*.

4. The following example releases a command device:

```
% aucmddev -unit array-name -rm -dev 1
Are you sure you want to release the command devices?
(y/n [n]): y
This operation may cause CCI, which is accessing this command
device, to freeze.
Stop the CCI, which is accessing this command device, before
performing this operation.
Are you sure you want to release the command devices? (y/n [n]): y
The specified command device will be released.
Are you sure you want to execute? (y/n [n]): y
The command devices have been released successfully.
%
```

5. To change an already set command device, release the command device, then change the LU number. The following example specifies LU 3 for command device 1.

```
% aucmddev -unit array-name -set -dev 1 3
Are you sure you want to set the command devices?
(y/n [n]): y
The command devices have been set successfully.
%
```

Setting LU mapping

If using iSCSI, use the `autargetmap` command instead of the `auhgmap` command used with fibre channel.

To set up LU Mapping

1. From the command prompt, register the array to which you want to set the LU Mapping, then connect to the array.
2. Execute the `auhgmap` command to set the LU Mapping. The following is an example of setting LU 0 in the array to be recognized as 6 by the host. The port is connected via target group 0 of port 0A on controller 0.

```
% auhgmap -unit array-name -add 0 A 0 6 0
Are you sure you want to add the mapping information?
(y/n [n]): y
The mapping information has been set successfully.
%
```

3. Execute the `auhgmap` command to verify that the LU Mapping is set. For example:

```
% auhgmap -unit array-name -refer
Mapping mode = ON
Port  Group  H-LUN    LUN
  0A      0      6        0
%
```

Defining the configuration definition file

The configuration definition file describes the system configuration. It is required to make CCI operational. The configuration definition file is a text file created and/or edited using any standard text editor. It can be defined from the PC where CCI software is installed.

A sample configuration definition file, HORCM_CONF, is included with the CCI software. It should be used as the basis for creating your configuration definition file(s). The system administrator should copy the sample file, set the necessary parameters in the copied file, and place the copied file in the proper directory. For details on configuration definition file, refer to the *Hitachi Adaptable Modular Storage Command Control Interface (CCI) User's Guide*.

The configuration definition file can be automatically created using the `mkconf` command tool. For details on the `mkconf` command, refer to the *Hitachi Adaptable Modular Storage Command Control Interface (CCI) Reference Guide*. However, the parameters, such as `poll(10ms)` must be set manually (see step 4 below).

To define the configuration definition file

The following is an example that manually defines the configuration definition file. The system is configured with two instances within the same Windows host.

1. On the host where CCI is installed, verify that CCI is not running. If CCI is running, shut it down using the `horcmshutdown` command.
2. In the command prompt, make two copies of the sample file (`horcm.conf`). For example:

```
c:\HORCM\etc> copy \HORCM\etc\horcm.conf \WINDOWS\horcm0.conf
c:\HORCM\etc> copy \HORCM\etc\horcm.conf \WINDOWS\horcm1.conf
```

3. Open `horcm0.conf` using the text editor.
4. In the **HORCM_MON** section, set the necessary parameters.
Important: A value more than or equal to 6000 must be set for `poll(10ms)`. Specifying the value incorrectly may cause resource contention in the internal process, resulting the process temporarily suspending and pausing the internal processing of the array.
5. In the **HORCM_CMD** section, specify the physical drive (command device) on the array. Figure C-1 and Figure C-2 show examples of the `horcm0.conf` file in which the ShadowImage P-VOL-to-S-VOL ratio is 1:1 and 1:3, respectively.




Figure C-1: Horcm0.conf Example 1 (P-VOL: S-VOL=1: 1)

```

horcm0.conf - Notepad
File Edit Search Help

HORCM_MON
#ip_address      service      poll(10ms)      timeout(10ms)
XXXXXXXXX        5000          12000          3000

HORCM_CMD
#dev_name        dev_name        dev_name
\\.\PHYSICALDRIVE1

HORCM_DEV
#dev_group      dev_name      port#      TargetID      LU#      MU#
UG01            oradb1        CL1-A      1             1       0
UG02            oradb2        CL1-A      1             1       1
UG03            oradb3        CL1-A      1             1       2

HORCM_INST
#dev_group      ip_address      service
UG01            XXXXXXXXX      5001

```

Figure C-2: Horcm0.conf Example 2 (P-VOL: S-VOL=1: 3)

6. Save the configuration definition file and use the `horcmstart` command to start CCI.
7. Execute the `raidscan` command and write down the target ID displayed in the execution result.
8. Shut down CCI and then open the configuration definition file again.
9. In the **HORCM_DEV** section, set the necessary parameters. For the target ID, set the ID of the `raidscan` result you wrote down. Also, the item MU# must be added after the LU#.
10. In the **HORCM_INST** section, set the necessary parameters, and then save (overwrite) the file.
11. Repeat Steps 3 to 10, using Figure C-3 and Figure C-4 for examples.

```

horcm1.conf - Notepad
File Edit Search Help

HORCM_MON
#ip_address      service      poll(10ms)      timeout(10ms)
XXXXXXXXX        5001          12000          3000

HORCM_CMD
#dev_name        dev_name        dev_name
\\.\PHYSICALDRIVE1

HORCM_DEV
#dev_group      dev_name      port#      TargetID      LU#      MU#
UG01            oradb1        CL1-A      1             2       0

HORCM_INST
#dev_group      ip_address      service
UG01            XXXXXXXXX      5000

```

Figure C-3: Horcm1.conf Example 3 (P-VOL: S-VOL=1: 1)

```

horcm1.conf - Notepad
File Edit Search Help

HORCH_MON
#ip_address      service      poll(10ns)      timeout(10ns)
XXXXXXXXX        5001          12000           3000

HORCH_CMD
#dev_name        dev_name        dev_name
\\.\PHYSICALDRIVE1

HORCH_DEV
#dev_group       dev_name        port#           TargetID        LUN#           MU#
UG01             oradb1          CL1-A           1               2              0
UG02             oradb2          CL1-A           1               3              0
UG03             oradb3          CL1-A           1               4              0

HORCH_INST
#dev_group       ip_address      service
UG01             XXXXXXXXX      5000

```

Figure C-4: Horcm1.conf Example 4 (P-VOL: S-VOL=1: 3)

12. Enter the following example lines in the command prompt to verify the connection between CCI and the array:

```

C:\>cd horcm\etc

C:\HORCM\etc>echo hd1-3 | .\ingraid
Harddisk 1 -> [ST] CL1-A Ser =85000174 LDEV = 0 [HITACHI ]
[DF600F-CM      ]
Harddisk 2 -> [ST] CL1-A Ser =85000174 LDEV = 1 [HITACHI ]
[DF600F          ]
                HORC = SMPL  HOMRCF[MU#0 = SMPL  MU#1 = NONE  MU#2
= NONE]
                RAID6[Group 1-0] SSID = 0x0000
Harddisk 3 -> [ST] CL1-A Ser =85000174 LDEV = 2 [HITACHI ]
[DF600F          ]
                HORC = SMPL  HOMRCF[MU#0 = SMPL  MU#1 = NONE  MU#2
= NONE]
                RAID6[Group 2-0] SSID = 0x0000

C:\HORCM\etc>

```

Setting the environment variable

To perform ShadowImage operations, you must set the environment variable for the execution environment. The following describes an example in which two instances are configured within the same Windows host.

1. Set the environment variable for each instance. Enter the following from the command prompt:

```
C:\HORCM\etc>set HORCMINST=0
```

2. To enable ShadowImage, the environment variable must be set as follows:

```
C:\HORCM\etc>set HORCC_MRCF=1
```

3. Execute the `horcmstart` script, and then execute the `pairedisplay` command to verify the configuration, as shown in the following example:

```

C:\HORCM\etc>horcmstart 0 1
starting HORCM inst 0
HORCM inst 0 starts successfully.
starting HORCM inst 1
HORCM inst 1 starts successfully.

C:\HORCM\etc>pairstat -g VG01
group    PairVOL(L/R) (Port#,TID, LU-M) ,Seq#,LDEV#.P/S,Status,
Seq#,P-LDEV# M
VG01     oradb1(L)      (CL1-A , 1,  1-0 )85000174      1.SMPL -----
,-----  ----- -
VG01     oradb1(R)      (CL1-A , 1,  2-0 )85000174      2.SMPL -----
,-----  ----- -

```

CCI setup for performing ShadowImage operations is now complete.

ShadowImage operations using CCI

Pair operation using CCI are shown in Figure C-5.




Figure C-5: ShadowImage Pair Status Transitions

Confirming pair status

Table C-1 shows the related CCI and Navigator 2 GUI pair status.

Table C-1: CCI/Navigator 2 GUI Pair Status

Description	CCI	Navigator 2
Status where a pair is not created.	SMPL	Simplex
Initial copy or resynchronization copy is in execution.	COPY	Synchronizing
Status that the copying is completed and the contents written to the P-VOL is reflected in the S-VOL.	PAIR	Paired
Status that the written contents are managed as difference data by split.	PSUS/ SSUS	Split
Status that the difference data is copied from the S-VOL to the P-VOL for restoration.	RCPY	Reverse Synchronizing
Status that suspends copying forcibly when a failure occurs.	PSUE	Failure

To confirm ShadowImage pairs

For the example below, the group name in the configuration definition file is VG01.

1. Execute the `pairdisplay` command to verify the pair status and the configuration. For example:

```
C:\HORCM\etc>pairdisplay -g VG01
Group  PairVol(L/R) (Port#,TID, LU-M) ,Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M
VG01   oradb1(L)   (CL1-A , 1, 1-0 )85000174 1.P-VOL PAIR,85000174 2 -
VG01   oradb1(R)   (CL1-A , 1, 2-0 )85000174 2.S-VOL PAIR,----- 1 -
```

The pair status is displayed. For details on the `pairdisplay` command and its options, refer to *Hitachi AMS Command Control Interface (CCI) Reference Guide*.

Creating pairs (paircreate)

To create ShadowImage pairs

1. Execute the `pairdisplay` command to verify that the status of the ShadowImage volumes is SMPL. The following example specifies the group name in the configuration definition file as VG01.

```
C:\HORCM\etc>pairdisplay -g VG01
Group  PairVol(L/R) (Port#,TID, LU-M) ,Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M
VG01   oradb1(L)   (CL1-A , 1, 1-0 )85000174 1.SMPL -----,----- ---- -
VG01   oradb1(R)   (CL1-A , 1, 2-0 )85000174 2.SMPL -----,----- ---- -
```

2. Execute the `paircreate` command, then execute the `pairevtwait` command to verify that the status of each volume is PAIR. When using the `paircreate` command, the `-c` option is copying pace, which can vary between 1-15. 6-10 (medium) is recommended. 1-5 is a slow pace, which is used when I/O performance must be prioritized. 11-15 is a fast pace, which is used when copying is prioritized. The following example shows the `paircreate` and `pairevtwait` commands.

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

3. Execute the `pairdisplay` command to verify the pair status and the configuration. For example:

```
C:\HORCM\etc>pairdisplay -g VG01
Group  PairVol(L/R) (Port#,TID, LU-M) ,Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M
VG01   oradb1(L)   (CL1-A , 1, 1-0 )85000174 1.P-VOL PAIR,85000174 2 -
VG01   oradb1(R)   (CL1-A , 1, 2-0 )85000174 2.S-VOL PAIR,----- 1 -
```

Pair creation using a consistency group

A consistency group insures that the data in two or more S-VOLs included in a group are of the same time. For more information, see [Consistency group \(CTG\) on page 1-11](#).

To create a pair using a consistency group

1. Execute the `pairdisplay` command to verify that the status of the ShadowImage volumes is SMPL. In the following example, the group name in the configuration definition file is VG01.

```
C:\HORCM\etc>pairdisplay -g VG01
Group  PairVol(L/R) (Port#,TID, LU-M) ,Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M
VG01   oradb1(L)   (CL1-A , 1, 1-0 )85000174 1.SMPL -----,----- ---- -
VG01   oradb1(R)   (CL1-A , 1, 2-0 )85000174 2.SMPL -----,----- ---- -
VG01   oradb2(L)   (CL1-A , 1, 3-0 )85000174 3.SMPL -----,----- ---- -
VG01   oradb2(R)   (CL1-A , 1, 4-0 )85000174 4.SMPL -----,----- ---- -
VG01   oradb3(L)   (CL1-A , 1, 5-0 )85000174 5.SMPL -----,----- ---- -
VG01   oradb3(R)   (CL1-A , 1, 6-0 )85000174 6.SMPL -----,----- ---- -
```

2. Execute the `paircreate -m grp` command, then execute the `pairevtwait` command to verify that the status of each volume is PAIR.

```
C:\HORCM\etc>paircreate -g VG01 -vl -m grp
C:\HORCM\etc>pairevtwait -g VG01 -s pair -t 300 10
pairevtwait : Wait status done.
```

3. Execute the `pairdisplay` command to verify the pair status and the configuration. For example:

```
C:\HORCM\etc>pairdisplay -g VG01
Group  PairVol(L/R) (Port#,TID, LU-M) ,Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M
VG01   oradb1(L)   (CL1-A , 1, 1-0 )85000174 1.P-VOL PAIR,85000174 2 -
VG01   oradb1(R)   (CL1-A , 1, 2-0 )85000174 2.S-VOL PAIR,----- 1 -
VG01   oradb2(L)   (CL1-A , 1, 3-0 )85000174 3.P-VOL PAIR,85000174 4 -
VG01   oradb2(R)   (CL1-A , 1, 4-0 )85000174 4.S-VOL PAIR,----- 3 -
VG01   oradb3(L)   (CL1-A , 1, 5-0 )85000174 5.P-VOL PAIR,85000174 6 -
VG01   oradb3(R)   (CL1-A , 1, 6-0 )85000174 6.S-VOL PAIR,----- 5 -
```

Splitting pairs (pairsplit)

To split ShadowImage pairs

1. Execute the `pairsplit` command to split the ShadowImage pair in the PAIR status. In the following example, the group name in the configuration definition file is VG01

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

2. Execute the `pairstatus` command to verify the pair status and the configuration.

```
C:\HORCM\etc>pairstatus -g VG01
Group   PairVol(L/R) (Port#,TID,LU-M) ,Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M
VG01    oradb1(L)   (CL1-A , 1, 1-0 )85000174      1.P-VOL PSUS,85000174      2  -
VG01    oradb1(R)   (CL1-A , 1, 2-0 )85000174      2.S-VOL SSUS,-----      1  -
```

When it is required to split two or more S-VOLs included in a group at the same time and to assure that data of the same time are stored in the S-VOLs, use the CTG. In order to use the CTG, create a pair adding the `-m grp` option with the `paircreate` command.

Resynchronizing pairs (pairresync)

To resynchronize ShadowImage pairs

1. Execute the `pairresync` command to resynchronize the ShadowImage pair, then execute the `pairevtwait` command to verify that the status of each volume is PAIR. When using the `-c` option (copy pace), see the explanation in [Creating pairs \(paircreate\) on page C-9](#). For the following example, the group name in the configuration definition file is VG01.

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

2. Execute the `pairstatus` command to verify the pair status and the configuration. For example:

```
C:\HORCM\etc>pairstatus -g VG01
Group   PairVol(L/R) (Port#,TID,LU-M) ,Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M
VG01    oradb1(L)   (CL1-A , 1, 1-0 )85000174      1.P-VOL PAIR,85000174      2  -
VG01    oradb1(R)   (CL1-A , 1, 2-0 )85000174      2.S-VOL PAIR,-----      1  -
```

Releasing pairs (pairsplit -S)

To release the ShadowImage pair and change the status to SMPL

1. Execute the `pairstatus` command to verify that the status of the ShadowImage pair is PAIR. In the following example, the group name in the configuration definition file is VG01.

```
C:\HORCM\etc>pairstatus -g VG01
Group   PairVol(L/R) (Port#,TID,LU-M) ,Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M
VG01    oradb1(L)   (CL1-A , 1, 1-0 )85000174      1.P-VOL PAIR,85000174      2  -
VG01    oradb1(R)   (CL1-A , 1, 2-0 )85000174      2.S-VOL PAIR,-----      1  -
```

2. Execute the `pairsplit (pairsplit -S)` command to release the ShadowImage pair.

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

3. Execute `pairstatus` command to verify that the pair status changed to SMPL. For example:

```
C:\HORCM\etc>pairstatus -g VG01
Group   PairVol(L/R) (Port#,TID, LU-M) ,Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M
VG01    oradb1(L)   (CL1-A , 1,  1-0 )85000174      1.SMPL -----,-----  ---- -
VG01    oradb1(R)   (CL1-A , 1,  2-0 )85000174      2.SMPL -----,-----  ---- -
```

Pair, group name differences in CCI and Navigator 2

Pairs and groups that were created using CCI will be displayed differently when status is confirmed in Navigator 2.

- Pairs created with CCI and defined in the configuration definition file display unnamed in Navigator 2.
- Groups defined in the configuration definition file are also different in Navigator 2.
- Pairs defined in a group on the configuration definition file using CCI are displayed in Navigator 2 as ungrouped.

For information about how to manage a group defined on the configuration definition file as a CTG, see the *Hitachi Adaptable Modular Storage Command Control Interface (CCI) Reference Guide*.

I/O switching mode feature

This appendix provides a description, specifications, and setup instructions for the I/O Switching Mode feature.

- ❑ [I/O switching mode feature](#)
- ❑ [Specifications](#)
- ❑ [Recommendations](#)
- ❑ [Enabling I/O switching mode](#)
- ❑ [Recovery from a drive failure](#)

I/O switching mode feature

When a P-VOL becomes inaccessible due to a drive failure, I/O Switching Mode automatically transfers input/output operations between the host and P-VOL to the host and S-VOL. I/O Switching Mode must be enabled.

This feature only operates under the following conditions:

- The pair must be in Paired status.
- Pair configuration must be P-VOL:S-VOL = 1:1.
- When the DP-VOLs created by Dynamic Provisioning are used for a ShadowImage P-VOL or S-VOL, the ShadowImage I/O Switching function cannot operate.

Figure D-1 illustrates I/O Switching Mode.




Figure D-1: I/O Switching Mode Function

The I/O Switching feature activates when a drive double failure (triple failures for RAID 6) occurs. At that time, the pair's status is changed to "Failure (S-VOL Switch)", and host read/write access is automatically transferred to the S-VOL.



NOTE: When I/O Switching is activated, all LUs in the associated RAID group become unformatted, whether or not they are in a ShadowImage pair.

Specifications

Table D-1 shows specifications for the ShadowImage I/O Switching Mode function.

Table D-1: I/O Switching Mode Specifications

Item	Specification
Preconditions for operation	• The function must be activated. • Pair status must be PAIR. • A pair must not be cascaded with a TrueCopy pair. • The pair configuration must be P-VOL: S-VOL=1:1. • The DP-VOL created by Dynamic Provisioning is not used for a P-VOL or S-VOL. • In ShadowImage I/O Switching function, DP-VOLs created by Dynamic Provisioning can be used for a P-VOL or an S-VOL of ShadowImage.
Scope of application	All ShadowImage pairs that satisfy the preconditions for operation. In ShadowImage I/O Switching function, DP-VOLs created by Dynamic Provisioning can be used for a P-VOL or an S-VOL of ShadowImage.
Access to a P-VOL	Execution of host I/O continues after a drive failure because a report is sent to the host from the S-VOL as if from the P-VOL.
Access to an S-VOL	An I/O instruction issued to an S-VOL results in an error.
Display of the status	In Navigator 2: When host I/O is switched to an S-VOL, pair status is displayed as PSUE (S-VOL Switch). When the pairresync command is executed, pair status changes to COPY (S-VOL Switch). In CCI (pairdisplay command): Even when host I/O is switched to an S-VOL, the pair status is displayed as PSUE. However, when the pairmon -allsnd -nowait command is issued, the code (internal code of the pair status) is displayed as 0x08. When the pairresync command is executed for a pair in the PSUE (S-VOL Switch) status, the pair status is displayed as RCPY.
Formatting	Quick formatting can only be performed when the pair status is PSUE (S-VOL Switch).
Notes	• The pairsplit, pairresync -restore or pairsplit -s commands cannot be performed when the status is PSUE (S-VOL Switch) or copying PSUE (S-VOL Switch) to P-VOL.

Recommendations

- Locate P-VOLs and S-VOLs in respective RAID groups. When both are located in the same RAID group, they can become unformatted in the event of a drive failure.
- When a pair is in the Paired status, as it must be for the I/O Switching Mode, performance is lower than when the pair is Split. Hitachi recommends assigning an LU that uses SAS/SSD drives to an S-VOL to assure best performance results.

Enabling I/O switching mode

To use CLI, see [Setting the ShadowImage I/O switching mode on page B-5](#).

The following procedure displays a Navigator 2 applet screen. If your system does not display the screen (it takes a minute or two to appear), see the GUI online Help for Advanced Settings.

To enable I/O Switching Mode

1. In Navigator 2, select the subsystem in which ShadowImage is to be operated, then click **Show & Configure Array**.
2. In the tree view, select the **Advanced Settings** icon.
3. Click **Open Advanced Settings**. The Array Unit screen displays, as shown in [Figure D-2](#). This may take a few minutes.

If you have problems with screen, see the following section.




Figure D-2: Array Unit Screen

4. Select **Configuration Settings**, then click **Set**. The Configuration Settings screen displays.
5. Click the **System Parameter** tab.
6. Click the **ShadowImage I/O Switch Mode** check box, then click **Apply**.
7. On the confirmation message, click **OK**.
8. Click **Close** on the **Configuration Settings** page.
9. Click **Close** on the subsequent message screen.



NOTE: When disabling I/O Switching Mode, pair statuses must be other than Failure (S-VOL Switch) and Synchronizing (S-VOL Switch).

About the array unit screen

This screen is an applet connected to the SNM2 Server. After 20 minutes elapses while displaying this applet screen, automatic logoff occurs. Therefore, when your operation is completed, close the screen.

If the applet screen does not display, the login to the SNM2 Server may have failed. In this case, the applet screen cannot be displayed again. The code: 0x000000000000b045 or "DMEG800003: The error occurred in connecting RMI server." is displayed on the applet screen. Take the following actions.

- Close the Web browser, stop the SNM2 Server, restart it, then navigate to the array.
- Close the Web browser; confirm the SNM2 Server is started. If it has stopped, start it and display the screen of the Array that you want to operate.
- Return to the Array screen after 20 minutes elapsed and display the screen of the Array you want to operate.

Recovery from a drive failure

When the I/O Switching Mode feature is used, recovery from a drive failure on which the P-VOL is located can be undertaken during the newly-established host read/write operations to the S-VOL. This section provides the basic procedure for recovery.

To recover from a drive failure after I/O Switching Mode is activated

1. After a drive double failure occurs (triple failure for RAID 6) in a P-VOL and I/O from the host is transferred via the I/O Switching Mode feature to the S-VOL, have the P-VOL drives replaced.
2. When the drives have been replaced, perform quick formatting of the P-VOL.
3. Perform a reverse-resync of the pair, copying S-VOL data to the P-VOL. Performance is lowered during the reverse-resync, but host I/O can be continued. When resynchronization is completed, the pair status becomes Paired.



Glossary

This glossary provides definitions for replication terms as well as terms related to the technology that supports your Hitachi modular array. Click the letter of the glossary section to display the related page.

A

array

A set of hard disks mounted in a single enclosure and grouped logically together to function as one contiguous storage space.

asynchronous

Asynchronous data communications operate between a computer and various devices. Data transfers occur intermittently rather than in a steady stream. Asynchronous replication does not depend on acknowledging the remote write, but it does write to a local log file. Synchronous replication depends on receiving an acknowledgement code (ACK) from the remote system and the remote system also keeps a log file.

B

background copy

A physical copy of all tracks from the source volume to the target volume.

bps

Bits per second, the standard measure of data transmission speeds.

A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z
-------------------	-------------------	-------------------	-------------------	-------------------	-------------------	-------------------	-------------------	-------------------	-------------------	-------------------	-------------------	-------------------	-------------------	-------------------	-------------------	-------------------	-------------------	-------------------	-------------------	-------------------	-------------------	-------------------	-------------------	-------------------	-------------------

C

cache

A temporary, high-speed storage mechanism. It is a reserved section of main memory or an independent high-speed storage device. Two types of caching are found in computers: memory caching and disk caching. Memory caches are built into the architecture of microprocessors and often computers have external cache memory. Disk caching works like memory caching; however, it uses slower, conventional main memory that on some devices is called a memory buffer.

capacity

The amount of information (usually expressed in megabytes) that can be stored on a disk drive. It is the measure of the potential contents of a device; the volume it can contain or hold. In communications, capacity refers to the maximum possible data transfer rate of a communications channel under ideal conditions.

CCI

See command control interface.

CLI

See command line interface.

cluster

A group of disk sectors. The operating system assigns a unique number to each cluster and then keeps track of files according to which clusters they use.

cluster capacity

The total amount of disk space in a cluster, excluding the space required for system overhead and the operating system. Cluster capacity is the amount of space available for all archive data, including original file data, metadata, and redundant data.

command control interface (CCI)

Hitachi's Command Control Interface software provides command line control of Hitachi array and software operations through the use of commands issued from a system host. Hitachi's CCI also provides a scripting function for defining multiple operations.

command devices

Dedicated logical volumes that are used only by management software such as CCI, to interface with the arrays. Command devices are not used by ordinary applications. Command devices can be shared between several hosts.

A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z
---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---

command line interface (CLI)

A method of interacting with an operating system or software using a command line interpreter. With Hitachi's Storage Navigator Modular Command Line Interface, CLI is used to interact with and manage Hitachi storage and replication systems.

concurrency of S-VOL

Occurs when an S-VOL is synchronized by simultaneously updating an S-VOL with P-VOL data AND data cached in the primary host memory. Discrepancies in S-VOL data may occur if data is cached in the primary host memory between two write operations. This data, which is not available on the P-VOL, is not reflected on to the S-VOL. To ensure concurrency of the S-VOL, cached data is written onto the P-VOL before subsequent remote copy operations take place.

concurrent copy

A management solution that creates data dumps, or copies, while other applications are updating that data. This allows end-user processing to continue. Concurrent copy allows you to update the data in the files being copied, however, the copy or dump of the data it secures does not contain any of the intervening updates.

configuration definition file

The configuration definition file describes the system configuration for making CCI operational in a TrueCopy Extended Distance Software environment. The configuration definition file is a text file created and/or edited using any standard text editor, and can be defined from the PC where the CCI software is installed. The configuration definition file describes configuration of new TrueCopy Extended Distance pairs on the primary or remote array.

consistency group (CTG)

A group of two or more logical units in a file system or a logical volume. When a file system or a logical volume which stores application data, is configured from two or more logical units, these multiple logical units are managed as a consistency group (CTG) and treated as a single entity. A set of volume pairs can also be managed and operated as a consistency group.

consistency of S-VOL

A state in which a reliable copy of S-VOL data from a previous update cycle is available at all times on the remote array. A consistent copy of S-VOL data is internally pre-determined during each update cycle and maintained in the remote data pool. When remote takeover operations are performed, this reliable copy is restored to the S-VOL, eliminating any data discrepancies. Data consistency at the remote site enables quicker restart of operations upon disaster recovery.

A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z
---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---

CRC

Cyclical Redundancy Checking. A scheme for checking the correctness of data that has been transmitted or stored and retrieved. A CRC consists of a fixed number of bits computed as a function of the data to be protected, and appended to the data. When the data is read or received, the function is recomputed, and the result is compared to that appended to the data.

CTG

See Consistency Group.

cycle time

A user specified time interval used to execute recurring data updates for remote copying. Cycle time updates are set for each array and are calculated based on the number of consistency groups CTG.

cycle update

Involves periodically transferring differential data updates from the P-VOL to the S-VOL. TrueCopy Extended Distance Software remote replication processes are implemented as recurring cycle update operations executed in specific time periods (cycles).

D

data pool

One or more disk volumes designated to temporarily store un-transferred differential data (in the local array or snapshots of backup data in the remote array). The saved snapshots are useful for accurate data restoration (of the P-VOL) and faster remote takeover processing (using the S-VOL).

data volume

A volume that stores database information. Other files, such as index files and data dictionaries, store administrative information (metadata).

differential data control

The process of continuously monitoring the differences between the data on two volumes and determining when to synchronize them.

differential data copy

The process of copying the updated data from the primary volume to the secondary volume. The data is updated from the differential data control status (the pair volume is under the suspended status) to the primary volume.

A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z
---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---

Differential Management Logical Unit (DMLU)

The volumes used to manage differential data in a array. In a TrueCopy Extended Distance system, there may be up to two DM logical units configured per array. For Copy-on-Write and ShadowImage, the DMLU is an exclusive volume used for storing data when the array system is powered down.

differential-data

The original data blocks replaced by writes to the primary volume. In Copy-on-Write, differential data is stored in the data pool to preserve the copy made of the P-VOL to the time of the snapshot.

disaster recovery

A set of procedures to recover critical application data and processing after a disaster or other failure. Disaster recovery processes include failover and failback procedures.

disk array

An enterprise storage system containing multiple disk drives. Also referred to as “disk array device” or “disk storage system.”

DMLU

See Differential Management-Logical Unit.

dual copy

The process of simultaneously updating a P-VOL and S-VOL while using a single write operation.

duplex

The transmission of data in either one or two directions. Duplex modes are full-duplex and half-duplex. Full-duplex is the simultaneous transmission of data in two direction. For example, a telephone is a full-duplex device, because both parties can talk at once. In contrast, a walkie-talkie is a half-duplex device because only one party can transmit at a time.

E

entire copy

Copies all data in the primary volume to the secondary volume to make sure that both volumes are identical.

extent

A contiguous area of storage in a computer file system that is reserved for writing or storing a file.

A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z
---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---

F

failover

The automatic substitution of a functionally equivalent system component for a failed one. The term failover is most often applied to intelligent controllers connected to the same storage devices and host computers. If one of the controllers fails, failover occurs, and the survivor takes over its I/O load.

fallback

Refers to the process of restarting business operations at a local site using the P-VOL. It takes place after the arrays have been recovered.

Fault tolerance

A system with the ability to continue operating, possibly at a reduced level, rather than failing completely, when some part of the system fails.

FC

See Fibre channel.

Fibre channel

A gigabit-speed network technology primarily used for storage networking.

firmware

Software embedded into a storage device. It may also be referred to as Microcode.

full duplex

The concurrent transmission and the reception of data on a single link.

G

Gbps

Gigabit(s) per second.

granularity of differential data

Refers to the size or amount of data transferred to the S-VOL during an update cycle. Since only the differential data in the P-VOL is transferred to the S-VOL, the size of data sent to S-VOL is often the same as that of data written to the P-VOL. The amount of differential data that can be managed per write command is limited by the difference between the number of incoming host write operations (inflow) and outgoing data transfers (outflow).

A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z
---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---

GUI

Graphical user interface.

I

I/O

Input/output.

initial copy

An initial copy operation involves copying all data in the primary volume to the secondary volume prior to any update processing. Initial copy is performed when a volume pair is created.

initiator ports

A port-type used for main control unit port of Fibre Remote Copy function.

IOPS

I/O per second.

iSCSI

Internet-Small Computer Systems Interface. A TCP/IP protocol for carrying SCSI commands over IP networks.

iSNS

Internet-Small Computer Systems Interface. A TCP/IP protocol for carrying SCSI commands over IP networks.

L

LAN

Local Area Network. A computer network that spans a relatively small area, such as a single building or group of buildings.

load

In UNIX computing, the system load is a measure of the amount of work that a computer system is doing.

logical

Describes a user's view of the way data or systems are organized. The opposite of logical is physical, which refers to the real organization of a system. A logical description of a file is that it is a quantity of data collected together in one place. The file appears this way to users. Physically, the elements of the file could live in segments across a disk.

A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z
---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---

logical unit

See logical unit number.

logical unit number (LUN)

An address for an individual disk drive, and by extension, the disk device itself. Used in the SCSI protocol as a way to differentiate individual disk drives within a common SCSI target device, like a disk array. LUNs are normally not entire disk drives but virtual partitions (or volumes) of a RAID set.

LU

Logical unit.

LUN

See logical unit number.

LUN Manager

This storage feature is operated through Storage Navigator Modular 2 software and manages access paths among host and logical units for each port in your array.

M

metadata

In sophisticated data systems, the metadata -- the contextual information surrounding the data -- will also be very sophisticated, capable of answering many questions that help understand the data.

microcode

The lowest-level instructions directly controlling a microprocessor. Microcode is generally hardwired and cannot be modified. It is also referred to as firmware embedded in a storage array.

Microsoft Cluster Server

Microsoft Cluster Server is a clustering technology that supports clustering of two NT servers to provide a single fault-tolerant server.

mount

To mount a device or a system means to make a storage device available to a host or platform.

mount point

The location in your system where you mount your file systems or devices. For a volume that is attached to an empty folder on an NTFS

A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z
---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---

file system volume, the empty folder is a mount point. In some systems a mount point is simply a directory.

P

pair

Refers to two logical volumes that are associated with each other for data management purposes (e.g., replication, migration). A pair is usually composed of a primary or source volume and a secondary or target volume as defined by the user.

pair splitting

The operation that splits a pair. When a pair is "Paired," all data written to the primary volume is also copied to the secondary volume. When the pair is "Split," the primary volume continues being updated, but data in the secondary volume remains as it was at the time of the split, until the pair is re-synchronized.

pair status

Internal status assigned to a volume pair before or after pair operations. Pair status transitions occur when pair operations are performed or as a result of failures. Pair statuses are used to monitor copy operations and detect system failures.

paired volume

Two volumes that are paired in a disk array.

parity

The technique of checking whether data has been lost or corrupted when it's transferred from one place to another, such as between storage units or between computers. It is an error detection scheme that uses an extra checking bit, called the parity bit, to allow the receiver to verify that the data is error free. Parity data in a RAID array is data stored on member disks that can be used for regenerating any user data that becomes inaccessible.

parity groups

RAID groups can contain single or multiple parity groups where the parity group acts as a partition of that container.

peer-to-peer remote copy (PPRC)

A hardware-based solution for mirroring logical volumes from a primary site (the application site) onto the volumes of a secondary site (the recovery site).

A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z
---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---

point-in-time logical copy

A logical copy or snapshot of a volume at a point in time. This enables a backup or mirroring application to run concurrently with the system.

pool volume

Used to store backup versions of files, archive copies of files, and files migrated from other storage.

primary or local site

The host computer where the primary volume of a remote copy pair (primary and secondary volume) resides. The term "primary site" is also used for host failover operations. In that case, the primary site is the host computer where the production applications are running, and the secondary site is where the backup applications run when the applications on the primary site fail, or where the primary site itself fails.

primary volume (P-VOL)

The storage volume in a volume pair. It is used as the source of a copy operation. In copy operations a copy source volume is called the P-VOL while the copy destination volume is called "S-VOL" (secondary volume).

P-VOL

See primary volume.

R**RAID**

Redundant Array of Independent Disks. A disk array in which part of the physical storage capacity is used to store redundant information about user data stored on the remainder of the storage capacity. The redundant information enables regeneration of user data in the event that one of the array's member disks or the access path to it fails.

Recovery Point Objective (RPO)

After a recovery operation, the RPO is the maximum desired time period, prior to a disaster, in which changes to data may be lost. This measure determines up to what point in time data should be recovered. Data changes preceding the disaster are preserved by recovery.

Recovery Time Objective (RTO)

The maximum desired time period allowed to bring one or more applications, and associated data back to a correct operational state. It defines the time frame within which specific business operations or data must be restored to avoid any business disruption.

A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z
---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---

remote or target site

Maintains mirrored data from the primary site.

remote path

A route connecting identical ports on the local array and the remote array. Two remote paths must be set up for each array (one path for each of the two controllers built in the array).

remote volume

In TrueCopy operations, the remote volume (R-VOL) is a volume located in a different array from the primary host array.

resynchronization

Refers to the data copy operations performed between two volumes in a pair to bring the volumes back into synchronization. The volumes in a pair are synchronized when the data on the primary and secondary volumes is identical.

RPO

See Recovery Point Objective.

RTO

See Recovery Time Objective.

S**SAS**

Serial Attached SCSI. An evolution of parallel SCSI into a point-to-point serial peripheral interface in which controllers are linked directly to disk drives. SAS delivers improved performance over traditional SCSI because SAS enables up to 128 devices of different sizes and types to be connected simultaneously.

SATA

Serial ATA is a computer bus technology primarily designed for the transfer of data to and from hard disks and optical drives. SATA is the evolution of the legacy Advanced Technology Attachment (ATA) interface from a parallel bus to serial connection architecture.

secondary volume (S VOL)

A replica of the primary volume (P-VOL) at the time of a backup and is kept on a standby array. Recurring differential data updates are performed to keep the data in the S-VOL consistent with data in the P-VOL.

A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z
---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---

SMPL

Simplex.

snapshot

A term used to denote a copy of the data and data-file organization on a node in a disk file system. A snapshot is a replica of the data as it existed at a particular point in time.

SNM2

See Storage Navigator Modular 2.

SSD

Solid State Disk (drive). A data storage device that uses solid-state memory to store persistent data. An SSD emulates a hard disk drive interface, thus easily replacing it in most applications.

Storage Navigator Modular 2

A multi-featured scalable storage management application that is used to configure and manage the storage functions of Hitachi arrays. Also referred to as "Navigator 2."

suspended status

Occurs when the update operation is suspended while maintaining the pair status. During suspended status, the differential data control for the updated data is performed in the primary volume.

S-VOL

See secondary volume.

S-VOL determination

Independent of update operations, S-VOL determination replicates the S-VOL on the remote array. This process occurs at the end of each update cycle and a pre-determined copy of S-VOL data, consistent with P-VOL data, is maintained on the remote site at all times.

T

target copy

A file, device, or any type of location to which data is moved or copied.

TrueCopy

Refers to the TrueCopy remote replication.

A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z
---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---

V

virtual volume (V-VOL)

In Copy-on-Write, a secondary volume in which a view of the primary volume (P-VOL) is maintained as it existed at the time of the last snapshot. The V-VOL contains no data but is composed of pointers to data in the P-VOL and the data pool. The V-VOL appears as a full volume copy to any secondary host.

volume

A disk array object that most closely resembles a physical disk from the operating environment's viewpoint. The basic unit of storage as seen from the host.

volume copy

Copies all data from the P-VOL to the S-VOL.

volume pair

Formed by pairing two logical data volumes. It typically consists of one primary volume (P-VOL) on the local array and one secondary volume (S-VOL) on the remote arrays.

V-VOL

See virtual volume.

V-VOLTL

Virtual Volume Tape Library.

W

WMS

Workgroup Modular Storage.

write order guarantee

Ensures that data is updated in an S-VOL, in the same order that it is updated in the P-VOL, particularly when there are multiple write operations in one update cycle. This feature is critical to maintain data consistency in the remote S-VOL and is implemented by inserting sequence numbers in each update record. Update records are then sorted in the cache within the remote system, to assure write sequencing.

write workload

The amount of data written to a volume over a specified period of time.

A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z
---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---

A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z
---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---

Index

A

assigning pairs to a consistency group [6-5](#)

B

backup requirements [2-3](#)

backup, protecting from read/write access [6-6](#)

business uses of S-VOLs [2-3](#)

C

Cascade Connection of ShadowImage with TrueCopy [1-16](#)

Cascade Restrictions with Data Pool of SnapShot [1-24](#)

Cascade Restrictions with ShadowImage P-VOL [1-21](#)

Cascade Restrictions with S-VOL of ShadowImage [1-23](#)

CCI

description [1-15](#)

CCI, using to

confirm pair status [C-8](#)

create pairs [C-9](#)

define the config def file [C-3](#)

release pairs [C-11](#)

resynchronize pairs [C-11](#)

set environment variable [C-6](#)

set LU mapping [C-3](#)

set the command device [C-2](#)

split pairs [C-11](#)

checking pair status [6-2, 7-2](#)

CLI

description [1-15](#)

CLI, using to

check pair status [B-5](#)

create a pair [B-6](#)

edit pair information [B-9](#)

enable, disable ShadowImage [B-3](#)

install ShadowImage [B-2](#)

release a pair [B-9](#)

restore the P-VOL [B-8](#)

resync a pair [B-8](#)

set up the DMLU [B-4](#)

split a pair [B-7](#)

uninstall ShadowImage [B-3](#)

Command Control Interface. See CCI

command device

recommendation for LUs [2-6](#)

releasing [C-2](#)

setting up [C-2](#)

Command Line Interface. See CLI

Concurrent Use with Dynamic Provisioning [2-10](#)

configuring ShadowImage [5-2](#)

Consistency Groups

creating, assigning pairs to [6-5](#)

description [1-11](#)

number allowed [A-3](#)

copy pace [6-3](#)

creating a pair [6-2](#)

D

definitions, pair status [7-2](#)

deleting a pair [6-9](#)

design workflow [2-2](#)

designating a command device [C-2](#)

Differential Management LUs. See DMLU

disabling ShadowImage [4-3](#)

DMLU

description [1-11](#)

recommendation for LUs [2-6](#)

setup [5-2](#)

setup, CLI [B-4](#)

E

editing pair information [6-9](#)

enabling ShadowImage [4-3](#)

G

Group Name, adding [6-5](#)

GUI, description [1-15](#)

GUI, using to

assign pairs to a Consistency Group [6-5](#)

- check pair status [6-2](#)
- create a pair [6-2](#)
- delete a pair [6-9](#)
- edit a pair [6-9](#)
- enable, disable ShadowImage [4-3](#)
- install ShadowImage [4-2](#)
- monitor pair status [7-2](#)
- restore the P-VOL [6-9](#)
- resync a pair [6-7](#)
- set up the DMLU [5-2](#)
- split a pair [6-6](#)
- uninstall ShadowImage [4-4](#)

H

- how long to keep S-VOL [2-3](#)
- how often to copy P-VOL [2-2](#)

I

- I/O Switching Mode
 - description [D-2](#)
 - enabling using GUI [D-4](#)
 - setup with CLI [B-5](#)
 - specifications [D-3](#)
- installation [4-2](#)
- interfaces for ShadowImage [1-15](#)

L

- license [A-4](#)
- lifespan, S-VOLs [2-3](#)

M

- monitoring ShadowImage [7-1](#)

N

- number of copies to make [2-4](#)

O

- overview [1-1](#)

P

- Pace field [6-3](#)
- pairs
 - assigning to a consistency group [6-5](#)
 - creating [6-2](#)
 - deleting [6-9](#)
 - editing [6-9](#)
 - number allowed [A-2](#)
 - recommendations for volumes [2-5](#)
 - resyncing [6-7](#)
 - splitting [6-6](#)
 - status definitions and checking [7-2](#)
- pairs resyncing [6-7](#)
- planning a ShadowImage system [2-1](#)
- planning workflow [2-2](#)
- platforms supported [3-4](#)
- primary volume [1-2](#)

P-VOL

- and S-VOL setup [5-2](#)
- and S-VOL, definition [1-3](#)
- restoring [6-9](#)

R

- RAID grouping for volume pairs [A-2](#)
- RAID level for volume pairs [A-2](#)
- releasing a command device [C-2](#)
- Replication Manager [1-15](#), [5-5](#), [5-6](#)
- requirements [3-2](#)
- restoring the P-VOL [6-9](#)
- resyncing a pair [6-7](#)

S

- secondary volume [1-2](#)
- ShadowImage
 - configuring [5-2](#)
 - enable, disable [4-3](#)
 - environment [1-2](#)
 - how it works [1-3](#)
 - installing [4-2](#)
 - interface [1-15](#)
 - maintaining [7-1](#)
 - overview [1-1](#)
 - plan and design [2-1](#)
 - specifications [A-2](#)
 - uninstalling [4-4](#)
 - using [6-1](#)
 - workflow [6-2](#)
- specifications [A-2](#)
- splitting the pair [6-6](#)
- status definitions [7-2](#)
- supported platforms [3-4](#)
- S-VOL
 - description [1-3](#)
 - frequency, lifespan, number of [2-2](#)
 - number allowed [A-2](#)
 - specifying as backup only [6-6](#)
 - updating [6-7](#)
 - using [6-10](#)
- system requirements [3-2](#)

U

- uninstalling ShadowImage [4-4](#)
- updating the S-VOL [6-7](#)
- using the S-VOL [6-10](#)

V

- volume pairs
 - description [1-4](#)
 - RAID levels and grouping [A-2](#)
 - setup recommendations [2-5](#)

Index-2

Hitachi Data Systems

Corporate Headquarters
750 Central Expressway
Santa Clara, California 95050-2627
U.S.A.
Phone: 1 408 970 1000
www.hds.com
info@hds.com

Asia Pacific and Americas

750 Central Expressway
Santa Clara, California 95050-2627
U.S.A.
Phone: 1 408 970 1000
info@hds.com

Europe Headquarters

Sefton Park
Stoke Poges
Buckinghamshire SL2 4HD
United Kingdom
Phone: + 44 (0)1753 618000
info.eu@hds.com



MK-97DF8129-09