

**Hitachi TagmaStore®
Adaptable Modular Storage
and Workgroup Modular Storage
ShadowImage™ In-System Replication
Software User's Guide**

© 2007 Hitachi Data Systems Corporation, ALL RIGHTS RESERVED

Notice: 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 Data Systems Corporation (hereinafter referred to as "Hitachi Data Systems").

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

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

Trademarks

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

TagmaStore is a trademark of Hitachi Data Systems Corporation.

Compaq, HP, HP-UX, and HP Tru64 are registered trademarks or trademarks of Hewlett-Packard Development Company, L.P.

IBM and AIX are registered trademarks of International Business Machines Corp. in the United States.

Linux is a registered trademark of Linus Torvalds.

Microsoft, Windows, Windows NT, and Microsoft Server are registered trademarks of Microsoft Corporation.

Red Hat is a registered trademark of Red Hat, Inc.

SGI and IRIX are trademarks or registered trademarks of Silicon Graphics Inc. in the United States.

Sun and Solaris are trademarks of Sun Microsystems, Inc.

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

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

Notice of Export Controls

Export of technical data contained in this document may require an export license from the United States government and/or the government of Japan. This technical data may not be used directly or indirectly (without U.S. Government permission), in the design, development, fabrication, production, testing or use of rocket systems, or unmanned air vehicles; or chemical or biological weapons; or nuclear weapons or explosives or operation of facilities or components thereof for chemical processing of irradiated special nuclear or source material, heavy water production, separation of isotopes of source and special nuclear material, or fabrication of nuclear reactor fuel containing plutonium, or in International Atomic Energy Agency (IAEA) un-safeguarded nuclear facilities.

Document Revision Level

Revision	Date	Description
MK-95DF709-00	June 2005	Initial Release
MK-95DF709-01	August 2005	Revision 1, supersedes and replaces MK-95DF709-00
MK-95DF709-02	January 2006	Revision 2, supersedes and replaces MK-95DF709-01
MK-95DF709-03	February 2006	Revision 3, supersedes and replaces MK-95DF709-02
MK-95DF709-04	April 2006	Revision 4, supersedes and replaces MK-95DF709-03
MK-95DF709-05	April 2006	Revision 5, supersedes and replaces MK-95DF709-04
MK-95DF709-06	May 2006	Revision 6, supersedes and replaces MK-95DF709-05
MK-95DF709-07	July 2006	Revision 7, supersedes and replaces MK-95DF709-06
MK-95DF709-08	November 2006	Revision 8, supersedes and replaces MK-95DF709-07
MK-95DF709-09	January 2007	Revision 9, supersedes and replaces MK-95DF709-08
MK-95DF709-10	May 2007	Revision 10, supersedes and replaces MK-95DF709-09
MK-95DF709-11	June 2007	Revision 11, supersedes and replaces MK-95DF709-10

Changes in this Revision

- Changed product version number to version 7.1 and higher
- Added Chapter 1 Overview of Hitachi Local Replication Solutions
- Added Chapter 2 Overview of ShadowImage In-system Replication Software
- Added Chapter 3 ShadowImage Requirements
- Added reference to Volume Migration product
- Added recommendations for setting differential management LUs in Table 3.1 and Table 3.2
- Added section 4.4 Concurrent use of Volume Migration
- Added Figure 4.11, Figure 4.12, Figure 4.13
- Added support and restrictions for concurrent use of Volume Migration in Table 3.2

Preface

This document describes and provides instructions for performing ShadowImage™ In-System Replication operations on the Hitachi TagmaStore® Adaptable Modular Storage and Workgroup Modular Storage array subsystem.

This document assumes the following:

- The user has a background in data processing and understands RAID storage subsystems and their basic functions,
- The user is familiar with the Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage array subsystem,
- The user has read and understands the *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Command Control Interface (CCI) User and Reference Guide*, MK-95DF701, and
- The user is familiar with Windows® 98, Windows NT®, and/or Windows® 2000 operating systems.

Notes:

- Throughout this manual, the term “ShadowImage” refers to the ShadowImage In-System Replication Software program.
- Throughout this manual, the term “Navigator” refers to the Storage Navigator Modular program.
- Throughout this manual, the term “TrueCopy” refers to the TrueCopy Synchronous Remote Replication Software program.
- Throughout this manual, the term “TCE” refers to the TrueCopy Extended Distance Remote Replication Software program.
- Throughout this manual, the term “SnapShot” refers to the Copy-on-Write SnapShot Software program.
- Throughout this manual, the term “Volume Migration” refers to the Modular Volume Migration.

Notice: The use of all Hitachi Data System products is governed by the terms of your agreement(s) with Hitachi Data Systems.

Software Revision

This document revision applies to Hitachi TagmaStore® Adaptable Modular Storage and Workgroup Modular Storage Products version 7.1 and higher.

Convention for Storage Capacity Values

Storage capacity values for hard disk drives (HDDs) are calculated based on the following values:

- 1 KB (kilobyte) = 1,000 bytes
- 1 MB (megabyte) = 1,000² bytes
- 1 GB (gigabyte) = 1,000³ bytes
- 1 TB (terabyte) = 1,000⁴ bytes

Referenced Documents

- *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Storage Navigator Modular Graphical User Interface (GUI) User's Guide*, MK-95DF711
- *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Storage Navigator Modular Command Line Interface (CLI)*, MK-95DF712
- *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Command Control Interface (CCI) User and Reference Guide*, MK-95DF701
- *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Copy-on-Write SnapShot Software User's Guide*, MK-95DF708
- *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage TrueCopy Synchronous Remote Replication Software User's Guide*, MK-95DF710

Comments

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

- E-mail: doc.comments@hds.com
- Fax: 858-695-1186
- Mail:
Technical Writing, M/S 35-10
Hitachi Data Systems
10277 Scripps Ranch Blvd.
San Diego, CA 92131

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

Contents

Chapter 1 Overview of Hitachi Local Replication Solutions

1.1	Replication & Disaster Preparedness Business Objectives	2
1.1.1	Disaster Scales	2
1.1.2	Disaster Preparedness Measures	3
1.1.3	RPO and RTO	3
1.1.4	Cost	4
1.2	Replication Concepts.....	5
1.2.1	Backup and Replication	5
1.2.2	Local Data Replication Options	6
1.2.3	Snapshots	7
1.2.4	Copy-on-Write Technology.....	7
1.2.5	Integrating Local and Remote Replication	8
1.3	Hitachi's Local Replication Product Portfolio.....	9
1.3.1	ShadowImage In-System Replication Software	9
1.3.2	Other Complementary Copy Solutions	9
1.4	ShadowImage Product Differentiation	11
1.4.1	Product Market Differentiation	11
1.4.2	Differentiation from Copy-on-Write SnapShot	12

Chapter 2 Overview of ShadowImage In-system Replication Software

2.1	Typical ShadowImage System Environment	16
2.2	ShadowImage Functional Components	17
2.2.1	ShadowImage Volume Pairs (P-VOLs and S-VOLs)	17
2.2.2	Consistency Group (CTG).....	18

Chapter 3 ShadowImage Requirements

3.1	ShadowImage Specifications	20
3.2	ShadowImage Requirements	22
3.3	ShadowImage Management Software.....	25
3.3.1	Command Control Interface Software (CCI).....	25
3.3.2	Storage Navigator Modular Software (SNM)	25
3.3.3	Optional Software	25

Chapter 4 Overview of ShadowImage Operations

4.1	ShadowImage Operations.....	28
4.1.1	Paircreate Operation	28
4.1.2	Pairsplit Operation	29
4.1.3	Pairresync Operations (Normal, Reverse)	29
4.1.4	Pairsplit -E Operation.....	31
4.1.5	Pairsplit -S Operation.....	32
4.2	ShadowImage Pair Status.....	32
4.3	Concurrent Use of LUN Size Expansion.....	34
4.4	Concurrent Use of Volume Migration	36
4.5	Cascade Connection of TrueCopy with ShadowImage	39
4.5.1	Cascade Restrictions with ShadowImage P-VOL	44
4.5.2	Cascade Restrictions with S-VOL of ShadowImage	46
4.5.3	Cascade Restrictions with POOL of SnapShot	47
4.5.4	Restrictions Configuration on Cascade of TrueCopy with ShadowImage	47
4.6	ShadowImage I/O Switching Function	48

4.6.1	ShadowImage I/O Switching Function Operation	48
4.6.2	Requirements for Using the ShadowImage I/O Switching Function	48
4.6.3	Recovery from a Double Drive Failure	49
4.6.4	Configuration and Operation	50
4.6.5	ShadowImage I/O Switching Function Specifications	51

Chapter 5 Preparing for ShadowImage Operations

5.1	Installing and Uninstalling ShadowImage	54
5.1.1	Installing ShadowImage	54
5.1.2	Uninstalling ShadowImage	58
5.1.3	Enabling or Disabling ShadowImage	59
5.2	Setting the Command Device	61
5.3	Setting the Differential Management LU	64
5.4	Setting the Target ID (LU Mapping)	66
5.4.1	Specifying Mapping Mode	66
5.4.2	Setting Mapping Information	67
5.5	Setting the ShadowImage I/O Switching Mode	71
5.6	Defining the Configuration Definition File	72
5.7	Setting the Environment Variable	75

Chapter 6 Performing ShadowImage Operations

6.1	Creating ShadowImage Pairs (paircreate)	78
6.2	Splitting ShadowImage Pairs (pairsplit)	79
6.3	Resynchronizing ShadowImage Pairs (pairresync)	80
6.4	Suspending ShadowImage Pairs (pairsplit -E)	80
6.5	Releasing ShadowImage Pairs (pairsplit -S)	82
6.6	Confirming the Status of Pairs	83
6.7	Troubleshooting	84

Chapter 7 ShadowImage Basic Specifications

7.1	Supported Capacities and Configurations	86
7.2	Functions	88
7.2.1	Pair Creation	89
7.2.2	Pair Splitting	89
7.2.3	Resynchronization (resync) from P-VOL to S-VOL	90
7.2.4	Resynchronization (reverse resync) from S-VOL to P-VOL	90
7.2.5	Pair Suspension	91
7.2.6	Pair Release	91
7.3	Resync Operation and Differential Data Copy Process	92
7.3.1	Managing Differential Data	92
7.3.2	Differential Data	92
7.4	Concurrent Use of LUN Expansion	93
7.5	Cascade Connection	96
7.5.1	Cascade Restrictions	103

Chapter 8 ShadowImage System Configuration

8.1	Path Connection	106
8.1.1	AMS500	106
8.2	Locating P-VOLs and S-VOLs	107
8.2.1	Location Recommendations for P-VOLs and S-VOLs	107
8.2.2	Locating Multiple LUs within Same Drive Column	108
8.2.3	Pair Status Differences when Setting Multiple Pairs	108

8.2.4	P-VOLs and S-VOLs in a RAID Configuration	109
8.2.5	Setting Multiple Command Devices	109
8.3	P-VOLS and S-VOLS in a RAID Configuration	110
8.4	Others (Example: Mode Setting)	111
8.5	Recommendations	112
8.5.1	Pair Assignment	112
8.5.2	Pair Creation	113
8.5.3	Resynchronization	113
Chapter 9	System Operation Example	
9.1	Full-Time Operation (Backup and Resync at Night)	116
9.2	Full-Time Split Operation	117
Chapter 10	Supported Platforms and Restrictions	
10.1	Supported Platforms	120
10.2	Restrictions	121
10.2.1	LU Mapping and ShadowImage Configuration	121
10.2.2	For VxVM and ShadowImage Configuration	121
10.2.3	MSCS and ShadowImage Configuration	121
10.2.4	AIX® and ShadowImage Configuration	121
10.2.5	Windows Server 2003 and ShadowImage TagmaStore Configuration	122
10.2.6	Linux and LVM Configuration	122
Chapter 11	Notes on System Configuration	
11.1	Previously Designed Procedure for Recovering Data during Malfunction	124
11.2	Applying ShadowImage to an Existing System	124
Appendix A	Operations Using CLI	
A.1	Installing ShadowImage	126
A.2	Uninstalling ShadowImage	127
A.3	Enabling or Disabling ShadowImage	128
A.4	Setting the Command Device	129
A.5	Setting Differential Management LU	131
A.6	Setting the Target ID (LU Mapping)	132
A.7	Setting the ShadowImage I/O Switching Mode	133
A.8	Confirming the Status of Pairs	134
Acronyms	135

List of Figures

Figure 1.1	Data Protection Options	5
Figure 1.2	SnapShot ShadowImage Product Differentiators	13
Figure 2.1	ShadowImage Environment	16
Figure 4.1	Adding a ShadowImage Pair	28
Figure 4.2	Normal Pairresync Operation	30
Figure 4.3	Reverse Pairresync Operation (pairresync-restore)	30
Figure 4.4	Reflecting Write Data to S-VOL during Reverse Pairresync Operation (pairresync-restore)	30
Figure 4.5	ShadowImage Pair Status Transitions	32
Figure 4.6	Example: A Unified LU Assigned to P-VOL or S-VOL	34
Figure 4.7	Example: Restriction Placed on Composing a Unified LU	34
Figure 4.8	Example: A Combination of RAID Levels	35
Figure 4.9	Example: A Combination of RAID Levels	35
Figure 4.10	Example: Restriction Placed on Unification of LUs Consisting of Different Drive Types	35
Figure 4.11	An Example in which the Number of Pairs of ShadowImage is Limited	36
Figure 4.12	An Example of a Case where the Copying Operation of ShadowImage is made to Wait	37
Figure 4.13	An Example of a Case where the Copying Operation of Volume Migration is made to Wait	38
Figure 4.14	Cascade Connection of TrueCopy with ShadowImage (P-VOL: S-VOL=1: 1)	39
Figure 4.15	Simultaneous Cascade Connection of TrueCopy with ShadowImage (P-VOL: S-VOL=1: 1)	40
Figure 4.16	Cascade Connection of TrueCopy with ShadowImage (P-VOL: S-VOL=1: 3)	42
Figure 4.17	Simultaneous Cascade Connection of TrueCopy with ShadowImage (P-VOL: S-VOL=1: 3)	43
Figure 4.18	Restrictions for Cascade Configuration of TrueCopy with ShadowImage	47
Figure 4.19	ShadowImage I/O Switching Function Operation	48
Figure 4.20	Recovering from a Double Drive Failure (Triple Failures for RAID 6)	49
Figure 5.1	Array System Viewer Window (Logical Status Tab)	55
Figure 5.2	Install/Unlock Options Dialog Box	55
Figure 5.3	Install/Unlock Options - Option Name Dialog Box	56
Figure 5.4	Result Dialog Box	57
Figure 5.5	Array System Viewer Window (Logical Status Tab, ShadowImage: Enable)	58
Figure 5.6	De-install/Lock Options Dialog Box	59
Figure 5.7	Array System Viewer Window (Logical Status Tab, ShadowImage: Disable)	60
Figure 5.8	Array System Viewer Window (Command Device Page: Before Setting)	61
Figure 5.9	Command Devices Settings Dialog Box (Before Setting)	62
Figure 5.10	Command Devices Settings Dialog Box (After Setting)	62
Figure 5.11	Array System Viewer Window (Command Device Page: After Setting)	63
Figure 5.12	Array System Viewer Window (Differential Management LU Page: Before Setting)	64
Figure 5.13	Select Logical Unit Dialog Box	64
Figure 5.14	Array System Viewer Window (Differential Management LU Page: After Setting)	65
Figure 5.15	Array System Viewer Window (Specifying Mapping Mode)	66
Figure 5.16	Mapping Mode Dialog Box	66
Figure 5.17	Array System Viewer Window (Setting Mapping Information)	67
Figure 5.18	Mapping Property Dialog Box (Before Setting)	68

Figure 5.19	Mapping Property Dialog Box (After Setting).....	69
Figure 5.20	Configuration Settings Dialog Box (System Parameter Page).....	71
Figure 5.21	Horcm0.conf Example (P-VOL: S-VOL=1: 1).....	73
Figure 5.22	Horcm0.conf Example (P-VOL: S-VOL=1: 3).....	73
Figure 5.23	Horcm1.conf Example (P-VOL: S-VOL=1: 1).....	74
Figure 5.24	Horcm1.conf Example (P-VOL: S-VOL=1: 3).....	74
Figure 6.1	Paircreate Command Example	78
Figure 6.2	Paircreate -m grp Command Example	79
Figure 6.3	Pairsplit Command Example.....	79
Figure 6.4	Pairresync Command Example	80
Figure 6.5	Pairsplit -E Command Example.....	80
Figure 6.6	Pairsplit -S Command Example.....	82
Figure 6.7	Array System Viewer Window (Pair Display)	83
Figure 7.1	Conceptual Configuration.....	86
Figure 7.2	Example of Pair Creation.....	89
Figure 7.3	Resynchronization (reverse resync) from S-VOL to P-VOL	90
Figure 7.4	Example of Differential Data	92
Figure 7.5	Example: Unified LU Assigned to P-VOL or S-VOL	93
Figure 7.6	Example: Restriction on a Number of LUs Composing a Unified LU	93
Figure 7.7	Example: Combination of RAID Levels.....	94
Figure 7.8	Example: Another Combination of RAID Levels Allowed	94
Figure 7.9	Example: Restriction on Unification of LUs Consisting of Different Types of Drives.....	95
Figure 7.10	Cascade Structure: ShadowImage P-VOL (P-VOL: S-VOL=1: 1)	96
Figure 7.11	Cascade Structure: ShadowImage P-VOL (P-VOL: S-VOL=1: 3)	97
Figure 7.12	Cascade Structure: ShadowImage S-VOL (P-VOL: S-VOL=1: 1)	98
Figure 7.13	Cascade Structure: ShadowImage S-VOL (P-VOL: S-VOL=1: 3)	99
Figure 7.14	Simultaneous Cascading with a ShadowImage P-VOL and S-VOL (P-VOL: S- VOL=1: 1)	101
Figure 7.15	Simultaneous Cascading with a ShadowImage P-VOL and S-VOL (P-VOL: S- VOL=1: 3)	102
Figure 7.16	Restrictions in Cascade Connection-1.....	103
Figure 7.17	Restrictions in Cascade Connection-2.....	103
Figure 7.18	Restrictions in Cascade Connection-3.....	104
Figure 8.1	Path Connection (AMS500).....	106
Figure 8.2	Locations for P-VOLs and S-VOLs (Not Recommended and Recommended)....	107
Figure 8.3	Locating Multiple LUs within Same Drive Column	108
Figure 9.1	Backup Using Tape	116
Figure 9.2	S-VOL (during split) connected to Tape for Backup.....	117

List of Tables

Table 1.1	Examples of Disaster Scales	3
Table 1.2	SnapShot-ShadowImage Differentiators	12
Table 3.1	External Specifications.....	20
Table 3.2	ShadowImage Requirements	22
Table 3.3	Characteristics of FC Drive and SATA Drive.....	24
Table 4.1	ShadowImage Pair Status	33
Table 4.2	A Read/Write Instruction to a P-VOL of ShadowImage on the Local Side.....	44
Table 4.3	A Read/Write Instruction to a P-VOL of ShadowImage on the Remote Side.....	45
Table 4.4	A Read/Write Instruction to an S-VOL of ShadowImage on the Local Side	46
Table 4.5	ShadowImage I/O Switching Function Specifications	51
Table 6.1	Icons for Pair Status	84
Table 7.1	Single Maximum Supported Capacity of ShadowImage (TB).....	87
Table 7.2	ShadowImage Functions	88
Table 8.1	P-VOLs and S-VOLs in a RAID Configuration.....	110
Table 8.2	Relationship between Copy Pace Parameters and Resync Process.....	111
Table 10.1	External Specifications.....	120

Chapter 1 Overview of Hitachi Local Replication Solutions

This chapter discusses local replication concepts; provides an overview of Hitachi's local replication solutions and an introduction to Hitachi's Copy-on-Write SnapShot solution. This chapter includes the following sections:

- Replication & Disaster Preparedness Business Objectives (see section 1.1)
- Replication Concepts (see section 1.2)
- Hitachi's Local Replication Product Portfolio (see section 1.3)
- ShadowImage Product Differentiation (see section 1.4)

1.1 Replication & Disaster Preparedness Business Objectives

Today, businesses are increasingly dependent on the availability of critical data resources. Most enterprises require access to business data 24/7 and cannot tolerate a downtime of more than a few hours, without serious impact to the bottom line. Beyond the loss of revenue, there are adverse headlines and the potential impact on company valuation to consider, along with lower employee productivity caused by sporadic outages.

Governments have also established mandates, requiring corporations to comply with business data protection regulations. For example, United States regulations on data protection now apply to health care (HIPAA), financial services (SEC 17a4), corporate accountability (Sarbanes-Oxley Act), life sciences (21 CFR Part 11), and government (DoD 5015.2-STD). Beyond government regulators, investors and even insurers insist that businesses should establish feasible disaster recovery and business continuity plans, to protect critical information. In addition, organizations constantly face increasing internal and external threats to the stability of information systems.

In today's extremely competitive business environment, organizations cannot afford to ignore the impact of all the above factors. Since business continuity and high availability have become the top concerns of businesses globally, it has become important for them to address the area of disaster preparedness and proactively plan for successful disaster recovery. Data protection and preparedness measures are often undertaken by analyzing the following business objectives:

1. Disaster scale: What are the assumptions for the scale of a disaster or failure?
2. Recovery Point Objective: Up to what point in time, should data be recovered in the event of a disaster?
3. Recovery Time Objective: How soon business operations should be resumed after the disaster?
4. Cost: How much can be budgeted for preparedness measures?
5. Backup window: How long can a backup process take to complete without impacting critical application processing?
6. Consistency groups: The number of multiple applications that need to be restored to the same point in time.

The overriding goals of any business continuity plan must be to survive a disaster and resume operations as quickly as possible. The best recovery plan to achieve these goals depends on how an organization chooses to prioritize the above business objectives. The following sections discuss each of the above business objectives in more detail.

1.1.1 Disaster Scales

Before determining necessary preparedness measures, assumptions must be made about the scale of the disaster. Disaster scales describe the scope of a disaster, in terms of damage to information systems and the overall business environment. These assumptions often determine how information systems need to be configured.

Table 1.1 shows examples of disaster scales. For example, when a file is accidentally erased, the damage is limited only to the specific file. On the other hand, in the event of a disaster such as an earthquake, the damage can easily extend to an entire city or region.

Table 1.1 Examples of Disaster Scales

Type of Disaster	Example	Disaster Scale
Data destruction	File erased by mistake	File
Drive failure	Dual failure on a drive on which RAID is being reconfigured	Logical Unit
Subsystem failure	Hardware failure in controller, etc.	Subsystem
Site disaster	Fire in building, etc	Site (within some hundreds of meters)
Natural disaster	Earthquake, hurricane, etc.	Extensive (hundreds of kilometers or more)

1.1.2 Disaster Preparedness Measures

Businesses are continuously faced with risks associated with catastrophes resulting from natural calamities, human operational errors, planned or unplanned downtime, security and terrorist incidents, content or application third-party dependencies, increased dependency on heterogeneous technologies etc. To eliminate these risks, it is necessary for organizations to undertake basic disaster preparedness measures. There are two fundamental types of disaster preparedness measures:

Data protection measures: These measures are steps taken to guard critical data against physical destruction (such as when a dual failure occurs on a drive), logical destruction (such as when a file is accidentally erased), site failures or natural disasters. One way to protect data is to periodically create data backups on physically independent secondary storage devices. If a primary data storage device is physically destroyed, most of the data can be retrieved from the latest backup data in the independent secondary storage device.

Another way to protect data is to create snapshots of primary data on the local storage device, at specific time intervals. In the case of logical destruction in the storage device, data can be restored from the point-in-time snapshots. Backup data can also be created at a remote location as a protection against site failures or natural disasters. This type of backup is called remote replication.

Disaster recovery measures: These measures are contingency plans that can be implemented in the event of a disaster. These measures help achieve a high RTO and may involve creating a clustering configuration to ensure work can be resumed using the copied data, immediately after a primary site failure is detected. Recovery measures may also involve establishing backup sites far from the location of a possible disaster.

Once the above business objectives are addressed and risks established, businesses can design resilient information systems, to ensure quick recovery time of critical application processing. Assessment based on the discussed business objectives also helps determine which applications require the most protection and the business impact of unexpected downtime.

1.1.3 RPO and RTO

Recovery Point Objective (RPO) and the Recovery Time Objective (RTO) are parameters that define how quickly an organization needs to recover to survive a disaster and how much data loss can be tolerated.

Recovery Point Objective (RPO) is defined by the Storage Networking Industry Association (SNIA) as follows:

- *The maximum desired time period prior to a failure or disaster during which changes to data may be lost as a consequence of recovery. Data changes preceding the failure or disaster by at least this time period are preserved by recovery. Zero is a valid value and is equivalent to a “zero data loss” requirement.*

In other words, RPO is a target value that specifies the point in time starting from which data needs to be recovered, in the event of a failure or disaster. It answers the question: How much data can your business afford to lose? For example, the RPO value for temporary data generated during a calculation is considered low because such data can be generated again if the calculation can be made again. On the other hand, if the transaction data of an Internet shopping site is lost, the image of the business suffers and, consequently, the RPO value is considered high.

Recovery Time Objective (RTO) is defined by the Storage Networking Industry Association (SNIA) as follows:

- *The maximum desired time period required to bring one or more applications and associated data back to a correct operational state.*

In other words, RTO defines the time frame within which specific business operations or data must be restored to avoid any business disruption. It answers the question: how long can a business afford to be down? RTO values for data are determined by how indispensable the data is for continuous site operation. For example, although auditing data is important for business operation and the RPO value is high, the RTO value of this data is considered low if the data is used only for auditing. On the other hand, the RTO value of the transaction data mentioned in the RPO example above is considered high because such data is indispensable to business continuity.

Both RTO and RPO influence the data replication and recovery options an organization needs to adopt. The Recovery Time Objective combined with the Recovery Point Objective of business data, can be mapped to a range of technical approaches, associated costs, and degrees of data protection.

1.1.4 Cost

Cost is the monetary requirement for building and operating a system which achieves its disaster scale, RPO, and RTO targets. Cost and the RPO and RTO values exist in a trade-off relationship. For example, cost increases as the loss of data in preparation for an extensive disaster is minimized and the target for a quick recovery is achieved. Recovery-time objective combined with a recovery point objective based on the value of particular data, map to a range of technical approaches, costs and degrees of data protection. The following figure illustrates the data protection choices based on cost versus value of data.

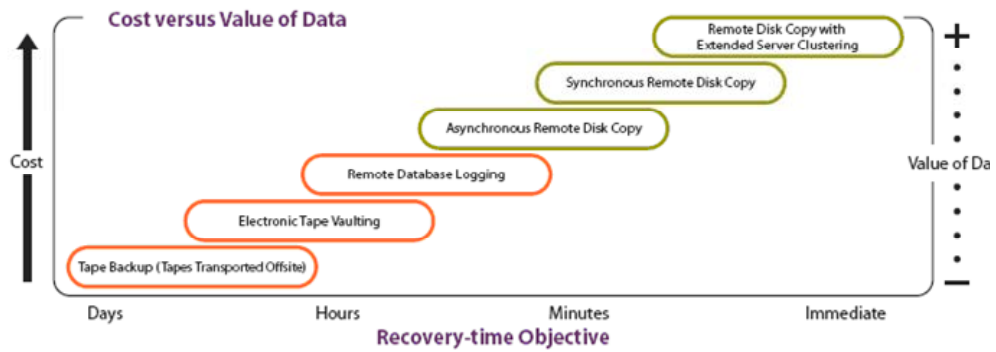


Figure 1.1 Data Protection Options

When determining cost it is also essential to analyze the cost associated with downtime. This cost can be estimated by determining the effects of downtime on factors such as employee productivity, business reputation, repair expenses, revenue and financial performance. Business processes and services must also be classified based on mission critical functions (customer-facing, revenue-producing, back-office etc) when determining business continuity and disaster recovery strategies.

1.2 Replication Concepts

An essential component of business continuity is the ability to quickly replicate and restore data. Replication provides the first line of protection and involves processes that create and maintain copies of data volume in the storage system. Replicated data can be leveraged for a variety of needs, well beyond those of data recovery, addressing such business challenges as backups, change/control, maintenance, for analytical purposes, etc. More importantly, using data replicas on a cyclical basis protects data from corruption, and allows for reverting to an uncorrupted version of data when needed.

Data replication increases data availability in the following ways:

- Allowing non-disruptive backup of current production data with no impact to the production application
- Speeding data restoration by replacing slow and labor-intensive tape-based restores with continuously available online backups
- Enabling frequent, non-disruptive disaster recovery testing with an online copy of current and accurate production data

The following sections discuss data protection concepts, local replication options and how local and remote replication solutions can be combined to provide near continuous data availability.

1.2.1 Backup and Replication

Traditional tape backup processes can be labor-intensive and span several hours. In a world of just-in-time delivery, global supply chains, around-the-clock customer demand, and competition, businesses cannot afford to gamble with the threat of planned or unplanned downtime. These tasks must be performed without impacting production data or causing business downtime.

Traditional tape backup also poses the following disaster recovery challenges:

- Inability to meet some requirements with really large databases and file systems
- Inability to restore multiple systems to the same point in time
- Difficult to back up remote offices
- Involve risks associated with physical media leaving business facility

In addition, most organizations deal with the following two conflicting demands:

- Necessity to backup increasing amounts of data
- Necessity to meet the business requirement for continuous data availability

The above challenges have led data protection and disaster recovery planners to switch to “replication” for protecting data and restoring critical applications. Addressing the need for continuous uptime and data availability, storage hardware manufacturers have responded with exceptionally fast server-less split mirror backup solutions. Software-based solutions using copy-on-write snapshot technology are also now available to address the same issues. These solutions are slower than hardware based backup and restore solutions but are faster than conventional backup procedures. Some backup implementations combine both software and hardware methods to provide more versatile solutions.

1.2.2 Local Data Replication Options

Local data and replication involves creating and maintaining replicas of primary data volumes on local storage subsystems. There are primarily two techniques of local data replication:

- Mirroring or Full volume replication
- Differential data replication

Full Volume Replication: Full volume replication uses split mirror technology which mirrors or clones a set of data volumes or drives to an existing alternate mirrored set of drives, often using a SAN. The mirrored data volumes can be rapidly synchronized and split from the original data volumes after replication. This replication technique creates multiple full volume point-in-time copies (snapshots) of critical information, on local storage subsystems or secondary storage mediums.

Differential Data Replication: Differential data replication uses copy-on-write snapshot technology to capture point-in-time images (snapshots) of differential data in a data volume. Differential data can be recently updated data (in a primary data volume) resulting from host I/O operations or residual data (in a secondary data volume) resulting from a previous copy operation. These point-in-time images are temporarily stored in a group of one or more designated disk volumes (data pools), and made available to replication functions (during periodic copy cycles) or to restoration functions (during recovery processes). This replication technique allows for accurate restoration of original data volumes or faster takeover processing using the mirrored volumes.

Your choice of data replication method should depend on the varying data protection requirements of your different applications and your organization’s data retention needs. When evaluating and understanding data protection techniques and alternatives, it is necessary to take the following factors into consideration:

- Applicable threats to your applications and data
- RTO and RPO requirements associated with your applications

- Hardware and software technologies already in place
- Usability of data protection solution's interface
- Budget, timeframe, tolerance to disruption and risk aversion

1.2.3 Snapshots

In a storage environment, a snapshot is a point-in-time copy of data generated for backup or replication purposes. Snapshots are images representing the state of a storage volume at a particular point in time. Snapshots can be either partial copies of changed data (copy-on-write snapshots) or full volume copies (split-mirror snapshots, clones or mirrored volumes).

Copy-on-Write Snapshots: Copy-on-write snapshots are instantaneous point-in-time copies of data that has changed on the original storage volumes. Copy-on-write snapshots are virtual images (views) of the original data volume, generated by referencing both the original data volume and the changed data blocks in the data pool. Each virtual snapshot comprises of references to data that exists partly in the data pool and partly in the original/primary volume.

It is required to maintain multiple generations of snapshots to effectively restore or recover data. Generations of copy-on-write snapshots can be stored in designated physical disk volumes (pools). A significant savings of storage space is realized when using copy-on-write snapshots instead of full cloning methods, since the amount of storage space required for retaining each snapshot is substantially smaller than the storage space required for storing full volume copies. Since snapshots are available instantly they are useful for speedy data restoration or recovery after data loss due to file system corruption, application malfunction or erroneous write operations.

Split Mirror Snapshots: Split-mirror snapshots are volume captures of all data in the original storage volume. All data blocks in the original data volume are sequentially copied onto a secondary volume. The original and secondary data volumes are synchronized so that all updates to the original data volume are continually mirrored to the secondary data volume. Full volume copies can be split from the original data volumes. When they are split, they are essentially mirror images of each other at that point in time. 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, these replicas allow for more simplified data recovery and archival processes.

It generally takes longer to create full volume captures and each split mirror snapshot requires more storage space than copy-on-write snapshots. Since they are physical copies of data volumes, they require the same amount of storage space as the original data volume.

1.2.4 Copy-on-Write Technology

Copy-on-Write SnapShot technology uses software utilities to continually monitor write operations to a storage subsystem, and copy differential data (resulting from writes to the original volumes) onto a designated storage volume (data pool). Virtual snapshots (views) are captured at particular instants in time when a snapshot creation instruction is processed. These snapshots are maintained in virtual volumes (V-VOLs) and comprise of pointer references to physical data in the original volume and the data pool.

This technique derives its name (copy-on-write) from the fact that there is a copy operation associated with every write operation received by the storage medium. This copy operation is performed first, even before the original data volume is updated with new or changed data. When a host issues a write operation to the original data volume, the changed data block or new data block is first copied to the data pool. Following this copy operation, pointers associated with existing (previous) snapshots in the V-VOLs, are updated to reference the new or changed data block in the data pool. The write operation is then completed by updating or creating the data block on the original data volume.

The copy-on-write replication technique replaces older virtual snapshot images on the virtual volumes when new snapshots are generated, maintaining only predetermined generations of point-in-time snapshots. This technique involves processes that run in the background without demanding server resources and impacting application performance. Copy-on-write technology is usually used in combination with other hardware replication options to implement backup to secondary storage volumes.

Snapshot technologies can be used for the following purposes:

- To simplify and speed up restoration after data loss, data corruption or failure to backup critical data
- To make copies of data for test purposes including software development, regression testing etc
- To make copies of data for application processing such as data warehousing, reporting and data mining
- To make copies for non-disruptive backups and data migration

1.2.5 Integrating Local and Remote Replication

Local replication functions (such as snapshot or image backups) can be combined with remote data replication functionality to provide extremely resilient and powerful replication solutions. The ability to integrate these functions differentiates local data replication technology from traditional backups. The static point-in-time frozen replicas (snapshots) of data can be transferred to the secondary storage mediums during different types of remote backup operations. The multiple snapshots created prior to the onset of a disaster, are useful for data recovery, particularly from rolling disasters. Frozen virtual copies created at frequent checkpoints, allow fast data restoration and rapid recovery of critical data.

Note: Copy-on-Write Snapshot and ShadowImage software can also be active simultaneously in the same subsystem.

1.3 Hitachi's Local Replication Product Portfolio

Hitachi's replication solutions enhance operational efficiency and resilience of your information systems. Using the available solutions you can implement customized business continuity solutions, addressing key business problems relating to disaster recovery, data protection and availability. Different solutions are available for the large-size TagmaStore Universal Storage Platform (USP/NSC) subsystems and the mid-range TagmaStore AMS/WMS subsystems.

Small and medium-sized businesses, facing the same business continuity challenges as their enterprise-level counterparts, can keep data assets protected and available with the local replication solutions available from Hitachi Data Systems for mid-range TagmaStore AMS/WMS subsystems. The portfolio of midrange local replication solutions includes the Copy-on-Write SnapShot solution (hereafter referred to as SnapShot) and the ShadowImage In-system Replication solution, either of which can be leveraged depending on an organization's data protection objectives. The following sections discuss the ShadowImage local replication solution for mid-range TagmaStore AMS/WMS subsystems and provide an overview of the other Hitachi copy solutions.

1.3.1 ShadowImage In-System Replication Software

Hitachi's ShadowImage is an in-system replication software that uses a high-speed, non-disruptive, local mirroring technology to rapidly create multiple full-volume disk-based copies or clones of mission-critical information within storage subsystems, while keeping data RAID-protected and fully recoverable. ShadowImage features enable you to maintain subsystem-internal copies of all user data on the TagmaStore AMS subsystem for purposes such as data backup or duplication. The duplicated volumes are created within the same TagmaStore AMS subsystem as secondary volumes. ShadowImage is used for UNIX[®]-based and PC server data and can provide up to three duplicates (one duplicate: the micro program is less than 0750/A version) for each primary volume.

ShadowImage can be combined with TrueCopy to provide a versatile remote replication solution with maximum flexibility in data backup and restoration. When integrated with advanced disaster recovery techniques, it provides the ability to create high-availability solutions for quick restoration of production environments, in the event of a disaster.

1.3.2 Other Complementary Copy Solutions

The following section provides a brief overview of some of Hitachi's other local and remote data replication technologies.

1.3.2.1 Hitachi TrueCopy™ Software

Hitachi TrueCopy™ software is the industry's first storage hardware-based, server-less, asynchronous, remote copy solution that replicates data over any distance. TrueCopy remote data replication technology can be used in conjunction with local replication functions such as Copy-on-Write SnapShot software or ShadowImage software, to provide versatile data protection and disaster recovery solutions.

1.3.2.2 Hitachi TrueCopy Extended Distance Software (TCE)

Hitachi has integrated its TrueCopy remote replication technology with Copy-on-write Snapshot operations, to provide TrueCopy Extended Distance (TCE), a continuous, non-disruptive, host-independent asynchronous remote replication and disaster recovery solution for midrange Hitachi TagmaStore® Adaptable Modular Storage models AMS500 and AMS1000. TrueCopy Extended Distance enables open-system users to perform remote copy operations between TagmaStore storage sub-systems in geographically distant locations, using fibre-channel interface connections.

1.3.2.3 Hitachi TrueCopy Synchronous Software

For distances within the same metropolitan area, TrueCopy Synchronous Remote Replication software (TrueCopy) provides a no data loss, rapid restart solution. TrueCopy provides the ability to copy data over short and medium distances, by utilizing local clones and data movement capabilities of TrueCopy™ technology. Combined with Universal Replicator and the Universal Storage Platform, TrueCopy allows for advanced three data center configurations for optimal data protection.

1.3.2.4 Hitachi Copy-on-Write SnapShot Software

Hitachi's Copy-on-Write Snapshot software uses a high-speed, non-disruptive, logical snapshot technology to rapidly create up to 15 point-in-time snapshot copies of any data volume within AMS/WMS subsystems. It is a local replication solution that creates instantaneous virtual backup copies without impacting host service or performance levels. These snapshots significantly reduce backup time and accelerate recovery processes and are suitable for immediate use in decision support, software testing and development, data backup, or rapid recovery operations.

1.3.2.5 ShadowImage FlashCopy® Extension

Based on ShadowImage software, the FlashCopy® Extension provides rapid replication capabilities that allow customers to physically create copies of a source volume (S-VOL) to a target volume (T-VOL). When the ShadowImage FlashCopy function is invoked with IBM Peer-to-Peer Remote Copy TSO commands or DFSMSdss commands, a relationship is established between the S-VOL and T-VOL. Once this relationship is established, the T-VOL (physical copy of the S-VOL) is available for both read and write operations.

When establishing a relationship for a ShadowImage FlashCopy pair, customers can specify a range of data to be copied, which is called the "extent". When the extent data copy is complete, the relationship ends automatically. The Hitachi ShadowImage FlashCopy Extension is functionally compatible with IBM FlashCopy.

1.4 ShadowImage Product Differentiation

Hitachi's ShadowImage is a data protection solution for the following demand drivers and customer requirements:

- Meet objectives for availability and performance
- Provide reliable access to, and the sharing and safeguarding of critical business information
- Protect more data with less effort
- Execute disk-based backups while critical applications remain productive and online
- Provide current test data for application development without disrupting production
- Cost effectively maintain frequent point-in-time copies on disk for faster recovery
- Provide nondisruptive copies of online databases for decision support

The following sections provide competitive benefits of ShadowImage software and differentiate ShadowImage from Hitachi's Copy-on-Write SnapShot software.

1.4.1 Product Market Differentiation

Hitachi's ShadowImage offers the following competitive business continuity benefits as well as productivity and process improvements:

Ensures Business Continuity and Data Availability

- Shortens restart and recovery times with the consistencygroup function, which provides multivolume, point-in-time copies for applications and databases that share or span multiple volumes
- Reduces recovery from data corruption time dramatically through the ShadowImage QuickRestore feature
- Replicates large data volumes without impacting service levels, timing out, or affecting performance levels
- Provides non-disruptive, volume "snapshots" which use less space than full copies or "clones"
- Allows frequent, cost-effective, point-in-time copies (up to 15 virtual copies of a primary volume)
- Shortens restart and recovery times with the consistency group function, which provides multivolume, point-in-time copies for applications and databases that share or span multiple volumes
- Dramatically reduces recovery from data corruption or human error through immediate restores from a disk-resident, point-in-time data copy
- Eliminates the backup window and enables immediate read/write access to time-critical information using virtual copies

Improves Productivity and Processes

- Reduces testing and deployment time and increases the accuracy of application development by providing always available copies of current production data

- Increases competitive advantage through the ShadowImage QuickSplit function by facilitating the sharing of and immediate access to time-critical information
- Enables normal backup operations on a copy of up-to-date production data while critical applications run unaffected
- Improves operational efficiency by allowing multiple processes to run in parallel with access to the same information
- Improves backup efficiency and eliminates costly downtime due to tape backups
- Offers shorter development cycles which translates into faster time-to-market for new business processes
- Accelerates application testing and development, and provides fast access to data for decision support

Reduces Operational and Capital Costs

- Allows business to remain online during data center activities, eliminating the need for 24/7 resources to perform these tasks
- Enables IT to refocus resources onto revenue-driven projects

1.4.2 Differentiation from Copy-on-Write SnapShot

Hitachi ShadowImage™ copies differ from Copy-on-Write snapshots in that they are physically separate copies of the data and no portion is shared between primary and secondary volumes, therefore performance and protection issues are resolved.

Table 1.2 SnapShot-ShadowImage Differentiators

	ShadowImage Clone	Copy-on-Write Virtual Volume
Features	Superior protection because it is a complete copy Immediate restore if P-VOL becomes corrupted	Provides a very quick copy because it is a virtual volume
Time to Create	From minutes to hours, depending on the size of the P-VOL	Instantaneous
Disk Space Used	Same amount as P-VOL	Size will vary depending upon rate of data change but will be far less than the P-VOL
Data Recovery Time After P-VOL is corrupted	Instantaneous	Two stage restore: 1. V-VOL verify, data not available to host 2. Copy back, data available

	ShadowImage	Snapshot
Size of Physical Volume	<p>P-Vol = S-Vol</p>	<p>P-Vol \geq Pool for one V-Vol</p>
Pair Configuration	<p>1:1</p>	<p>1:14</p>
Restore	<p>P-Vol can be restored from S-Vol</p>	<p>Restore from any V-Vol</p>

Figure 1.2 SnapShot ShadowImage Product Differentiators

For more detailed product specifications and function differentiation, see 1.4.2 SnapShot and ShadowImage Differentiation.

Note on product availability: Copy-on-Write SnapShot and ShadowImage solutions can be purchased separately or as a competitively priced in-system replication bundle.

Chapter 2 Overview of ShadowImage In-system Replication Software

This chapter presents an overview of ShadowImage software and discusses the functional and operational details of the product. This chapter includes the following sections:

- Typical ShadowImage System Environment (see section 2.1)
- ShadowImage Functional Components (see section 2.2)

2.1 Typical ShadowImage System Environment

A typical ShadowImage hardware configuration includes a TagmaStore AMS subsystem (local subsystem or remote subsystem), a host connected to the TagmaStore subsystem and a management server. The host is connected to the TagmaStore AMS subsystem via fibre channel connections. The management server is connected to the TagmaStore AMS subsystems via a management LAN.

The logical configuration of the TagmaStore AMS subsystem includes a command device, a differential management logical unit (DM-LU), primary data volumes (P-VOLs) belonging to the same consistency group (CTG), and a logical unit for the data pool.

In addition to the above components, the ShadowImage system architecture also includes the Storage Navigator Modular (SNM) software and the Command Control Interface (CCI) software. Storage Navigator Modular is installed on the management server and is used to configure the ShadowImage system environment. Command Control Interface software is installed on the host and manages ShadowImage volume pair operations.

The following illustration shows a typical ShadowImage system configuration.

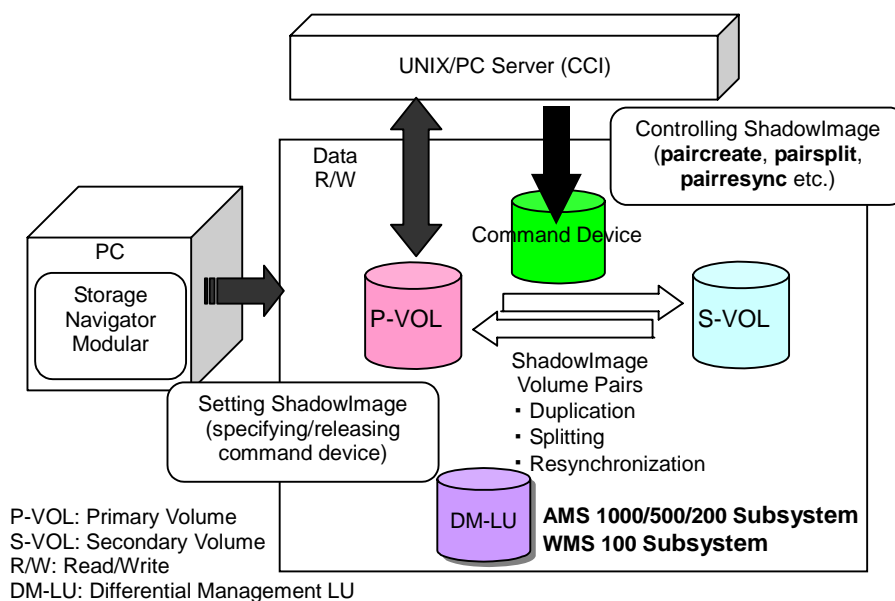


Figure 2.1 ShadowImage Environment

2.2 ShadowImage Functional Components

ShadowImage functional components include:

- ShadowImage volume pairs (P-VOLs and S-VOLs) (see section 2.2.1)
- Consistency Group (CTG) (see section 2.2.2)

2.2.1 ShadowImage Volume Pairs (P-VOLs and S-VOLs)

The TagmaStore subsystem contains and manages both the original and copied ShadowImage data. ShadowImage supports a maximum of 255 pairs (WMS100/AMS200), 1,023 pairs (AMS500), or 2,047 pairs (AMS1000).

ShadowImage performs internal copy operations for logical volume pairs established by the user. Each ShadowImage pair consists of one primary volume (P-VOL) and up to three (one: micro program less than 0750/A version) secondary volumes (S-VOLs), these are located in the same TagmaStore subsystem. The ShadowImage P-VOLs are the primary volumes, which contain the original data. The ShadowImage S-VOLs are the secondary or mirrored volumes, which contain the duplicated data. During ShadowImage operations, the P-VOLs remain available to all hosts for read and write I/O operations (except during reverse resync). The S-VOLs become available for host access only after the pair has been suspended or released.

When a ShadowImage volume pair is created, the data on the P-VOL is copied to the S-VOL to synchronize the volumes. During this initial copy operation and after the pair is synchronized, all write operations to the S-VOL are prohibited. If you need to access an S-VOL, you can “suspend” the pair to make the S-VOL accessible (the P-VOL is accessible except during reverse resync). While a ShadowImage pair is suspended, the TagmaStore subsystem keeps track of all changes to the P-VOL and S-VOL. When you “resync” the pair, the differential data in the P-VOL (due to P-VOL updates) is copied to the S-VOL so that the S-VOL is again identical to the P-VOL. For reverse resync, the differential data in the S-VOL (due to S-VOL updates) is copied to the P-VOL so that the P-VOL is now identical to the S-VOL.

Note: For performance reasons, ShadowImage P-VOLs and S-VOL pairs should not be from the same RAID group. ShadowImage pairs for which ShadowImage operations are performed simultaneously should be on different controllers, but the P-VOL and S-VOL of each pair must be on the same controller. If ShadowImage pairs are concentrated in the same controller, the host I/O performance may be impacted.

The TagmaStore is capable of performing only one `paircreate` command at a time per controller.

2.2.2 Consistency Group (CTG)

A file system can store the application data or a logical volume on an OS that is configured with two or more logical units. In this case, you must assure that the data of those logical units are of the same time. In order to assure that the data of two or more S-VOLs are of the same time and to manage them as a group, the TagmaStore provides the consistency group (CTG). The TagmaStore assures that data of the S-VOLs in the group are of the same time by splitting the pairs in the group using the CTG. When the CTG is not used, the timing to split each S-VOL varies because the pairs in the group are split in order. Therefore, if an I/O instruction is received from a host in the middle of the splitting of the pairs in the group, data of the different time is stored in each S-VOL.

To use the CTG when splitting a pair, you must create the pair and add the `-m grp` option at the time of the pair creation.

Chapter 3 ShadowImage Requirements

This chapter describes ShadowImage operational requirements and provides an overview of ShadowImage management software. This chapter includes the following sections:

- ShadowImage Specifications (see section 3.1)
- ShadowImage Requirements (see section 3.2)
- ShadowImage Management Software (see section 3.3)

3.1 ShadowImage Specifications

This section describes the external specifications for ShadowImage.

Table 3.1 External Specifications

No.	Item	TagmaStore
1	Applicable model	WMS100/AMS200/AMS500/AMS1000 (For dual configuration only).
2	Number of pairs	AMS1000: 2,047 (maximum) AMS500: 1,023 (maximum) AMS200: 255 (maximum) WMS100: 255 (maximum) When one P-VOL configures a pair with three S-VOLs, the number of pairs becomes three.
3	Number of command devices	Two LUs
4	Pair structure (number of S-VOLs per P-VOL)	1:3 (1:1 micro program less than 0750/A version)
5	P-VOL and S-VOL arrangement (SMPL)	Within the same CTL
6	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.
7	RAID configuration for P-VOL and S-VOL	P-VOL and S-VOL can be paired between different RAIDs. The number of data disks do not need to be the same.
8	Types of the drive for a P-VOL/S-VOL	Both LUs consisting of FC drives and SATA drives can be assigned to any P-VOL and S-VOL. However, assign an LU with FC drives to a P-VOL. When integrating the LUs configured in the SATA drive, the use conditions of the SATA drive may be different (see Table 3.3).
9	Consistency Group (CTG) number	Max 128/subsystem WMS100/AMS200: 255 AMS500: 1,023 AMS1000: 2,047
10	Mixing ShadowImage and non-ShadowImage	Available (mixing within CTL, and between CTLs).
11	Host interface	Fibre, iSCSI, or NAS
12	Size of P-VOL and S-VOL	P-VOL = S-VOL
13	Deleting during Coupling (RAID group, P-VOL, S-VOL) Formatting during Coupling	Not available. However, when the pair status is PSUE (S-VOL Switch) only a P-VOL can be formatted.
14	Concurrent use of TrueCopy	TrueCopy can be cascaded with ShadowImage. For details, see section 7.5.
15	Concurrent use of SnapShot	SnapShot and ShadowImage can be used together at the same time, but cascade between SnapShot and ShadowImage is not supported. The number of CTG at the time of using SnapShot and ShadowImage together is limited to the maximum of 128 combining that of SnapShot and ShadowImage.
16	Concurrent use of TCE	TCE and ShadowImage can be used together at the same time, but cascade between TCE and ShadowImage is not supported.

17	Concurrent use of LUN Expansion	Available (For details, see section 7.4.)
18	Concurrent use of Volume Migration	Available. However, a P-VOL, an S-VOL, and reserved LU of Volume Migration cannot be specified as a P-VOL or an S-VOL of ShadowImage. Besides, 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. For details, see section 4.4.
19	Concurrent use of Cache Residency Manager	The LU specified for Cache Residency (LU cache residence) cannot be set as an S-VOL or a P-VOL.
20	Concurrent use of SNMP Agent	Available
21	Concurrent use of LUN Manager	Available
22	Concurrent use of Password Protection	Available
23	Concurrent use of Data Retention	Available
24	Concurrent use of Power Saving	Available. However, when a P-VOL or an S-VOL is included in a RAID group, for which the Power Saving has been specified, no pair operation can be performed except the pair split and the pair release.
25	ShadowImage I/O switching function	Available. For details, see section 4.6.
26	Maximum supported capacity value of S-VOL (TB)	Maximum supported capacity value of S-VOL (TB) (Maximum ShadowImage single capacity - (Total P-VOL and S-VOL capacity of TrueCopy) - 17 x (Total P-VOL capacity of SnapShot) - 51 x (Total P-VOL and S-VOL capacity of TCE)) ÷ 2
27	License	ShadowImage becomes usable through entry of the key code.
28	Size for management of difference data	Fixed at 1 MB.
29	Location for storing a bit map	1 GB area at the top of the Differential Management LU.
30	Number of Differential Management LUs	One or two (One of them is for mirroring) The Differential Management LU size must be greater than or equal to 5 GB. It is recommended that two Differential Management LUs are set according to following conditions <ul style="list-style-type: none"> ▪ To be created in different RAID group ▪ To be allocated in different controllers.
31	Point in time Split	All pairs can be defined as one group.

3.2 ShadowImage Requirements

ShadowImage operations provide subsystem-internal copies of UNIX/PC server volumes (LUs) on the TagmaStore subsystem. Table 3.2 lists and describes the operational requirements for ShadowImage.

Table 3.2 ShadowImage Requirements

Parameters	Requirements
User interface	The following software is necessary: <ul style="list-style-type: none"> ▪ Hitachi Command Control Interface (CCI): Used for issuing commands to create and split ShadowImage pairs. ▪ Navigator: Used for specifying and releasing command devices.
Number of controllers	2
Max number of command devices	2
Max number of Differential Management LUs	2 The Differential Management LU size must be greater than or equal to 5 GB. It is recommended that two Differential Management LUs are set according to following conditions. <ul style="list-style-type: none"> ▪ To be created in different RAID group ▪ To be allocated in different controllers
Unit of pair management	LUs are the target of ShadowImage pairs, and are managed per logical unit.
Max number of pairs	WMS100/AMS200: 255 AMS500: 1,023 AMS1000: 2,047 When one P-VOL configures a pair with three S-VOLs, the number of pairs becomes three.
Pair structure	Micro program version 0750/A or later: Up to three (S-VOLs) per P-VOL. Micro program less than 0750/A version: One copy (S-VOL) per P-VOL. When one P-VOL configures a pair with the maximum of three S-VOLs, only a set of ShadowImage can be in the PAIR status, the COPY status or the PSUE(R) status.
Supported RAID level	RAID 0 (2D to 16D), RAID 1 (1D+1D), RAID 5 (2D+1P to 15D+1P), RAID 6 (2D+2P to 28D+2P), RAID 1+0 (2D+2D to 8D+8D) (Using the P-VOL and S-VOL both which have a redundant RAID level is recommended. However, RAID 0 cannot be set for the SATA disk drive.)
Combination of RAID levels	All combinations supported. The number of data disks does not need to be the same.
Size of LU	The P-VOL size must equal the S-VOL LU size in blocks. The Command Device LU size must be greater than or equal to 33 MB.
Types of the drive for a P-VOL/S-VOL	Both LUs with FC drives and SATA drives can be assigned to any P-VOL and/or S-VOL. However, it is recommended to assign an LU consisting of FC drives to a P-VOL. When integrating the LUs configured in the SATA drive, the use conditions of the SATA drive may be different (see Table 3.3).
Allocation of P-VOL and S-VOL	The P-VOL and S-VOL can only be allocated in the same controller.
Consistency Group (CTG) number	Max 128/subsystem WMS100/AMS200: 255 AMS500: 1,023 AMS1000: 2,047
Host interface	Fibre, iSCSI, or NAS.

Initial copy	When receiving the paircreate command, the entire copy is executed in the direction of P-VOL to S-VOL.
Split	When receiving the pairsplit command, splitting P-VOL and S-VOL is executed. Splitting/suspending is completed immediately (within 10ms).
Differential management	When the P-VOL and S-VOL is in suspend status (PSUS), write operations received individually will be managed as the differential data of the P-VOL and the S-VOL. When one P-VOL configures a pair with the maximum of three S-VOLs, the difference is managed for each pair.
Resynchronization	When receiving the pairresync command, the differential copy is executed in the direction of P-VOL to S-VOL or S-VOL to P-VOL.
S-VOL data creation	The S-VOL data in cache is created during the write operation. <ol style="list-style-type: none"> 1. Receiving the write command from the host. 2. Creates the P-VOL data. 3. Reports "GOOD" to host. 4. Creates the S-VOL data.
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.
Concurrent use of LUN Expansion	Available. For details, see section 4.3.
Concurrent use of Data Retention Utility	When the S-VOL Disable is set for an LU, a pair creation using the LU as an S-VOL is suppressed. An S-VOL Disable setting of a volume that has already become an S-VOL is not suppressed only when the pair status is PSUS.
Concurrent use of Cache Residency Manager	The LU specified for Cache Residency (LU cache residence) cannot be set as a S-VOL or P-VOL.
Concurrent use of SNMP Agent	Available
Concurrent use of Password Protection	Available
Concurrent use of LUN Manager	Available
Concurrent use of TrueCopy	TrueCopy can be cascaded with ShadowImage. For details, see section 4.5.
Concurrent use of TCE	TCE and ShadowImage can be used together at the same time, but a cascade between TCE and ShadowImage is not supported.
Concurrent use of SnapShot	SnapShot and ShadowImage can both run concurrently on the same array with the following restrictions: <ul style="list-style-type: none"> ▪ You cannot take a SnapShot of a ShadowImage volume (P-VOL or S-VOL). ▪ You cannot make a SnapShot P-VOL or V-VOL a member of a ShadowImage pair. The number of CTG at the time of using SnapShot and ShadowImage together is limited to the maximum of 128 combining that of SnapShot and ShadowImage.
Concurrent use of Cache Partition Manager	Available
Concurrent use of Volume Migration	Available. However, a P-VOL, an S-VOL, and reserved LU of Volume Migration cannot be specified as a P-VOL or an S-VOL of ShadowImage. Besides, 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. For details, see section 4.4.
Concurrent use of Power Saving	Available. However, when a P-VOL or an S-VOL is included in a RAID group, for which the Power Saving has been specified, no pair operation can be performed except the pair split and the pair release.

Failures	When a failure of the copy operation from P-VOL to S-VOL occurs, ShadowImage will suspend the pair (PSUE). If a LU failure occurs, ShadowImage suspends the pair. If a drive failure occurs, the ShadowImage pair status is not affected because of the RAID architecture.
Parameters	Requirements
ShadowImage I/O switching function	This function can switch the object of a host I/O instruction to an S-VOL when a double failure (triple failures for RAID 6) occurs in the P-VOL drives. For details, see section 4.6.
Deleting a RAID group, deleting an LU, and formatting a paired P-VOL or S-VOL	You cannot delete a RAID group, delete an LU, and format a paired P-VOL or S-VOL. However, when the pair status is PSUE (S-VOL Switch), a P-VOL can be formatted.
Restriction for formatting the volumes	Do not execute ShadowImage operations while formatting the volume. Formatting takes priority and the ShadowImage operations will be suspended. The pair operation of ShadowImage can be performed unless the volume to be targeted is being formatted.
Reduction of memory	Memory cannot be reduced when the ShadowImage, SnapShot, TrueCopy, or TCE function is validated. Make the reduction after invalidating the function.

Table 3.3 Characteristics of FC Drive and SATA Drive

Items	SATA Drive		FC Drive
	ATE250R/ATE400R	Others	
Drives presumed operation environment	The Read/Write access takes 330 hours/month or less on the average.	The Read/Write access takes 720 hours/month.	The Read/Write access takes 720 hours/month.

3.3 ShadowImage Management Software

3.3.1 Command Control Interface Software (CCI)

CCI software is used to display ShadowImage volume information, create and manage ShadowImage pairs, and issue commands for replication operations. CCI software resides on the UNIX®/ MS Windows® management server and interfaces with the TagmaStore AMS subsystems through dedicated logical volumes. CCI commands can be issued from the UNIX command line or using a script file. For more information on CCI, see *Hitachi TagmaStore™ Adaptable Modular Storage and Workgroup Modular Storage Command Control Interface (CCI) User and Reference Guide* (MK-95DF701).

3.3.2 Storage Navigator Modular Software (SNM)

Storage Navigator Modular software GUI displays detailed ShadowImage information and is used for configuring the ShadowImage environment. SNM software communicates directly with the TagmaStore AMS subsystems via a local area network (LAN). SNM software supports both a GUI and CLI user interface to perform configuration, monitoring, maintenance and system reduction.

ShadowImage configuration tasks include setting up data pool, cycle time, logical units and path for remote ShadowImage operations, enabling or releasing a command device etc. SNM interfaces display important failure information for subsystems and maintenance information including the amount of differential data, the time difference between the P-VOL and the S-VOL, and the used capacity of a data pool. Navigator can be used to increase the pool capacity by adding more LUs to the pool.

In the event of a system failure or disaster at the primary site, the SNM software simplifies and expedites disaster recovery procedures. System Reduction tasks such as deleting command devices and data pools can also be performed using SNM software.

For more information on SNM software, see *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Storage Navigator Modular Graphical User Interface (GUI) User's Guide* (MK-95DF711) or *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Storage Navigator Modular Command Line Interface (CLI) User's Guide* (MK-95DF712).

3.3.3 Optional Software

The following optional management software is also used in the ShadowImage system environment:

- HiCommand Device Manager (see section 3.3.3.1)
- HiCommand Tuning Manager (see section 3.3.3.2)

3.3.3.1 HiCommand Device Manager

HiCommand Device Manager is a storage management product which enables users to consolidate storage operations and manage capacity in systems containing multiple storage subsystems. Targeted for users managing multiple storage subsystems in open or shared environments, Device Manager quickly discovers the key configuration attributes of storage subsystems and allows storage administrators to begin proactively managing complex and heterogeneous storage environments using its easy-to-use, browser-based GUI. Device Manager's built-in report facility compiles and presents key information in preformatted reports (HTML) and as comma-separated values for export.

Device Manager provides link-and-launch integration with other software such as HiCommand Tuning Manager, HiCommand Protection Manager and Storage Navigator Modular (for Web) software. For more information on HiCommand Device Manager, see *Hitachi HiCommand Device Manager Web Client User's Guide (MK-91HC001)* or *Hitachi HiCommand Device Manager Command Line Interface (CLI) User's Guide (MK-91HC007)*.

3.3.3.2 HiCommand Tuning Manager

HiCommand Tuning Manager is a storage resource management product which enables centralized management of a storage area network environment. This software monitors and analyzes the performance information and capacity information of the host, fibre channel switches, and storage devices. Changes in host I/O performance and capacity over time are displayed graphically and useful for planning future system expansion and reduction. For more information on HiCommand Tuning Manager, see *Hitachi HiCommand Tuning Manager User's Guide (MK-92HC022)*.

Chapter 4 Overview of ShadowImage Operations

ShadowImage is a storage-based hardware solution for duplicating logical volumes that reduce backup time and provide point-in-time backup. The ShadowImage primary volumes (P-VOLs) contain the original data; the ShadowImage secondary volume(s) (S-VOLs) contain(s) the duplicate data. Since each P VOL is paired with its S-VOL independently, each volume can be maintained as an independent copy set that can be split (**pairsplit**), resynchronized (**pairresync**), and released (**pairsplit -S**) separately.

ShadowImage operations are performed from the UNIX and/or PC-server host using the Hitachi Data Systems Command Control Interface (CCI) software. ShadowImage commands are issued from CCI through command devices in the TagmaStore subsystem. These command devices can be specified or released by the user by using the Navigator.

This chapter includes the following:

- ShadowImage Operations (see section 4.1)
- ShadowImage Pair Status (see section 4.2)
- Concurrent Use of LUN Size Expansion (see section 4.3)
- Concurrent Use of Volume Migration (see section 4.4)
- Cascade Connection of TrueCopy with ShadowImage (see section 4.5)
- ShadowImage I/O Switching Function (see section 4.6)

4.1 ShadowImage Operations

ShadowImage operations can be performed from the UNIX/PC server host using Hitachi Data Systems Command Control Interface (CCI) software. For further information on using the Hitachi Data Systems CCI, refer to the *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Command Control Interface (CCI) User and Reference Guide (MK-95DF701)*.

4.1.1 Paircreate Operation

The ShadowImage **paircreate** operation establishes the newly specified ShadowImage pair(s). The volumes, which will become the P-VOL and S-VOL, must both be in the SMPL (simplex) state before becoming a ShadowImage pair. Also, when the P-VOL to be an operation target configures a ShadowImage pair in the PAIR status, the COPY status or the PSUE(R) status with another S-VOL, a pair cannot be newly created using the **paircreate** command.

4.1.1.1 Initial Copy Operation

The ShadowImage initial copy operation takes place when you create a new ShadowImage pair (see Figure 4.1). 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 COPY(PD) (PD = pending) while the initial copy operation is in progress. The pair status changes to PAIR 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 pace is divided on the scale of 1 to 15. 1 to 5: slow, 6 to 10: medium, 11 to 15: fast. The default value is 3.

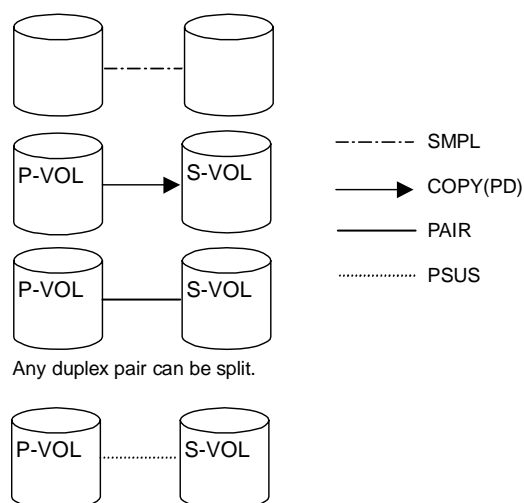


Figure 4.1 Adding a ShadowImage Pair

For details on **paircreate** operation, see section 6.1.

4.1.2 Pairsplit Operation

The ShadowImage **pairsplit** operation performs all pending S-VOL updates (those issued prior to the `split` command and recorded in the P-VOL track map) to make the S-VOL identical to the state of the P-VOL when the suspend command was issued, and then provides full read/write access to the split S-VOL. You can split existing pairs as needed, and you can also use the **pairsplit** operation to create and split pairs in one step. 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 PSUS, and you have full read/write access to the split S-VOL. While the pair is split, the TagmaStore 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 the **pairsplit** operation. **Pairsplit** operations cannot be performed on suspended (PSUE) pairs.

For details on pairsplit operation, see section 4.2.

4.1.3 Pairresync Operations (Normal, Reverse)

The ShadowImage **pairresync** operation resynchronizes the suspended pairs (PSUS) or the suspended on error pairs (PSUE). When the **pairresync** operation starts, the pair status changes to COPY(RS) or COPY(RS-R). The pair status changes to PAIR when the **pairresync** operation completes.

ShadowImage allows you to perform two types of **pairresync** operations:

- **Normal.** The normal **pairresync** operation (see Figure 4.2) resynchronizes the S-VOL with the P-VOL. The copy direction for a normal **pairresync** operation is P-VOL to S-VOL. The pair status during a normal resync operation is COPY(RS). The S-VOL becomes inaccessible to all hosts for write operations, and the P-VOL remains accessible to all hosts for both read and write operations during a normal **pairresync**. The normal **pairresync** operation can be executed for pairs with the status PSUS and PSUE. However, when the P-VOL that will be an operation target configures a ShadowImage pair in the PAIR status and the COPY status or the PSUE(R) status with another S-VOL, the **pairresync** operation cannot be executed. When the **pairresync** operation is executed for a pair in the PSUE (S-VOL Switch) status, the reverse resynchronization operation is performed.
- **Reverse.** The reverse **pairresync** (**pairresync -restore**) operation (see Figure 4.3) 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 PAIR status and the COPY status or the PSUE(R) status with another S-VOL, the **pairresync -restore** operation cannot be executed. The copy direction for a reverse **pairresync** operation is S-VOL to P-VOL. The pair status during a reverse resync operation is COPY(RS-R), and the S-VOL becomes inaccessible to all hosts for write operations during a reverse **pairresync** operation. The P-VOL remains accessible for both read and write operations, and the write operations on P-VOL will always be reflected to S-VOL (see Figure 4.4).




Figure 4.2 Normal Pairresync Operation




Figure 4.3 Reverse Pairresync Operation (pairresync-restore)

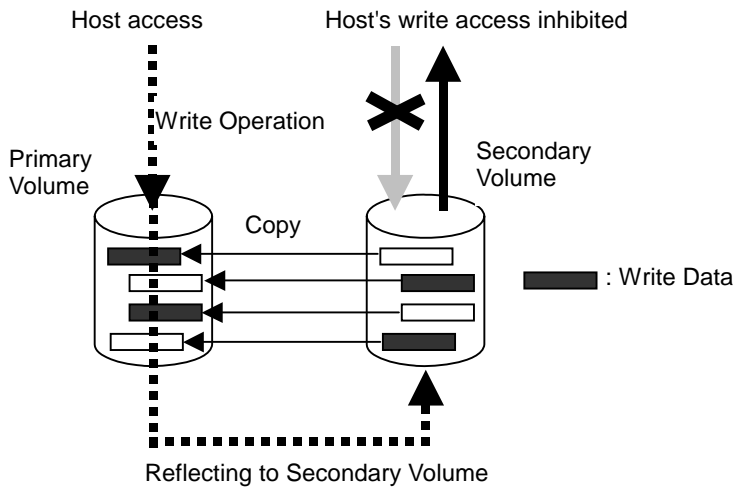


Figure 4.4 Reflecting Write Data to S-VOL during Reverse Pairresync Operation (pairresync-restore)

ShadowImage allows you to perform normal and quick **pairresync** operations on split and suspended pairs:

- **Pairresync for a suspended pair.** When a **pairresync** operation is performed on a suspended pair (status = PSUS), the subsystem merges the S-VOL differential track map into the P-VOL differential track map and then copies all flagged data from the P-VOL to the S-VOL. When a reverse **pairresync** operation is performed on a suspended pair, the subsystem 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.
- **Pairresync for suspended pair.** When a **pairresync** operation is performed on a suspended pair (status = PSUE), the subsystem 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. The **pairresync** operation takes as long as the ShadowImage initial copy operation.

For details on **pairresync** operation, see section 6.3.

4.1.4 Pairsplit -E Operation

The ShadowImage **pairsplit-E** operation 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 = PSUE), the subsystem 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 **pairresync** operation is performed on a suspended pair, the entire P-VOL is copied to the S-VOL (when a **pairresync -restore** operation is performed, the entire S-VOL is copied to the P-VOL). While the **pairresync** operation for a split ShadowImage pair can be very fast, the **pairresync** operation for a suspended on error pair will take as long as the initial copy operation.

The TagmaStore subsystem will automatically suspend a ShadowImage pair when it cannot keep the pair mirrored for any reason. When the TagmaStore suspends a pair, a file is output to the system log or event log to notify the host. The TagmaStore will automatically suspend a pair under the following conditions:

- When the ShadowImage volume pair has been suspended or deleted from the UNIX/PC server host using the Hitachi Data Systems CCI.
- When the TagmaStore detects an error condition related to an initial copy operation. When a volume pair with COPY(PD) status is suspended on error, the TagmaStore aborts the initial copy operation, changes the status of the P-VOL and S-VOL to PSUE and accepts all subsequent write I/Os to the P-VOL.

For details on **pairsplit -E** operation, see section 6.4.

4.1.5 Pairsplit -S Operation

The ShadowImage **pairsplit-S** 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 SMPL status. In both ShadowImage volumes, the status will change to SMPL.

For details on **pairsplit -S** operation, see section 6.5.

4.2 ShadowImage Pair Status

The ShadowImage displays the pair status of all ShadowImage volumes (LUs). Figure 4.5 shows the ShadowImage pair status transitions and the relationship between the pair status and the ShadowImage operations.

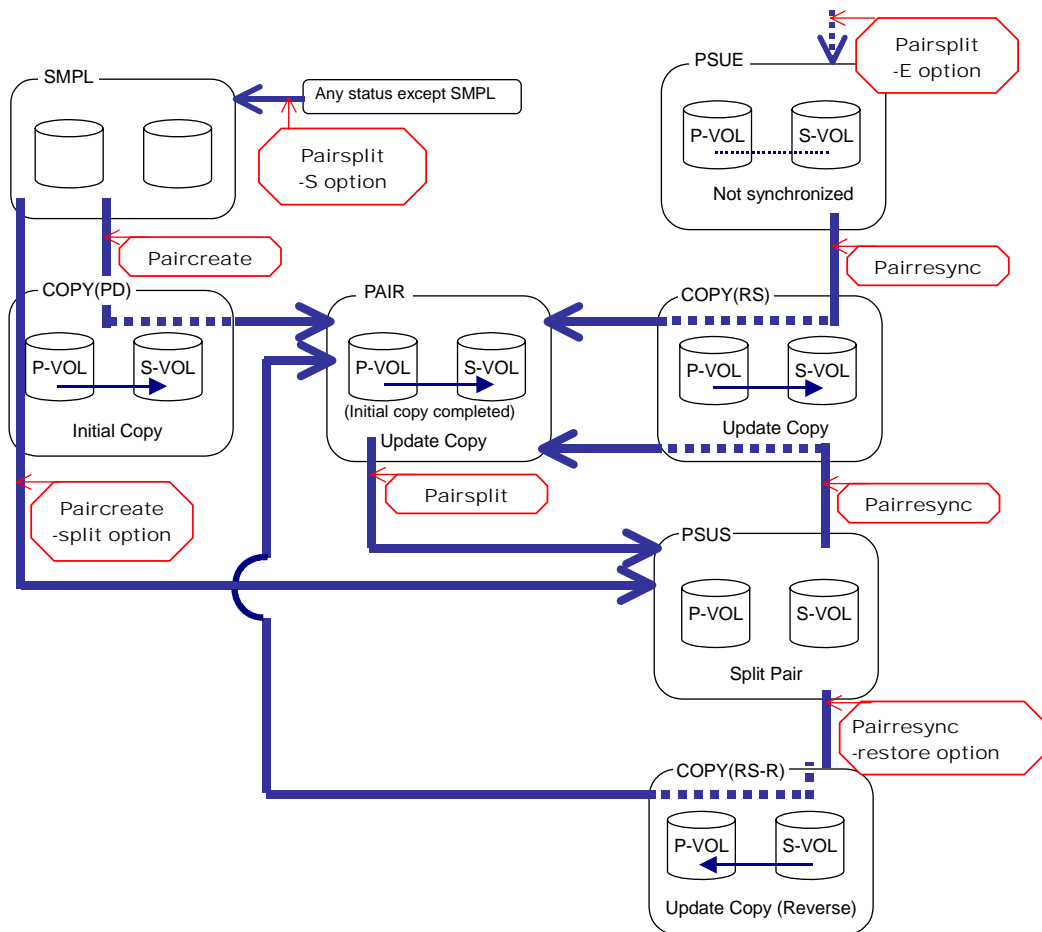


Figure 4.5 ShadowImage Pair Status Transitions

Table 4.1 lists and describes the ShadowImage pair status conditions. If a volume is not assigned to a ShadowImage pair, its status is SMPL. When you create a ShadowImage pair, the status of the P-VOL and S-VOL changes to COPY(PD). When the initial copy operation is complete, the pair status becomes PAIR. If the TagmaStore cannot maintain PAIR status for any reason or if you suspend on error the pair (**pairsplit -E**), the pair status changes to PSUE. When you suspend a pair (**pairsplit**), the pair status changes to COPY(SP). When the **pairsplit** operation is complete, the pair status changes to PSUS to enable you to access the suspended S-VOL. When you start a **pairresync** operation, the pair status changes to COPY(RS). When you specify reverse mode for a **pairresync** operation (**pairresync -restore**), the pair status changes to COPY(RS-R) (data is copied in the reverse direction from the S-VOL to the P-VOL). When the **pairresync** operation is complete, the pair status changes to PAIR. When you split/release a pair (**pairsplit -S**), the pair status changes to SMPL.

Table 4.1 ShadowImage Pair Status

Pair Status	Description	S-VOL Access
SMPL	The volume is not assigned to a ShadowImage pair. The TagmaStore accepts read and write I/Os for all SMPL volumes.	N/A (there is no S-VOL yet)
COPY(PD)	The initial copy operation is in progress. The TagmaStore continues to accept read and write operations for the P-VOL but does not accept write operations for the S-VOL.	Read only
PAIR	The initial copy operation is complete, and the TagmaStore 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 (PAIR status) is identical. The TagmaStore rejects all write I/Os for S-VOLs with the status PAIR.	Read only
PSUS	The TagmaStore starts accepting write I/Os for PSUS S-VOLs. The TagmaStore keeps track of all updates to the suspended P-VOL and S-VOL, so that the pair can be resynchronized quickly.	Read and write The S-VOL can be mounted.
COPY(RS)	The TagmaStore does not accept write I/Os for COPY(RS) S-VOLs. When a suspended pair is resynchronized in normal mode, the TagmaStore copies only the P-VOL differential data to the S-VOL. When a PSUE (suspended under error) pair is resynchronized, the TagmaStore copies the entire P-VOL to the S-VOL.	Read only
COPY(RS-R)	The TagmaStore does not accept write I/Os for COPY(RS-R) S-VOLs. When a split pair is resynchronized in reverse mode, the TagmaStore copies only the S-VOL differential data to the P-VOL.	Read only
PSUE	The TagmaStore continues accepting read and write I/Os for a PSUE (suspended under error) P-VOL (however, if the status transits from COPY(RS-R), all access to P-VOL is disabled). The TagmaStore marks the entire P-VOL track map as difference data, so that the entire P-VOL is copied to the S-VOL when the PSUE pair is resumed. Use the pairresync command to resume a PSUE pair.	Read only (If the status transits from COPY(RS-R), all access to P-VOL is disabled.)
(S-VOL Switch)	This is a state in which a double failure (triple failures for RAID 6) of drives occurred in a P-VOL and the P-VOL was switched to an S-VOL internally. This state is displayed as PSUE with CCI. For details, see section 4.6.	Read/Write is not available.

For details on **pairedisplay** operation, see section 6.6.

4.3 Concurrent Use of LUN Size Expansion

In ShadowImage, a unified LU can be used as a P-VOL or an S-VOL. Preconditions and restrictions of concurrent use with ShadowImage and LUN Size Expansion are as follows:

- P-VOL and S-VOL Capacities:
 - Capacities of the P-VOL and S-VOL must be the same. The P-VOL and S-VOL can be paired as long as their capacities are the same, even if numbers of LUs composing them (unified LUs) are different from each other.

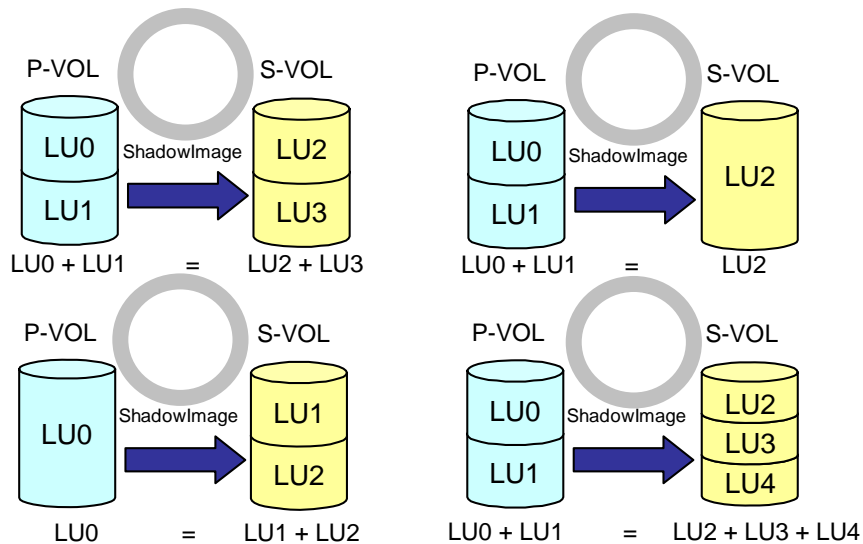


Figure 4.6 Example: A Unified LU Assigned to P-VOL or S-VOL




Figure 4.7 Example: Restriction Placed on Composing a Unified LU

- P-VOL and S-VOL RAID Levels:
 - All the LUs including the unified LU(s) assigned to the P-VOL and S-VOL must not be on the same RAID level and have the same number of data disks.

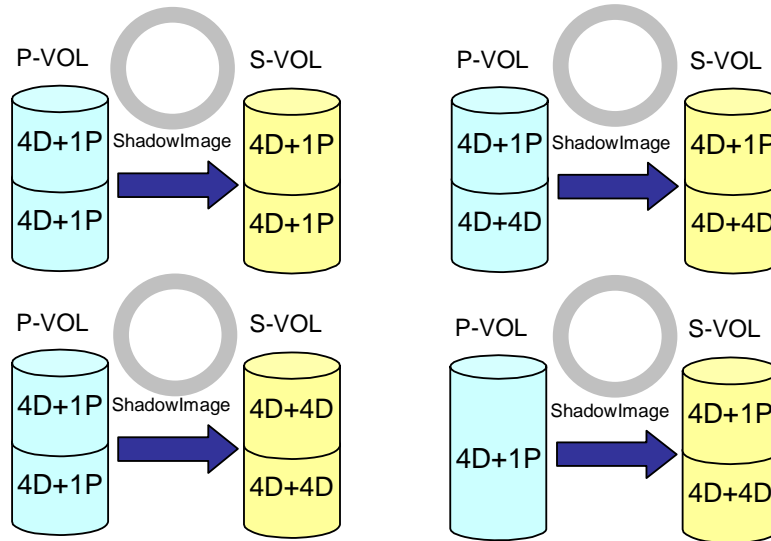


Figure 4.8 Example: A Combination of RAID Levels




Figure 4.9 Example: A Combination of RAID Levels

- Others:
 - A P-VOL or an S-VOL belonging to a ShadowImage pair cannot be used to compose a unified LU. Unify it after splitting the pair once.
 - An LU consisting of an FC drive with a SATA drive cannot be unified. Unify LUs consisting of drives of the same type.




Figure 4.10 Example: Restriction Placed on Unification of LUs Consisting of Different Drive Types

4.4 Concurrent Use of Volume Migration

When ShadowImage is used together with Volume Migration, the maximum number of pairs and the copying operations are limited. The notes when ShadowImage is used together with Volume Migration are shown here.

- Maximum number of pairs

The disk subsystem limits the maximum number of ShadowImage pairs and Volume Migration pairs to 255 (AMS200S/WMS100, 1,023 (AMS500), 2,047 (AMS1000)) together. The numbers of ShadowImage pairs that can be executed are calculated by subtracting the number of migration pairs from the maximum number of pairs.

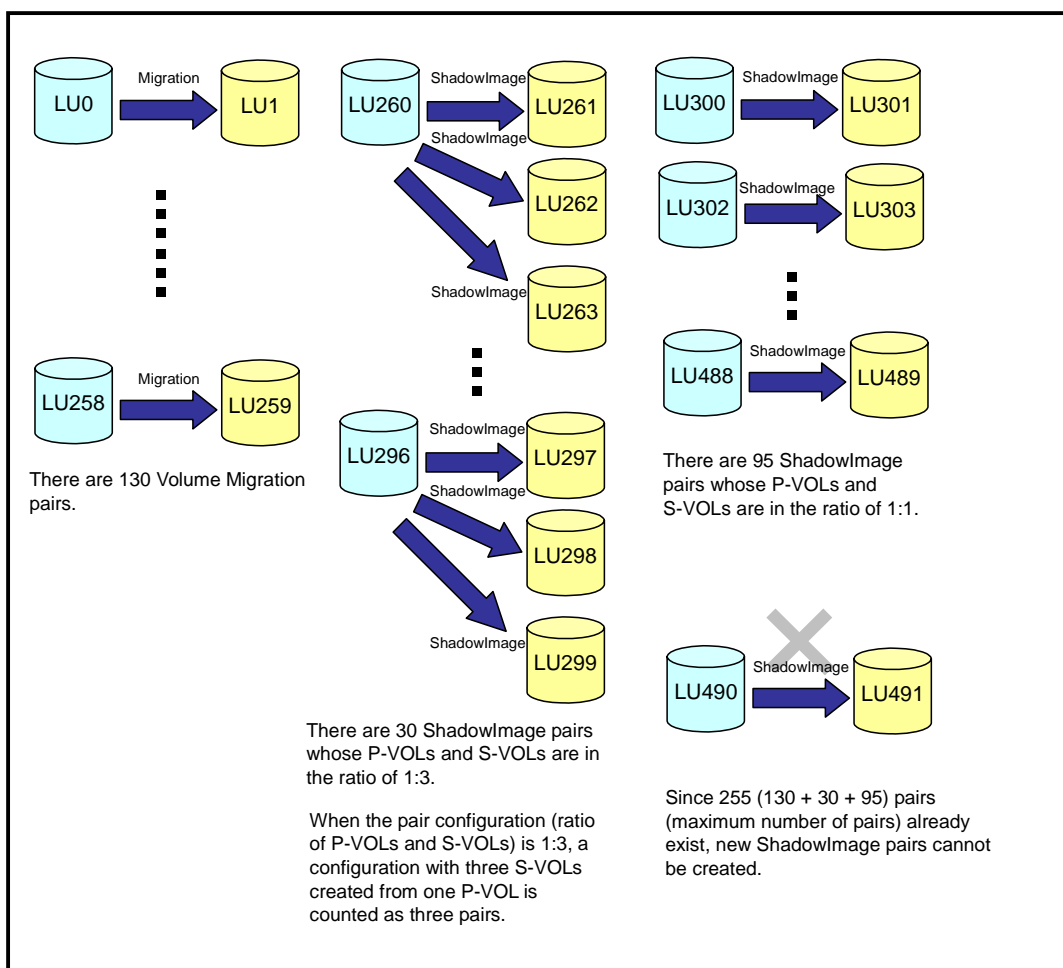


Figure 4.11 An Example in which the Number of Pairs of ShadowImage is Limited

- The number of pairs that data copying can be performed in the background

The number of copying operations that can be performed at the same time in the background is called the copying multiplicity. The disk subsystem limits the copying multiplicity of both the Volume Migration pairs and ShadowImage pairs together to four per controller. When ShadowImage is used together with Volume Migration, the copying multiplicity of ShadowImage becomes smaller than the maximum number (two per controller) because ShadowImage and Volume Migration share the copying multiplicity.

Because the disk subsystem basically executes the copying operations conducted by ShadowImage and Volume Migration in order as instructed, there may be a case where the copying is not started immediately after ShadowImage instructs the execution of the pair creating or resynchronization.

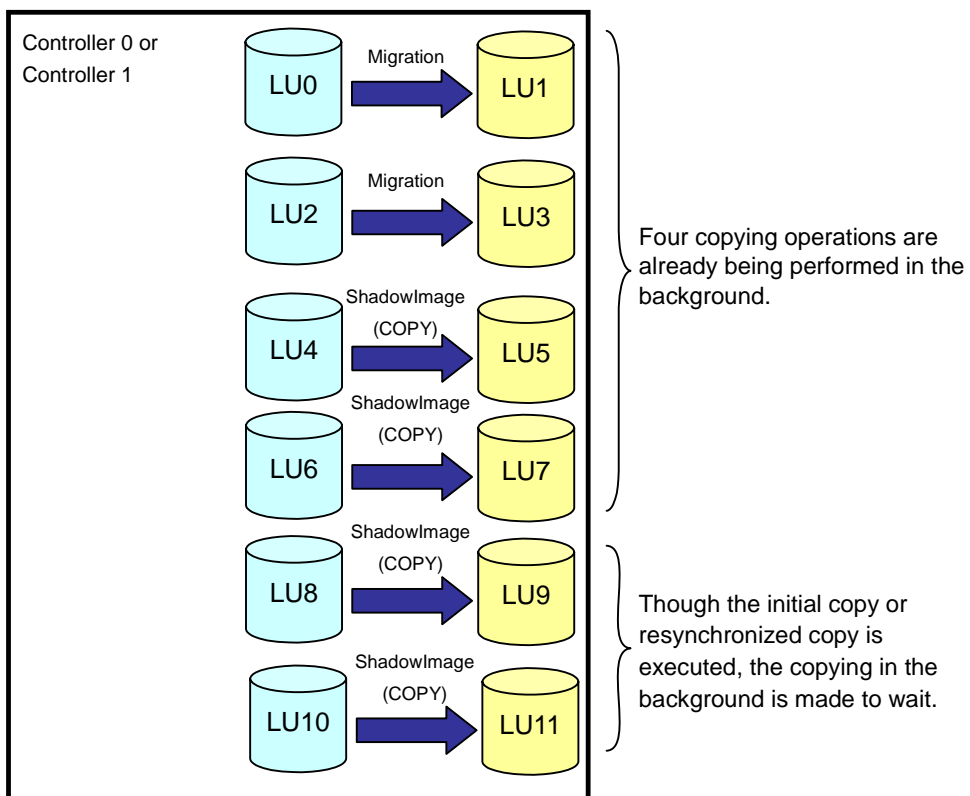


Figure 4.12 An Example of a Case where the Copying Operation of ShadowImage is made to Wait

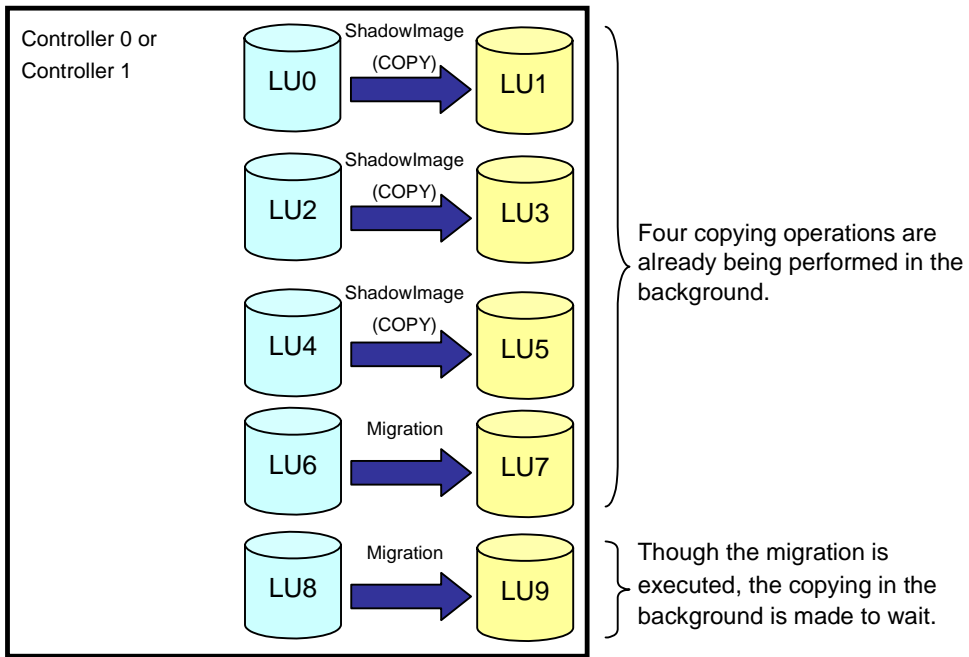
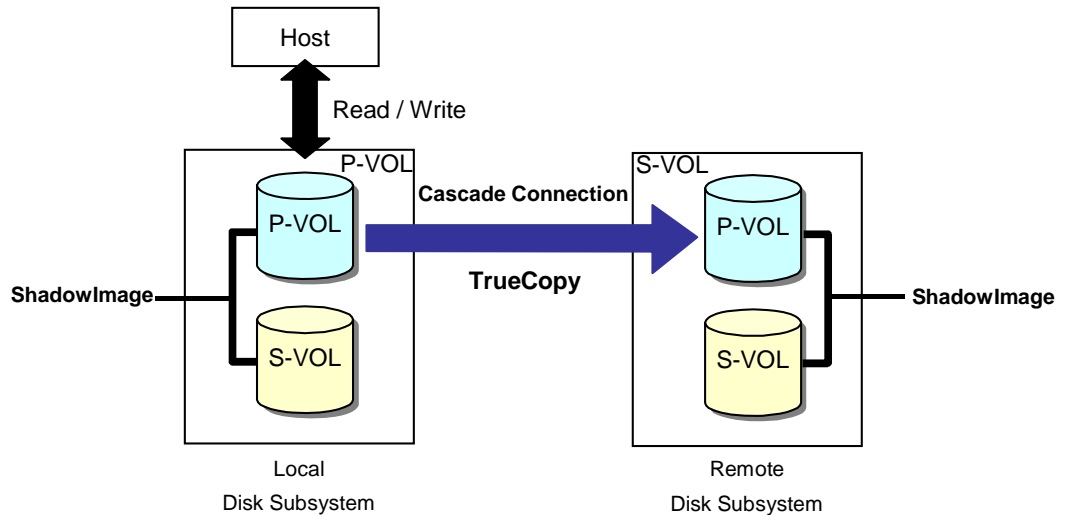


Figure 4.13 An Example of a Case where the Copying Operation of Volume Migration is made to Wait

4.5 Cascade Connection of TrueCopy with ShadowImage

Volumes of TrueCopy can be cascaded with those of ShadowImage as shown in Figure 4.14. TrueCopy does not always need ShadowImage. SnapShot/TCE cannot be cascaded with ShadowImage.

- Cascade with a P-VOL of ShadowImage



- Cascade with an S-VOL of ShadowImage

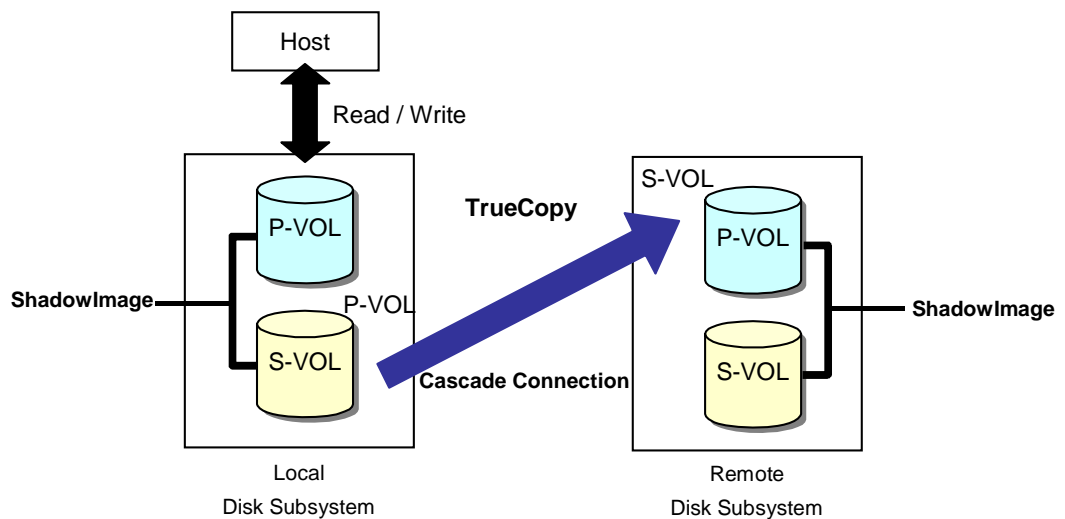


Figure 4.14 Cascade Connection of TrueCopy with ShadowImage (P-VOL: S-VOL=1: 1)

- Cascade with a P-VOL and an S-VOL of ShadowImage
 (Following configuration is supported when the micro program version is more than 0772/A)

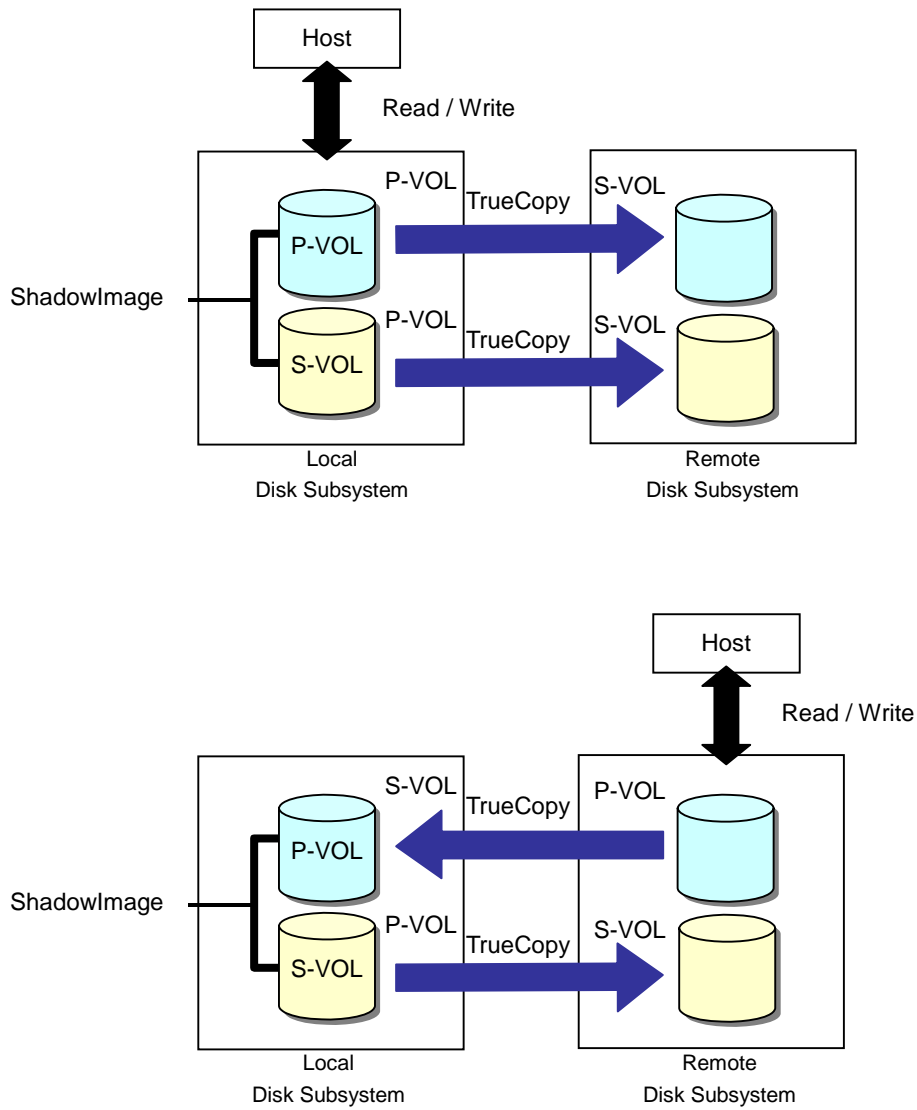
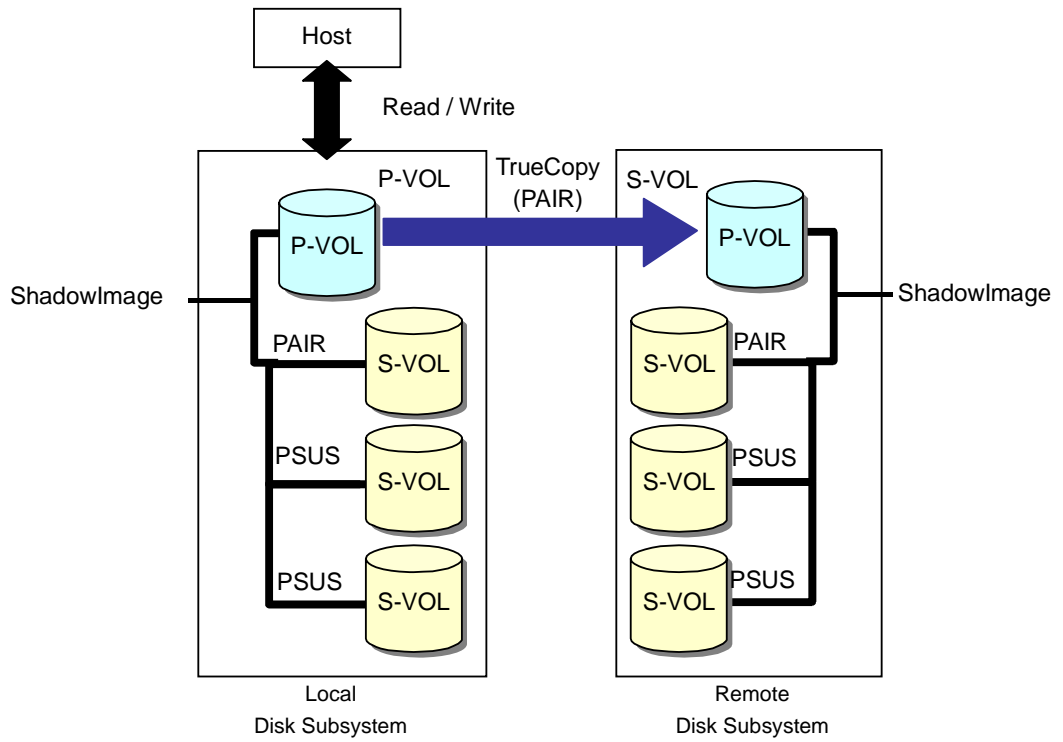


Figure 4.15 Simultaneous Cascade Connection of TrueCopy with ShadowImage (P-VOL: S-VOL=1: 1)

■ Cascade with a P-VOL of ShadowImage



■ Cascade with an S-VOL of ShadowImage

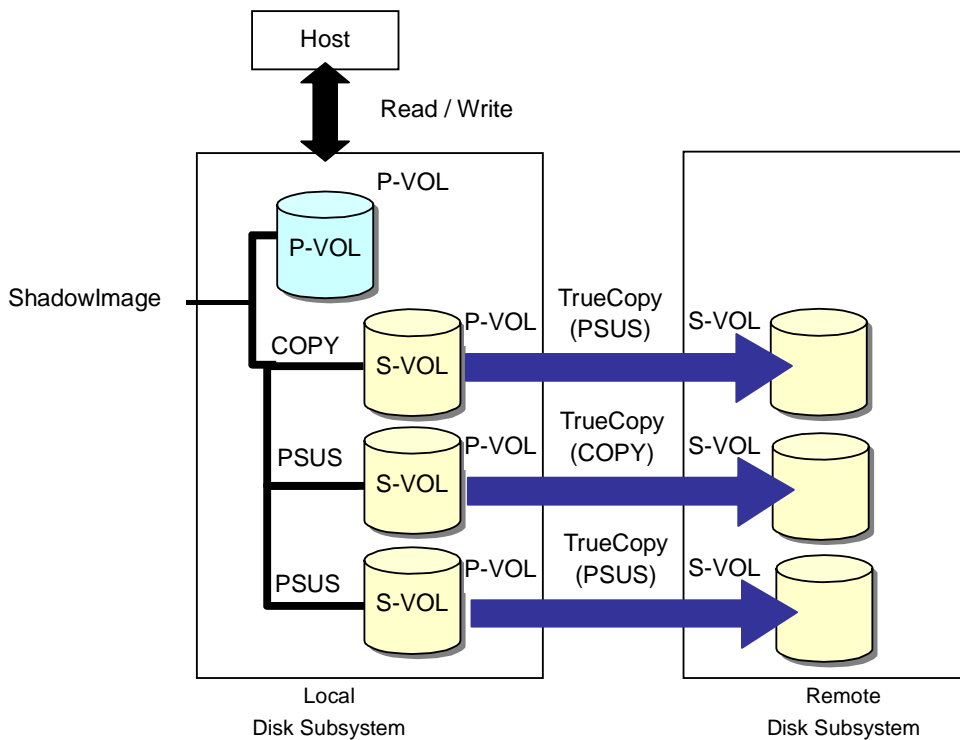
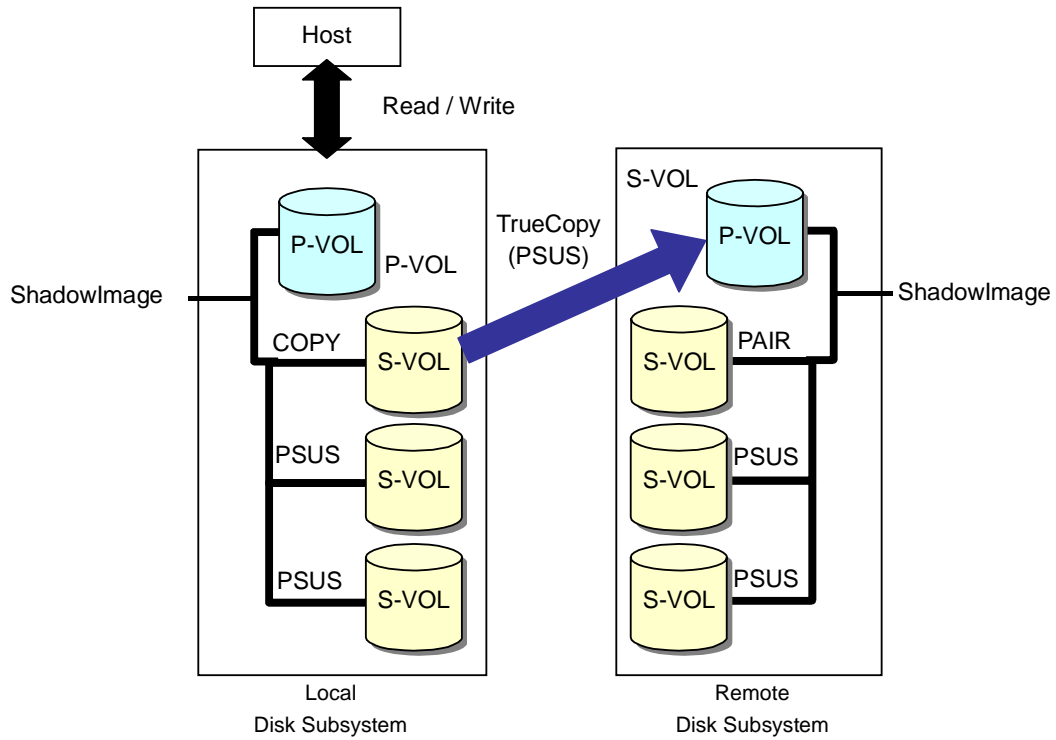


Figure 4.16 Cascade Connection of TrueCopy with ShadowImage (P-VOL: S-VOL=1: 3)

- Cascade with a P-VOL and an S-VOL of ShadowImage
 (Following configuration is supported when the micro program version is more than 0772/A)

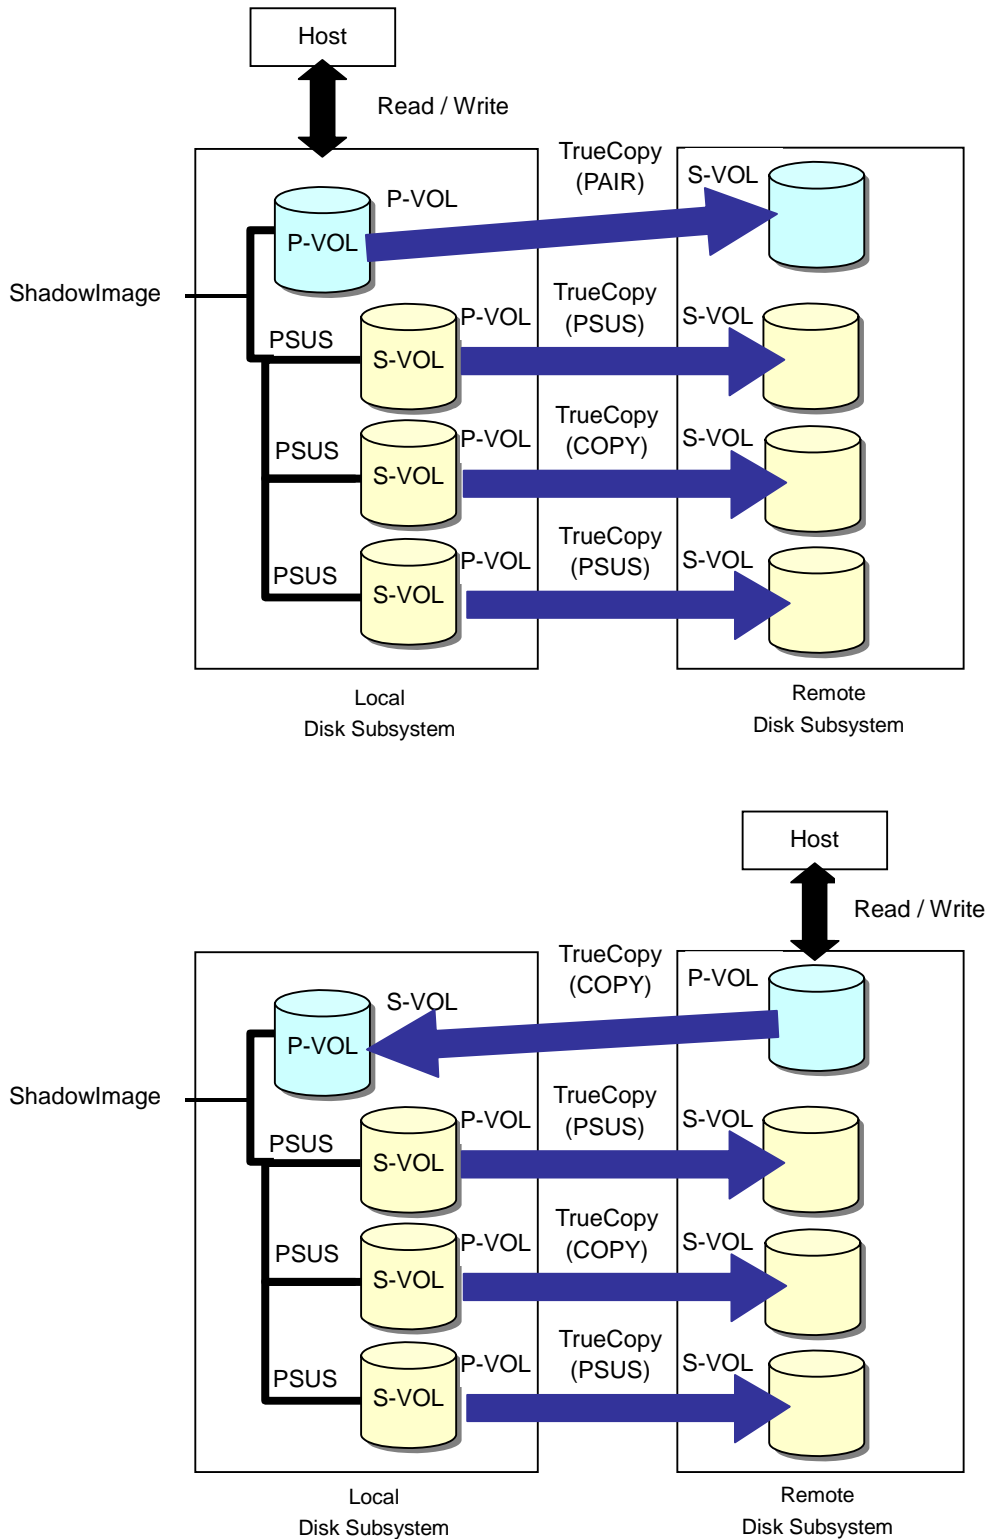


Figure 4.17 Simultaneous Cascade Connection of TrueCopy with ShadowImage (P-VOL: S-VOL=1: 3)

4.5.1 Cascade Restrictions with ShadowImage P-VOL

When using TrueCopy, if the P-VOL on ShadowImage is cascaded, due to reverse-resync being executed by ShadowImage, the data of the P-VOL on the local side might not equal the data of P-VOL on the remote side. When reverse-resync using ShadowImage is executed, TrueCopy must have a PSUS status. If reverse-resync using ShadowImage is executed in TrueCopy COPY status or PAIR status, 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 equality.

When reverse-resync using ShadowImage is executed, TrueCopy must have a PSUS status. If reverse-resync using ShadowImage is executed in COPY status or PAIR 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 equality.

- LU Shared with P-VOL on ShadowImage and P-VOL on TrueCopy:

Table 4.2 displays whether or not a read/write from/to a P-VOL of ShadowImage on the local side is possible when a P-VOL of ShadowImage and a P-VOL of TrueCopy are the same LU.

Table 4.2 A Read/Write Instruction to a P-VOL of ShadowImage on the Local Side

		ShadowImage P-VOL						
		PAIR	COPY	RCPY		PSUS	PSUE	PSUE (Reverse Resync)
				R/W	W			
TrueCopy P-VOL	PAIR	○ R/W	○ R/W	×	×	○ R/W	○ R/W	×
	COPY	○ R/W	○ R/W	×	×	○ R/W	○ R/W	×
	PSUS R/W	○ R/W	○ R/W	○ R/W	○ W	○ R/W	○ R/W	△ RAW
	PSUS R	○ R	○ R	×	×	○ R	○ R	×
	PSUE R/W	○ R/W	○ R/W	△ R/W	△ W	○ R/W	△ R/W	△ RAW
	PSUE R	○ R	○ R	×	×	○ R	△ R	×
	PSUE RAW	○ RAW	○ RAW	×	×	○ RAW	△ RAW	×

○: Possible. ×: Not Possible. △: When a pair operation causes an error (can occur as a result of a change of the pair status to PSUE) R/W: Read/Write by a host is possible.

R: Read by a host is possible but write is not possible.

W: Write by a host is possible but read is not possible.

Note 1: PSUE in this table excludes conditions in which LU access is not possible (e.g., LU blockage).

Note 2: 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 PSUS status, the item of PSUS is applied. ② If all the pairs that the P-VOL concerned configures are in the PSUS status or the PSUE status, the item of PSUS is applied. However, when including the pair that became PSUE during the reverse resynchronization, the items of PSUE(R) are applied. ③ If a pair in the PAIR status, the COPY status, or the RCPY status is included in the pair that the P-VOL concerned configures, the item of PAIR, COPY, and RCPY is applied, respectively. (Two or more pairs in the PAIR status, the COPY status, and the RCPY status are never included in the pair that the P-VOL concerned configures.)

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

Table 4.3 displays whether or not a read/write from/to a P-VOL of ShadowImage on the remote side is possible when a P-VOL of ShadowImage and an S-VOL of TrueCopy are the same LU.

Table 4.3 A Read/Write Instruction to a P-VOL of ShadowImage on the Remote Side

		ShadowImage P-VOL						
		PAIR	COPY	RCPY		PSUS	PSUE	PSUE (Reverse Resync)
				R/W	W			
TrueCopy S-VOL	PAIR	○ R	○ R	×	×	○ R	○ R	×
	COPY	○ R	○ R	×	×	○ R	○ R	×
	PSUS R/W	○ R/W	○ R/W	○ R/W	○ W	○ R/W	○ R/W	△ R/W
	PSUS R	○ R	○ R	×	×	○ R	○ R	×
	PSUE	○ R	○ R	×	×	○ R	△ R	×

○: Possible. ×: Not Possible. △: When a pair operation causes an error (can occur as a result of a change of the pair status to PSUE)

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

R: Read by a host is possible but write is not possible.

W: Write by a host is possible but read is not possible.

R/W: Read/Write by a host is not possible.

Note 1: PSUE in this table excludes conditions in which LU access is not possible (e.g., LU blockage).

Note 2: 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 PSUS status, the item of PSUS is applied. ② If all the pairs that the P-VOL concerned configures are in the PSUS status or the PSUE status, the item of PSUS is applied. However, when including the pair that became PSUE during the reverse resynchronization, the items of PSUE(R) are applied. ③ If a pair in the PAIR status, the COPY status, or the RCPY status is included in the pair that the P-VOL concerned configures, the item of PAIR, COPY, and RCPY is applied, respectively. (Two or more pairs in the PAIR status, the COPY status, and the RCPY status are never included in the pair that the P-VOL concerned configures.)

4.5.2 Cascade Restrictions with S-VOL of ShadowImage

Cascading for an LU of TrueCopy with an S-VOL of ShadowImage is supported only when the S-VOL of ShadowImage and a P-VOL of TrueCopy are the same LU. Operations of ShadowImage and TrueCopy pairs are restricted, depending on the status of the pairs.

When cascading volumes of TrueCopy with an S-VOL of ShadowImage, a ShadowImage pair is created first. 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 PSUS or PSUE. When changing the status of a TrueCopy pair, the status of a ShadowImage pair must be PSUS.

Table 4.4 displays whether or not a read/write from/to an S-VOL of ShadowImage on the local side is possible when an S-VOL of ShadowImage and a P-VOL of TrueCopy are the same LU.

Table 4.4 A Read/Write Instruction to an S-VOL of ShadowImage on the Local Side

		PAIR	COPY	RCPY	PSUS	PSUE	PSUE (Reverse Resync)
TrueCopy P-VOL	PAIR	×	×	×	○ R/W	×	×
	COPY	×	×	×	○ R/W	×	×
	PSUS R/W	○ R	○ R	○ R	○ R/W	△ R	△ R/W
	PSUS R	○ R	○ R	○ R	○ R	△ R	△ R/W
	PSUE R/W	○ R/W	○ R	○ R	○ R/W	△ R	△ R/W
	PSUE R	○ R	○ R	○ R	○ R	△ R	△ R/W
	PSUE R/W	○ R/W	○ R/W	○ R/W	○ R/W	△ R/W	△ R/W

O: Possible. ×: Not Possible. Δ: When a pair operation causes an error (can occur as a result of a change of the pair status to PSUE)

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

R: Read by a host is possible but write is not possible.

W: Write by a host is possible but read is not possible.

R/W: Read/Write by a host is not possible.

Note: PSUE in this table excludes conditions in which LU access is not possible (e.g., LU blockage).

4.5.3 Cascade Restrictions with POOL of SnapShot

ShadowImage pair cannot be created using a POOL.

4.5.4 Restrictions Configuration on Cascade of TrueCopy with ShadowImage

The following is an example of a configuration in which restrictions are placed on cascading TrueCopy with ShadowImage.

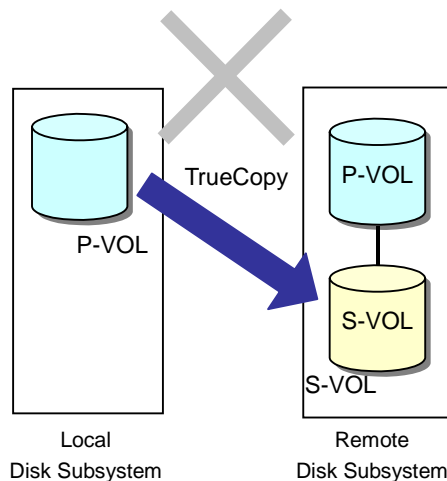


Figure 4.18 Restrictions for Cascade Configuration of TrueCopy with ShadowImage

4.6 ShadowImage I/O Switching Function

This function switches an object of an I/O instruction issued by a host from a P-VOL to an S-VOL automatically. This makes access from the host continuous when a P-VOL of ShadowImage becomes inaccessible due to a drive double failure (triple failures for RAID 6). (Hereafter, this function is referred to as the ShadowImage I/O Switching Function.)

This function operates only when the pair configuration is P-VOL: S-VOL=1:1.

4.6.1 ShadowImage I/O Switching Function Operation

When the ShadowImage I/O Switching Function operates during a drive double failure (triple failures for RAID 6), pair status is changed to "PSUE (S-VOL Switch)" and the object of a host I/O instruction is switched from a P-VOL to an S-VOL. Since a report is made as if sent from a P-VOL, a job can continue without interruption.

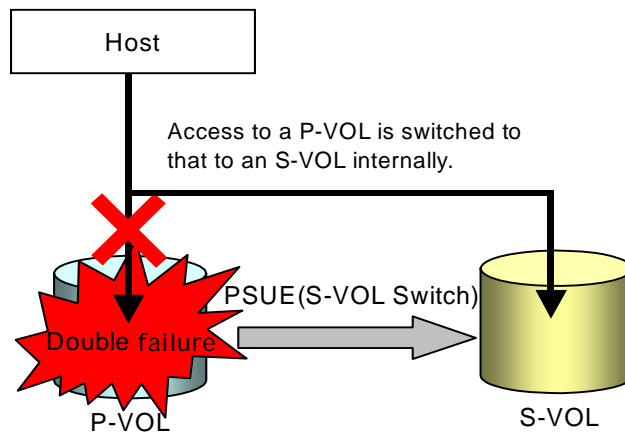


Figure 4.19 ShadowImage I/O Switching Function Operation

4.6.2 Requirements for Using the ShadowImage I/O Switching Function

Requirements for using the ShadowImage I/O Switching Function are shown below:

- Turn on the ShadowImage I/O Switching mode through Navigator. For details, see section 5.5.
- It operates only when the pair is in the PAIR status.
- A pair cannot be cascaded with a TrueCopy pair. When it is cascaded with a TrueCopy pair, it is placed in the ordinary PSUE status. Also, when a host I/O instruction is switched from a P-VOL to an S-VOL, it cannot be cascaded with a TrueCopy pair.
- The pair configuration needs to be P-VOL: S-VOL=1:1. Also, a S-VOL cannot be created newly from the P-VOL whose I/O is switched to the S-VOL.

Note 1: If a P-VOL or S-VOL of ShadowImage exists in a RAID Group in which a double (triple, in the case of RAID 6) drive failure occurred when the ShadowImage I/O Switching mode is turned on in the case where the micro program version is 0751/A or more, all the LUs in the RAID Group become unformatted irrespective of whether they are ShadowImage pairs or not.

Note 2: When ShadowImage I/O Switching mode is turned on and the micro program version is less than 0751/A, all LUs of the RAID group in which the drive double failure (triple failures for RAID 6) occurred, become unformatted. This will occur regardless if they are, or are not, ShadowImage pair LU's.

4.6.3 Recovery from a Double Drive Failure

When the ShadowImage I/O Switching Function is used, an operation for recovery from a failure can be performed while the job is being continued. The procedure for recovery from a failure is shown in Figure 4.20:

1. When a double failure of drives occurs (triple failures for RAID 6) in a P-VOL and an object is switched from an I/O instruction to an S-VOL, call the service personnel and request that the drives composing the P-VOL in which the failure occurred be replaced.
2. Perform quick formatting for the P-VOL in which the double failure (triple failures for RAID 6) occurred.
3. When the `pairresync` command is executed after the quick formatting is completed, S-VOL data is copied to the P-VOL. Though the performance is lowered while the `pairresync` command is being executed, execution of a host I/O can be continued.

Note: The `pairresync` command cannot be executed until the quick formatting performed in the background is completed.

Note: When the pair status is PSUE (S-VOL Switch), execution of the `pairresync-restore` command will result in an error.

4. When execution of the `pairresync` command is completed, the pair enters the PAIR status, a normal status.

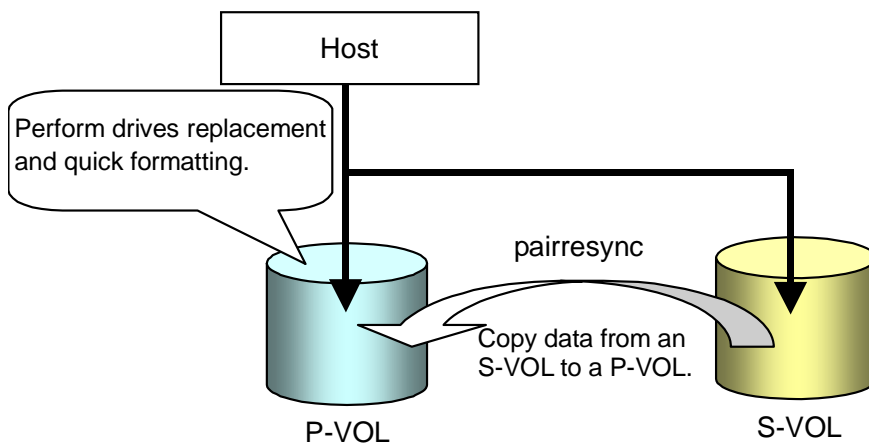


Figure 4.20 Recovering from a Double Drive Failure (Triple Failures for RAID 6)

4.6.4 Configuration and Operation

- Locating P-VOLs and S-VOLs:

When the same RAID group is defined as both a P-VOL and an S-VOL of ShadowImage and a double drive failure occurs in the RAID group, both the P-VOL and the S-VOL become unformatted and their pair statuses are changed to PSUE because an object of an I/O instruction cannot be switched to an S-VOL. Therefore, define respective RAID groups as a P-VOL and an S-VOL.

- Performance of the Host:

When using the ShadowImage I/O Switching Function, the pair status must be PAIR when doing a typical job (operation with the pair in the PAIR status). Note that the performance of an operation with the pair in PAIR status is inferior to that of an operation with a split pair, in which the pair status is changed to PSUS, when performing a typical job. It is recommended that you assign an LU using the FC drives to an S-VOL because the performance is remarkably lowered when an LU using the SATA drives is assigned to an S-VOL.

- Backup:

Back up all LUs in the subsystem; an LU of a RAID group in which a double-drive failure occurs becomes unformatted regardless of whether or not it is the mate to a ShadowImage pair.

4.6.5 ShadowImage I/O Switching Function Specifications

Table 4.5 shows a list of specifications for ShadowImage I/O Switching Function.

Table 4.5 ShadowImage I/O Switching Function Specifications

Items	Requirements
Preconditions for operation	<ul style="list-style-type: none"> ▪ The ShadowImage I/O Switching mode must be turned on. ▪ The pair status must be PAIR. ▪ A pair must not be cascaded with a Synchronous True Copy pair. ▪ The pair configuration needs to be P-VOL: S-VOL=1:1.
Scope of application	All ShadowImage pairs that satisfy the preconditions for operation.
Access to a P-VOL	Execution of a host I/O can be continued because a report is sent to the host as if a sender were a P-VOL.
Access to an S-VOL	An I/O instruction issued to an S-VOL results in an error.
Display of the status	<p>For Navigator:</p> <ul style="list-style-type: none"> ▪ When an object of a host I/O is switched to an S-VOL, the pair status is displayed as PSUE (S-VOL Switch). ▪ When the pairresync command is executed for a pair in the PSUE (S-VOL Switch) status, the pair status is displayed as COPY (S-VOL Switch). <p>For CCI (pairdisplay command):</p> <ul style="list-style-type: none"> ▪ Even when an object of a 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 in the case where an object of a host I/O has been switched to an S-VOL, 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	<ul style="list-style-type: none"> ▪ If a P-VOL or S-VOL of ShadowImage and a NAS system LU are in the same RAID Group when the ShadowImage I/O Switching mode is turned on when the micro program version is 0751/A or later, the NAS system LU becomes unformatted due to a double (triple, in the case of RAID 6) drive failure. When the ShadowImage I/O Switching mode is turned on, create the P-VOL or S-VOL of ShadowImage and the NAS system LU in separate RAID Groups. ▪ The pairsplit, pairresync -restore or pairsplit -S command cannot be performed when the status is PSUE (S-VOL Switch) or copying PSUE (S-VOL Switch) to P-VOL.

Chapter 5 Preparing for ShadowImage Operations

The following preparation is required for ShadowImage operations:

- Installing and Uninstalling ShadowImage (see section 5.1)
- Setting the Command Device (see section 5.2)
- Setting the Differential Management LU (see section 0)
- Setting the Target ID (LU Mapping) (see section 5.3)
- Setting the ShadowImage I/O Switching Mode (see section 5.4)
- Defining the Configuration Definition File (see section 5.5)
- Setting the Environment Variable (see section 5.6)

5.1 Installing and Uninstalling ShadowImage

Since ShadowImage is an extra-cost option, ShadowImage cannot usually be selected (locked) when first using the TagmaStore. To make ShadowImage available, you must install the ShadowImage and make its function selectable (unlocked).

ShadowImage can be installed from the Navigator. This section describes the installation/uninstallation procedures performed by using the Navigator via the Graphical User Interface (GUI). For details on the GUI version of Navigator, refer to the *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Navigator Modular Graphical User Interface (GUI) User's Guide* (MK-95DF711).

For procedures performed by using the Command Line Interface (CLI) of Navigator, see Appendix A.

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

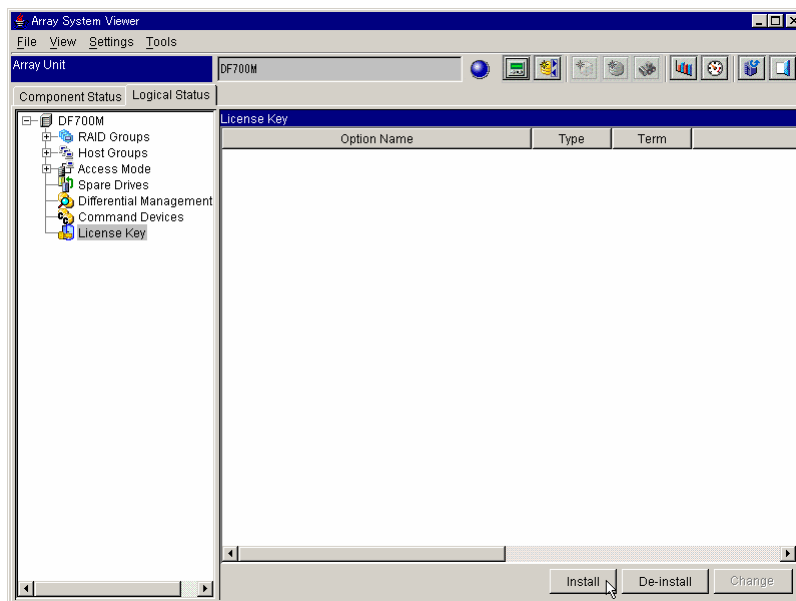
5.1.1 Installing ShadowImage

To install ShadowImage, the key code or key file provided with the optional feature is required. The following describes the installation procedure:

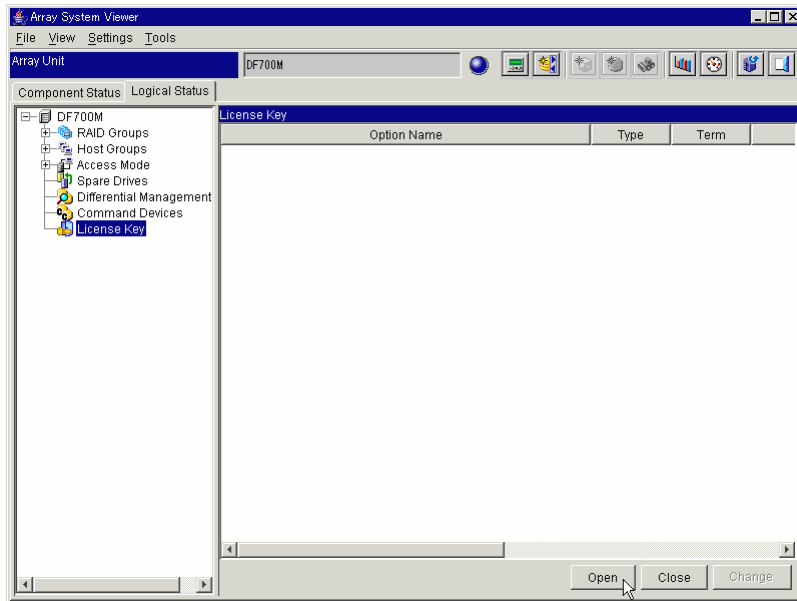
1. Start Navigator and switch to **Management Mode** (administrator mode).
2. Register the subsystem (array unit) in which ShadowImage is to be installed, and connect to the subsystem.

The **Array System Viewer** window and the connected subsystem displays.

3. Click the **Logical Status** tab.
4. Click the **License Key** icon.



Navigator: Version 5.00 or later




Navigator: Less than 5.00 version


Figure 5.1 Array System Viewer Window (Logical Status Tab)

5. Click **Install**. The **Install Options** dialog box displays. (Navigator: Version 5.00 or later)

Click **Open**. The **Unlock Options** dialog box displays. (Navigator: Less than 5.00 version)



Navigator: Version 5.00 or later



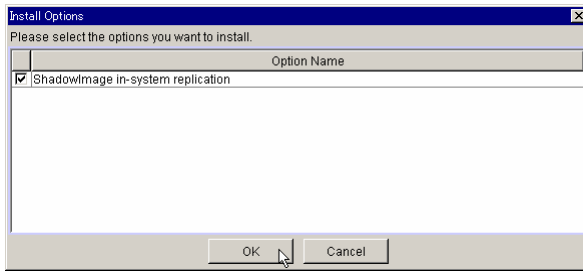
Navigator: Less than 5.00 version

Figure 5.2 Install/Unlock Options Dialog Box

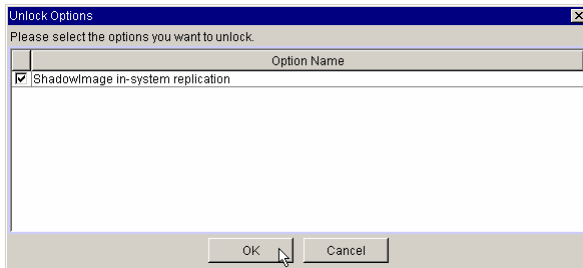
6. When you install the option using the key code, click the **Key Code** radio button, and then set up the key code. When you install the options using the key file, click the **Key File** radio button, and then set up the path for the key file. Click **OK**.

Note: **Browse** is used to set the path to a key file correctly.

7. When you install the options using the key file, the **Install/Unlock Options - Option Name** dialog box displays. Select the appropriate **Option Name**, and then click **OK**.




Navigator: Version 5.00 or later




Navigator: Less than 5.00 version

Figure 5.3 Install/Unlock Options – Option Name Dialog Box

8. A message box displays asking you to confirm that you want to install ShadowImage. Click OK.




Navigator: Version 5.00 or later

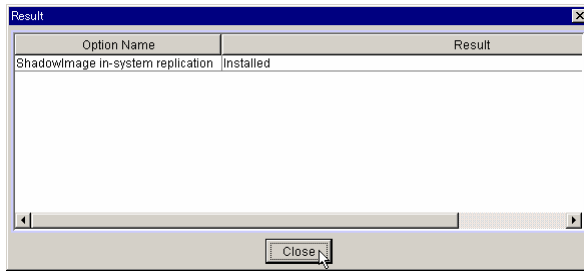


Navigator: Less than 5.00 version

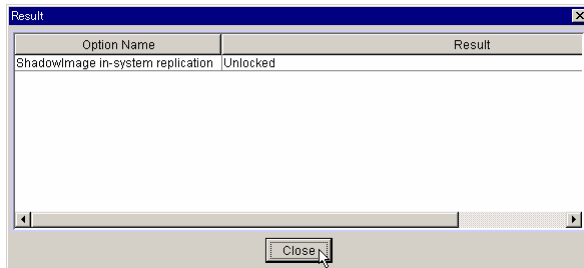
9. When Navigator version is 3.00 or later and Cache Partition Manager is enabled, the following message is displayed. Since ShadowImage does not use the data pool, click the OK button at this point without doing anything else.



10. When you install the options using the key file, the Result dialog box displays. Click Close.



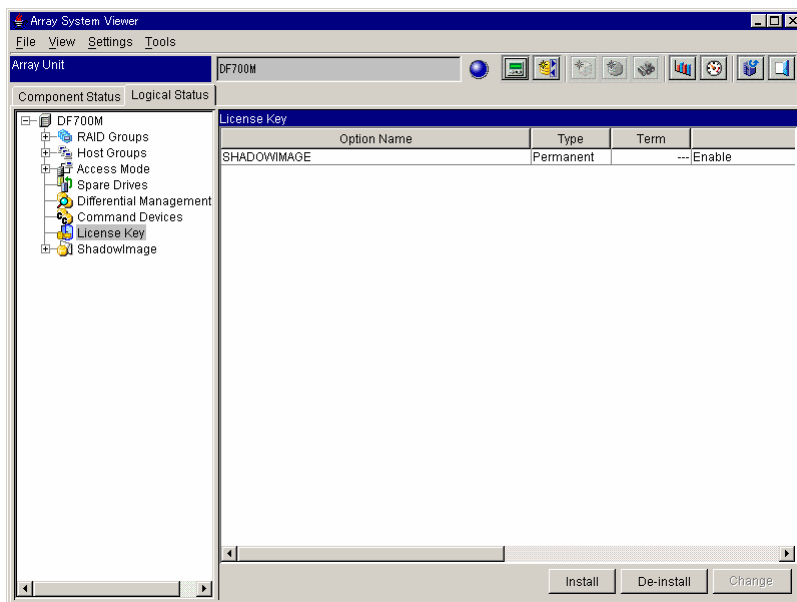
Navigator: Version 5.00 or later



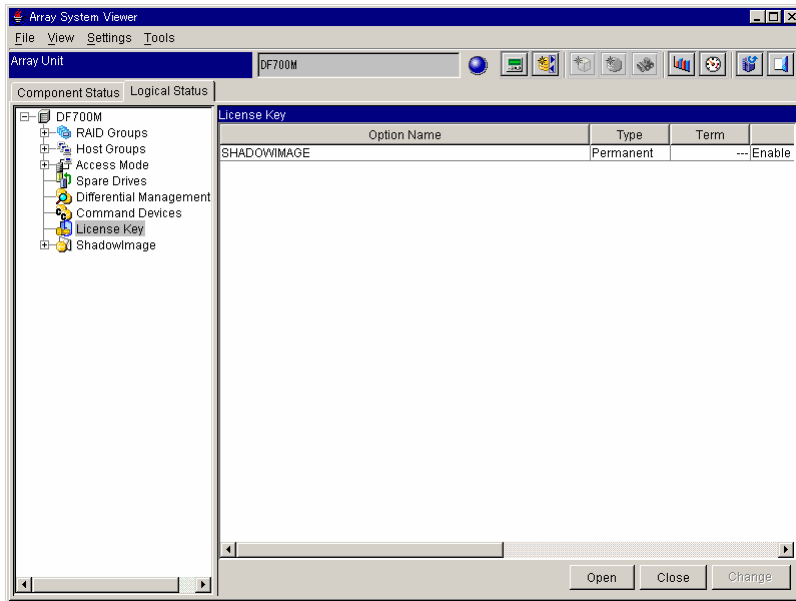
Navigator: Less than 5.00 version

Figure 5.4 Result Dialog Box

Figure 5.5 shows the Array System Viewer Window that displays when installation of ShadowImage is complete.



Navigator: Version 5.00 or later



Navigator: Less than 5.00 version

Figure 5.5 Array System Viewer Window (Logical Status Tab, ShadowImage: Enable)

5.1.2 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.

Note: The following conditions must be satisfied in order to uninstall ShadowImage.

All ShadowImage pairs must be released (the status of all LUs are SMPL).


To uninstall ShadowImage:

1. Start Navigator and switch to Management Mode (administrator mode).
2. Register the subsystem in which ShadowImage is to be uninstalled, then connect to the subsystem.


The Array System Viewer window opens displaying the connected subsystem.

3. Click the Logical Status tab.
4. Click the License Key icon.
5. Click De-install. The De-install Options dialog box displays. (Navigator: Version 5.00 or later)

Click Close. The Lock Options dialog box displays. (Navigator: Less than 5.00 version)



Navigator: Version 5.00 or later



Navigator: Less than 5.00 version

Figure 5.6 De-install/Lock Options Dialog Box

6. Enter the **Key Code**, and then click **OK**.
7. A message box displays asking you to confirm that you want to uninstall ShadowImage. Click **OK**.

ShadowImage is now uninstalled.

5.1.3 Enabling or Disabling ShadowImage

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

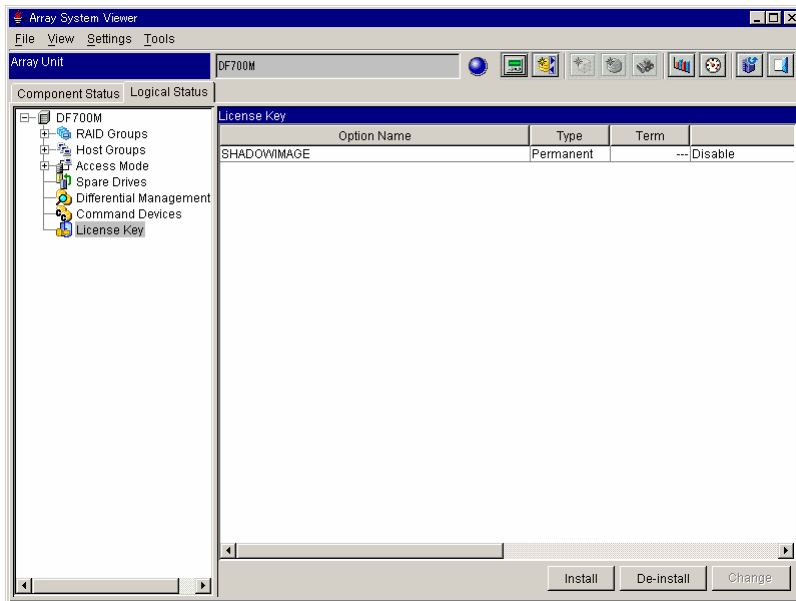
Note: The following conditions must be satisfied in order to disable ShadowImage.

All ShadowImage pairs must be released (that is, the status of all LUs are SMPL).

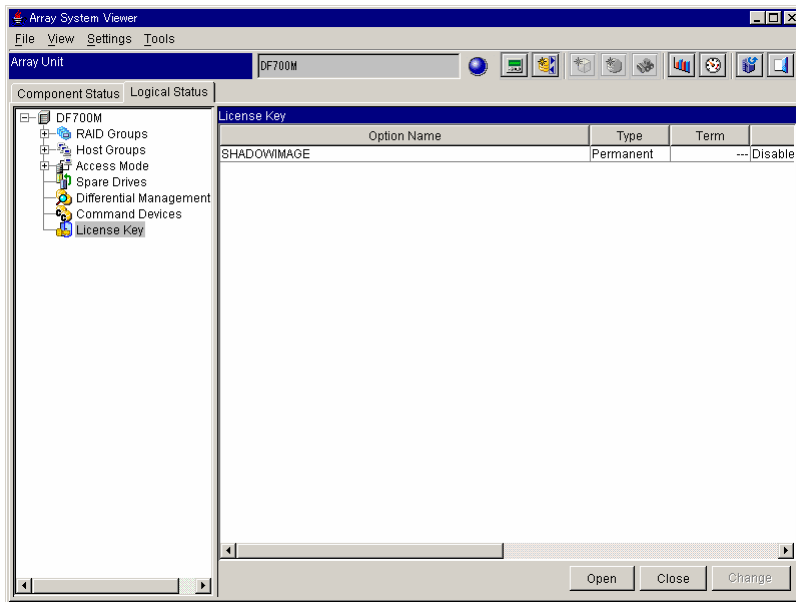
The following describes the enabling/disabling procedure.

1. Start Navigator and switch to **Management Mode** (administrator mode).
2. Register the subsystem (array unit) from which the status of ShadowImage is to be changed; then connect to the subsystem.
The **Array System Viewer** window opens displaying the connected subsystem.
3. Click the **Logical Status** tab.
4. Click the **License Key** icon.
5. From **Option Name**, click **SHADOWIMAGE**, and then click **Change**.
6. A message box displays asking you to confirm that you want to change the status (enable or disable). Click **OK**.

Figure 5.7 shows the Array System Viewer window that displays when enabling or disabling ShadowImage is complete.



Navigator: Version 5.00 or later



Navigator: Less than 5.00 version

Figure 5.7 Array System Viewer Window (Logical Status Tab, ShadowImage: Disable)

5.2 Setting the Command Device

The command device is a user-selected, dedicated logical volume on the TagmaStore subsystem, which functions as the interface to the CCI software. The ShadowImage commands are issued by CCI (HORCM) to the TagmaStore command device.

A command device must be designated in order to issue ShadowImage commands. The command device must be defined in the HORCM_CMD section of the configuration definition file for the CCI instance on the attached host. Two command devices can be designated for the TagmaStore subsystem. You can designate command devices using Navigator.

Note 1: LUs set for command devices must be recognized by the host. The Command Device LU size must be greater than or equal to 33 MBs.

Note 2: The following restrictions apply when either a pair of ShadowImage, SnapShot, TrueCopy or TCE exists, or the path of True Copy or TCE is defined.

- When two command devices are set, only one command device can be released.
- When only one command device is set, the command device cannot be released.

To designate command device(s):

1. Start Navigator and switch to **Management Mode** (administrator mode).
2. Connect to the subsystem. The **Array System Viewer** window opens displaying the connected subsystem.
3. Click the **Logical Status** tab.
4. Click the **Command Devices** icon.

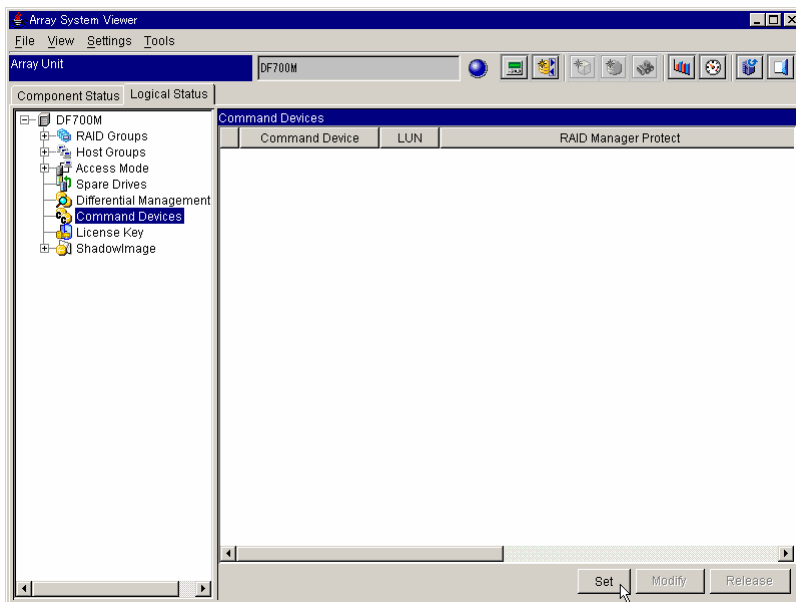


Figure 5.8 Array System Viewer Window (Command Device Page: Before Setting)

- Click **Set**. The **Command Devices Settings** dialog box displays.

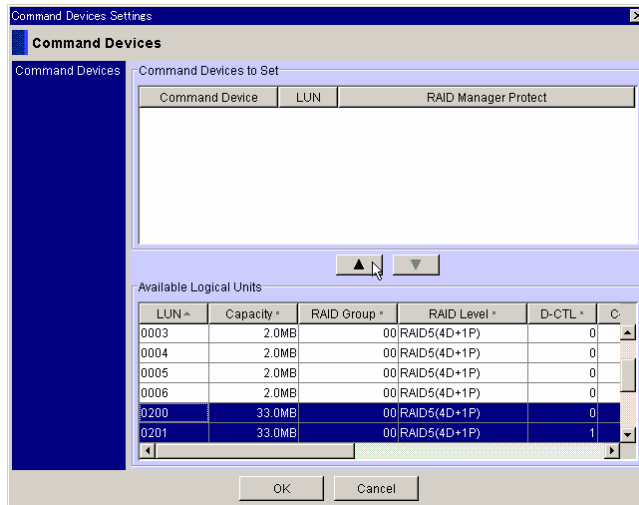


Figure 5.9 Command Devices Settings Dialog Box (Before Setting)

- In the **Available Logical Units** list, select the LUN you want to set the command devices, and then click the button.

The selected LUN moves to the **Command Devices to Set** list.

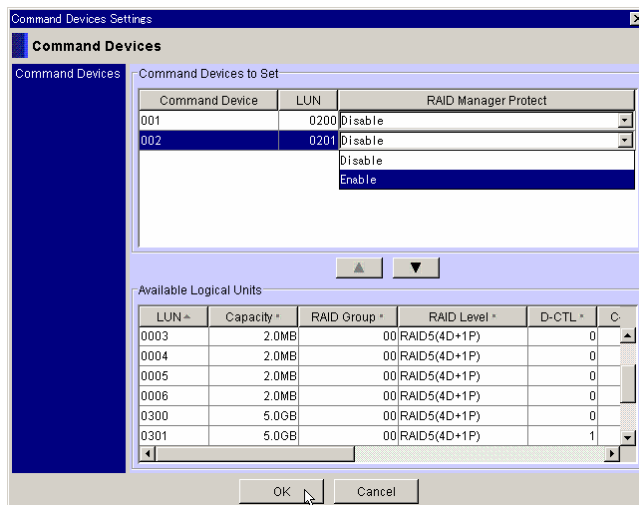


Figure 5.10 Command Devices Settings Dialog Box (After Setting)

- When you want to use the protection function of CCI, select **Disable** or **Enable** from the **RAID Manager Protect** drop-down list.

Note: When you select **Enable**, the **RAID Manager Protect** function is not available on NAS.

When you want to change the already set command devices, select the LUN on the **Command Devices to Set** list, and then click the button.

The selected LUN moves to the **Available Logical Units** list.

- Click OK. The setting information is displayed.

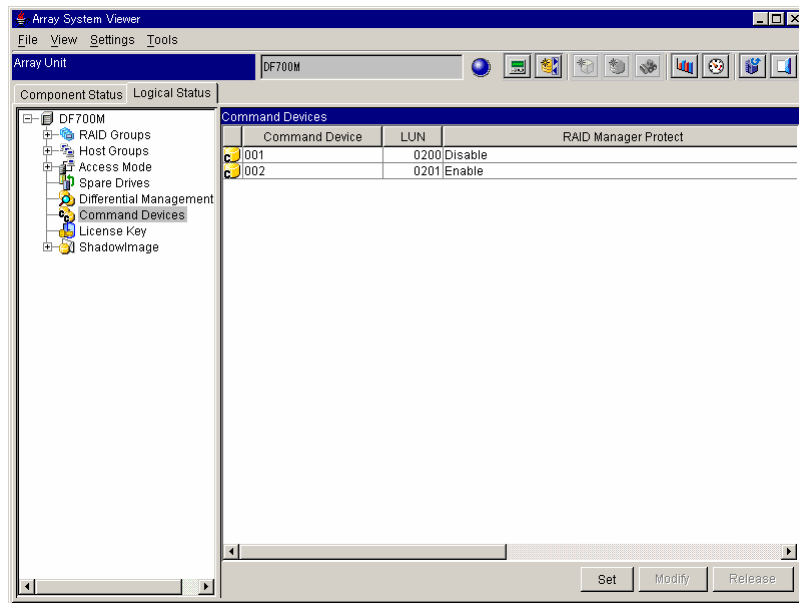


Figure 5.11 Array System Viewer Window (Command Device Page: After Setting)

Note: To use the alternate command device function or to avoid data loss and subsystem downtime, designate two command devices. When two command devices are set within one disk subsystem, assign them to the respective RAID groups; if they are assigned to the same RAID group, both command devices become unavailable due to an error (such as a drive failure). For details on the alternate command device function, refer to the *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Command Control Interface (CCI) User and Reference Guide* (MK-95DF701).

- To release an already set command device, select the LU number of the command device you want to release, and then click Release.
- A message box displays. Click OK.

Designating the command device is now complete. Leave the Array System Viewer window open and go to section 5.3 Setting the Target ID (LU Mapping).

Setting the Differential Management LU

When the Differential Management LU is not set, create it. The Differential Management LU is an exclusive logical unit for storing the differential data at the time when the volume is copied. The Differential Management LU in the disk subsystem is treated in the same way as the other logical units. However, a logical unit that is set as the Differential Management LU is not recognized by a host (it is hidden).

Set a logical unit with a size of 5 GB minimum as the Differential Management LU. Up to two Differential Management LUs can be set. The second one is used for the mirroring. It is recommended that two Differential Management LUs are set.

To designate Differential Management LUs:

1. In the **Array System Viewer** window, click the **Logical Status** tab.
2. Click the **Differential Management LU** icon.

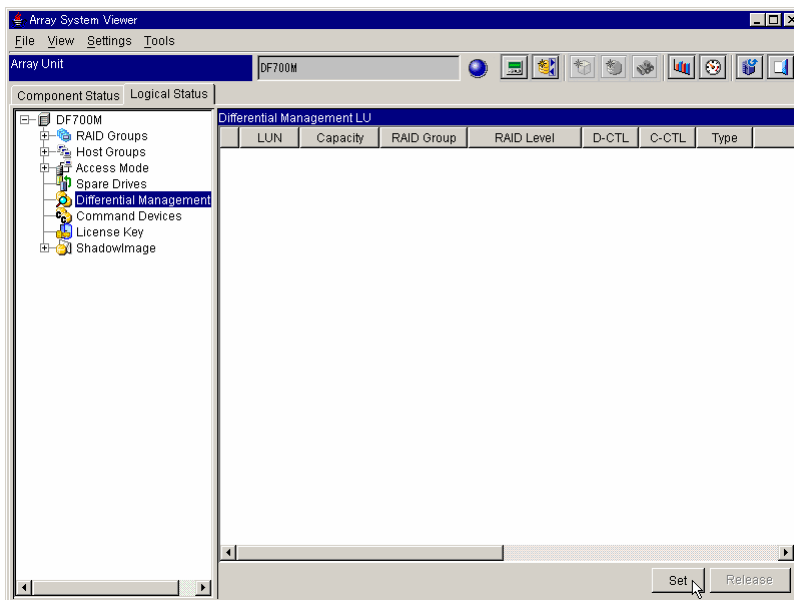


Figure 5.12 Array System Viewer Window (Differential Management LU Page: Before Setting)

3. Click **Set**. The **Select Logical Unit** dialog box displays.

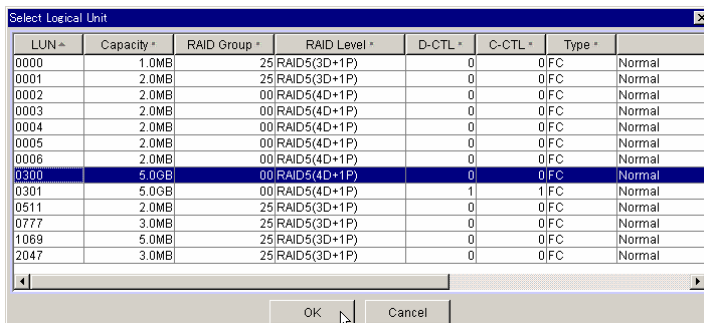


Figure 5.13 Select Logical Unit Dialog Box

4. Select the LUN you want to set the differential management LU, then click OK.
5. A message box displays. Click OK. The setting information displays in the Array System Viewer window.

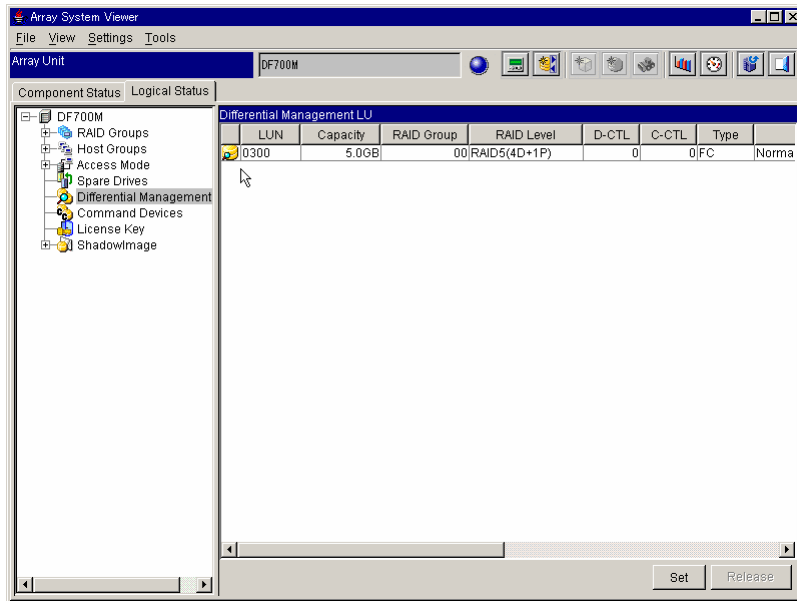


Figure 5.14 Array System Viewer Window (Differential Management LU Page: After Setting)

Designating the differential management LU is now complete. Leave the Array System Viewer Window open, and go on to section 5.3.

Note: The following restrictions apply when either a pair of ShadowImage, SnapShot, TrueCopy or TCE exists, the path of True Copy or TCE is defined, or POOL of SnapShot or TCE is defined:

- When two differential management LUs are set, only one differential management LU can be released.
- When only one differential management LU is set, the differential management LU cannot be released.

5.3 Setting the Target ID (LU Mapping)

The following is the procedure for mapping LUs. The target ID is specified using the Navigator. This section includes the following:

- Specifying Mapping Mode (see section 5.3.1)
- Setting Mapping Information (see section 5.3.2)

5.3.1 Specifying Mapping Mode

1. On the Array System Viewer Window, click the Logical Status tab.
2. Click the Access Mode plus signs next to the Mapping Mode.

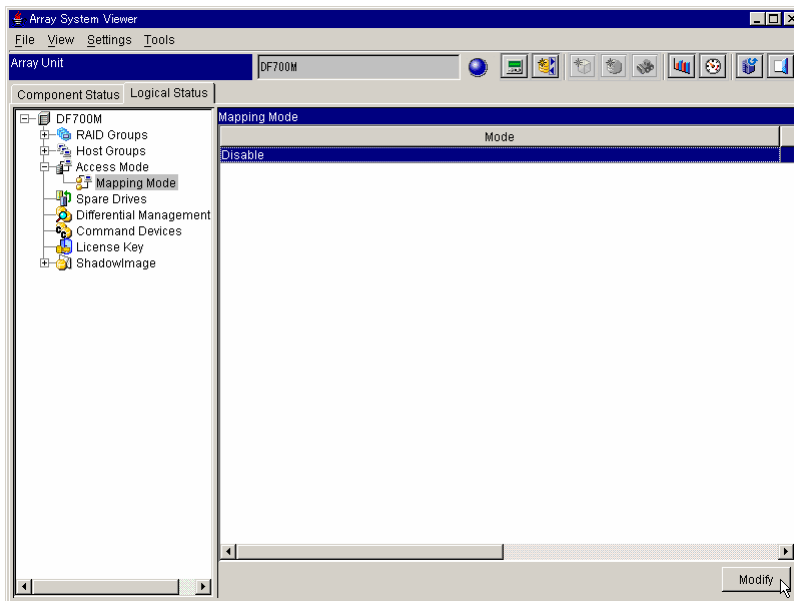


Figure 5.15 Array System Viewer Window (Specifying Mapping Mode)

3. Select **Disable** in the **Mapping Mode** list, and then click **Modify**. The **Mapping Mode** dialog box is displays.

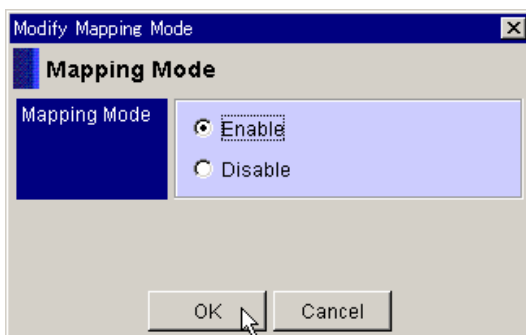
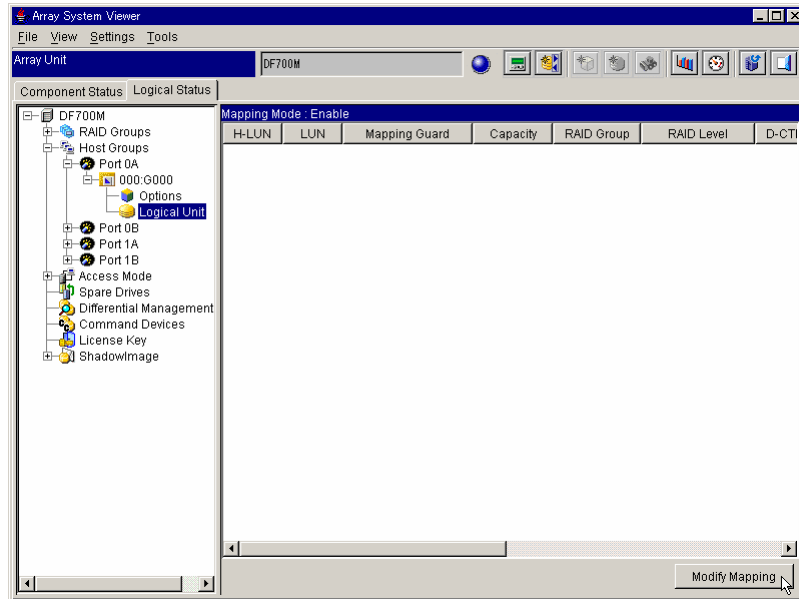


Figure 5.16 Mapping Mode Dialog Box

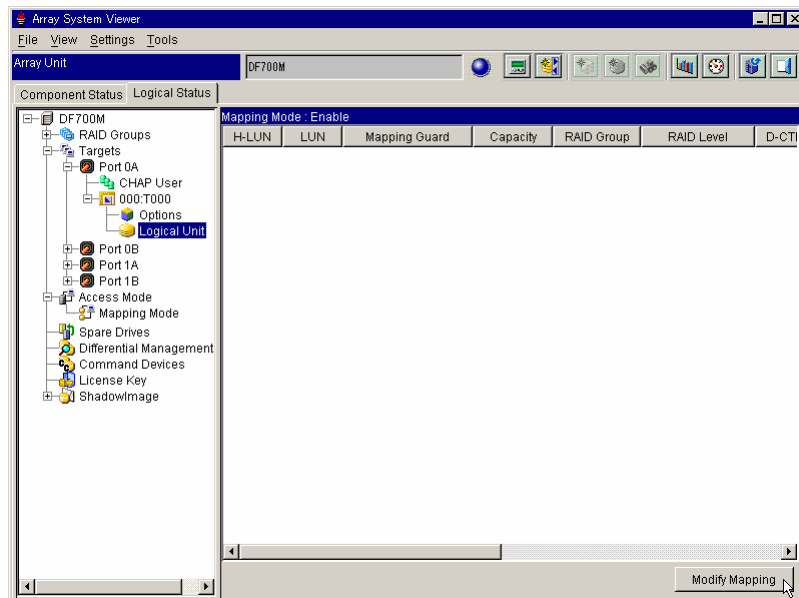
4. Select the **Enable** radio button, and then click **OK**.
5. A confirmation message box displays. Click **OK** three times.

5.3.2 Setting Mapping Information

1. Click **Port** plus signs next to 000:G000 or 000:T000, and then select 000:G000 or 000:T000.
2. Click the **Logical Unit** icon. Click **Modify Mapping**.



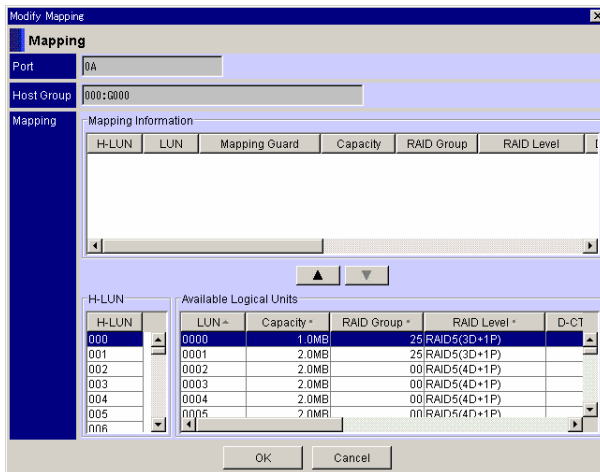
(Host Group)



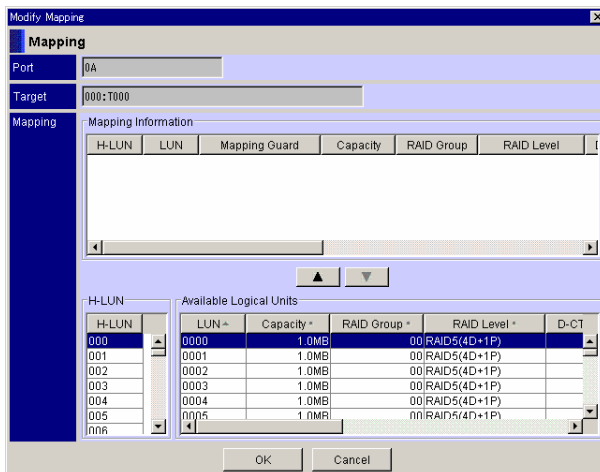
(Target)

Figure 5.17 Array System Viewer Window (Setting Mapping Information)

The Modify Mapping dialog box displays.




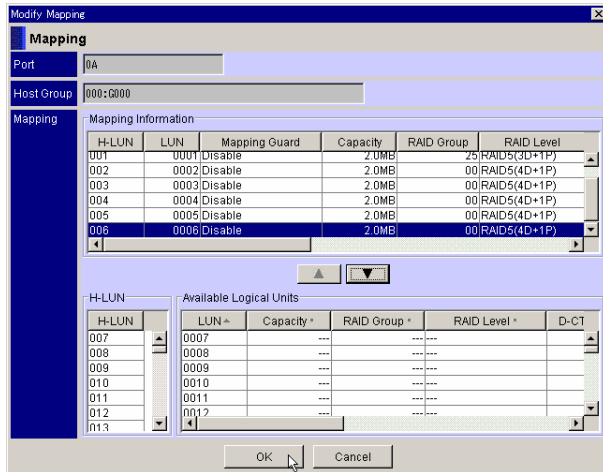
(Host Group)



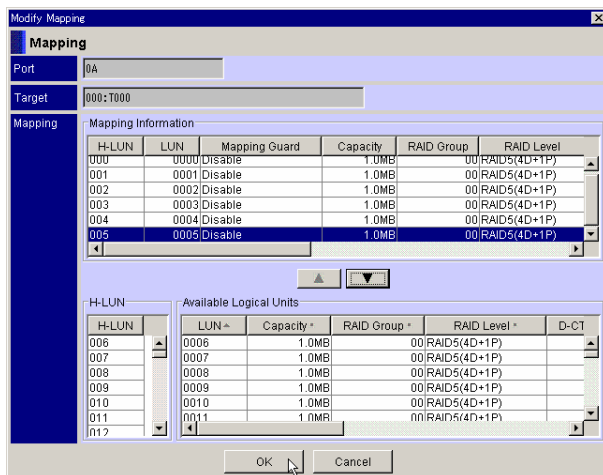
(Target)

Figure 5.18 Mapping Property Dialog Box (Before Setting)

3. Select one H-LUN to be added. Click LUN, and click the  button. The added contents display in the Mapping Information list:
 - For H-LUN, select the LU number that the host can recognize.
 - For LUN, select the LU number of the subsystem.



(Host Group)

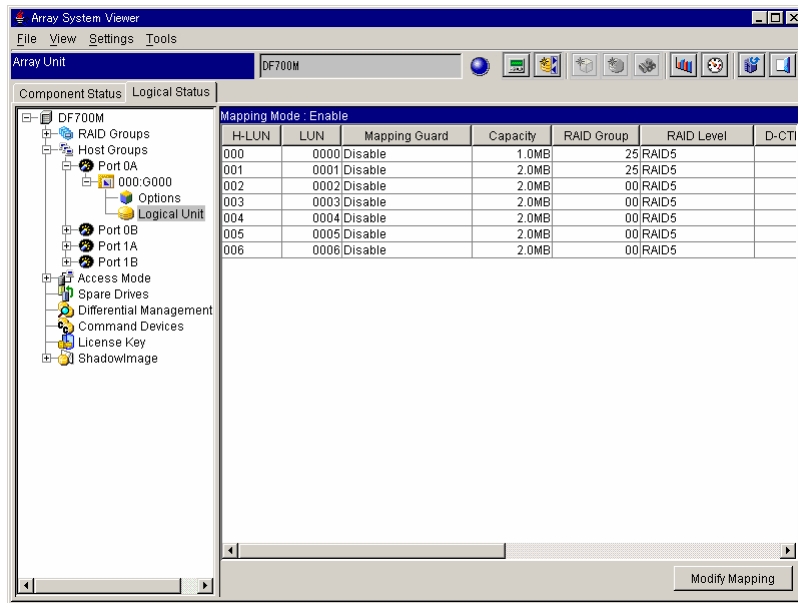


(Target)

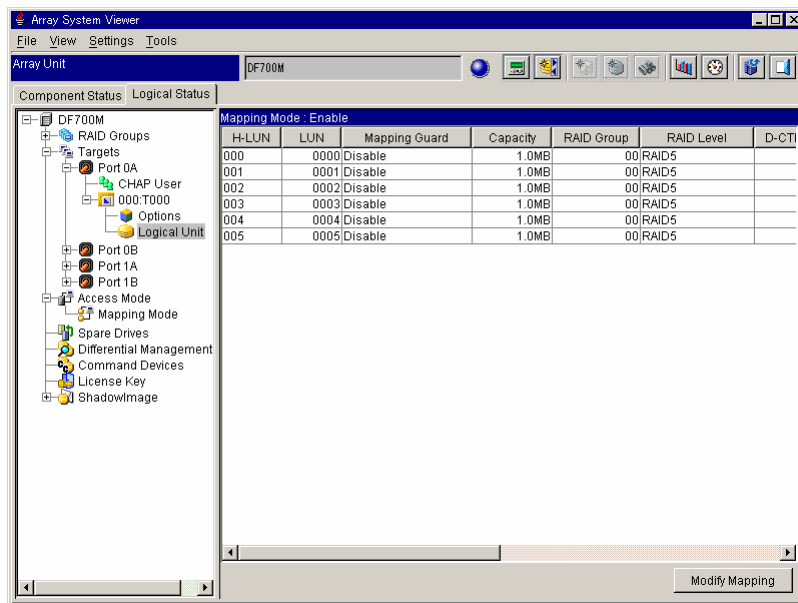
Figure 5.19 Mapping Property Dialog Box (After Setting)

4. Click OK.
5. A confirmation message box displays. Click OK.

- The set mapping information is updated and Array System Viewer window displays.



(Host Group)




(Target)

Setting the target ID is now complete.

5.4 Setting the ShadowImage I/O Switching Mode

When using the ShadowImage I/O Switching Function, set the ShadowImage I/O Switching mode. Set it with the Navigator.

To set the ShadowImage I/O Switching Mode:

1. Start the Navigator and switch to **Management Mode** (administrator mode).
2. Connect to the subsystem. The **Array System Viewer** window opens and displays the connected subsystem.
3. From the **Tools** menu, select **Configuration Settings**. Alternatively, from the tool bar, click the **Configuration Settings**  button. The **Configuration Settings** dialog box displays.
4. On the **Configuration Settings** dialog box, click the **System Parameter** tab. The **System Parameter** dialog box displays.

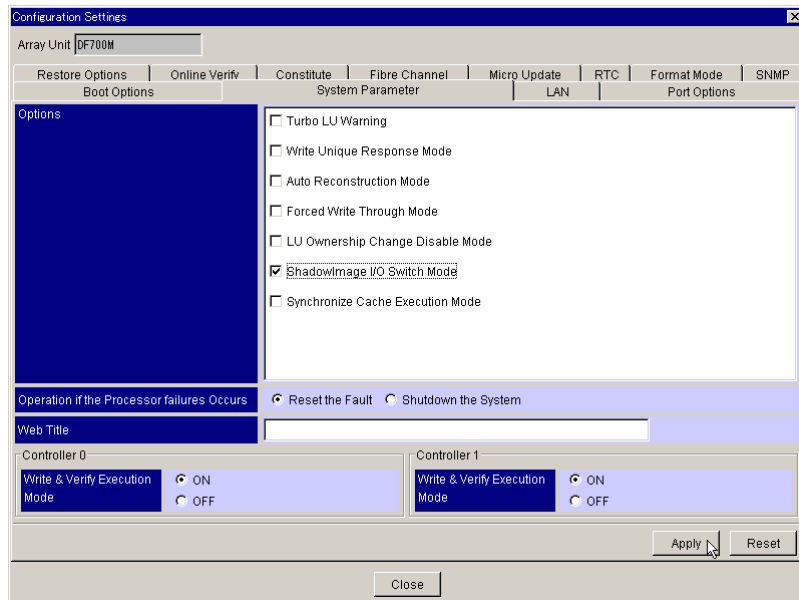


Figure 5.20 Configuration Settings Dialog Box (System Parameter Page)

5. Click the **ShadowImage I/O Switch Mode** in the **Options**.
6. Click **Apply**.
7. A confirmation message box displays. Click **OK**.
8. Click **Close**.

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 PSUS (S-VOL Switch) and COPY (S-VOL Switch).

5.5 Defining the Configuration Definition File

The configuration definition file describes the system configuration in order to make CCI operational. 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. This sample configuration definition file (HORCM_CONF), is included with the CCI software, and this file should be used as the basis for creating your configuration definition file(s). The system administrator should copy the sample file, set the necessary parameters in the copied file, and place the copied file in the proper directory. For details on configuration definition file, refer to the *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Command Control Interface (CCI) User and Reference Guide* (MK-95DF701).

The configuration definition file can be automatically created using the `mkconf` command tool. However, the parameters, such as `poll` (10ms) must be set manually (see step 4 below). For details on the `mkconf` command tool, refer to the *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Command Control Interface (CCI) User and Reference Guide* (MK-95DF701).

The following describes an example for manually defining the configuration definition file, when the system is configured with two instances within the same server (Windows NT/2000/Windows Server 2003):

1. On the host where CCI is installed, verify that the CCI is not running. If the CCI software is still running, shut down the CCI software using the `horcmshutdown` command. For details on `horcmshutdown`, refer to the *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Command Control Interface (CCI) User and Reference Guide* (MK-95DF701).
2. In the command prompt, make two copies of the sample file (`horcm.conf`).

Example:

```
c:\HORCM\etc> copy \HORCM\etc\horcm.conf \WINNT\horcm0.conf
c:\HORCM\etc> copy \HORCM\etc\horcm.conf \WINNT\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)`. Refer to the *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Command Control Interface (CCI) User and Reference Guide* (MK-95DF701) for details on calculating the `poll(10ms)` value. Specifying the value incorrectly may cause resource contention in the internal process; the process is temporarily suspended and pauses the internal processing of the subsystem. For more details on configuration parameters, refer to the *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Command Control Interface (CCI) User and Reference Guide* (MK-95DF701).

- In the HORCM_CMD section, specify the physical drive (command device) on the subsystem.

```

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

HORCM_INST
#dev_group       ip_address    service
UG01             XXXXXXXXX    5001

```

Figure 5.21 Horcm0.conf Example (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 5.22 Horcm0.conf Example (P-VOL: S-VOL=1: 3)

- Save the configuration definition file and use the `horcmstart` command to start the CCI software. For details on `horcmstart`, refer to the *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Command Control Interface (CCI) User and Reference Guide (MK-95DF701)*.
- Execute the `raidscan` command and write down the target ID displayed in the execution result. For details on the `raidscan` command, refer to the *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Command Control Interface (CCI) User and Reference Guide (MK-95DF701)*.
- Shut down the CCI software and then open the configuration definition file again.
- 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#, and the value must be set as 0 (zero).
- In the HORCM_INST section, set the necessary parameters, and then save (overwrite) the file.
- Repeat steps 3 to 10 for the `horcm1.conf` file (Figure 5.23).

```

horcm1.conf - Notepad
File Edit Search Help
HORCM_MON
#ip_address      service      poll(10ms)    timeout(10ms)
XXXXXXXXXX      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 5.23 Horcm1.conf Example (P-VOL: S-VOL=1: 1)

```

horcm1.conf - Notepad
File Edit Search Help
HORCM_MON
#ip_address      service      poll(10ms)    timeout(10ms)
XXXXXXXXXX      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
UG02            oradb2        CL1-A         1              3             0
UG03            oradb3        CL1-A         1              4             0

HORCM_INST
#dev_group      ip_address    service
UG01            XXXXXXXXX    5000

```

Figure 5.24 Horcm1.conf Example (P-VOL: S-VOL=1: 3)

- Enter the following in the command prompt to verify the connection between CCI and the subsystem.

Example:

```

C:\>cd horcm\etc

C:\horm\etc>echo hd1-3 | .\inraid
Harddisk 1 -> [ST] CL1-A Ser =75000174 LDEV = 0 [HITACHI ] [DF600F-CM ]
Harddisk 2 -> [ST] CL1-A Ser =75000174 LDEV = 1 [HITACHI ] [DF600F ]
HORC = SMPL HOMRCF[MU#0 = SMPL MU#1 = NONE MU#2 = NONE]
RAID5[Group 1-0] SSID = 0x0000
Harddisk 3 -> [ST] CL1-A Ser =75000174 LDEV = 2 [HITACHI ] [DF600F ]
HORC = SMPL HOMRCF[MU#0 = SMPL MU#1 = NONE MU#2 = NONE]
RAID5[Group 2-0] SSID = 0x0000

C:\horm\etc>

```

For details on configuration definition file, refer to the *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Command Control Interface (CCI) User and Reference Guide (MK-95DF701)*.

5.6 Setting the Environment Variable

To perform ShadowImage operations, you must set the environment variable for the execution environment. The following describes an example supposing the system configuration with two instances within the same server (Windows NT/2000/Windows Server 2003).

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

Example:

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

2. To use ShadowImage, you must set the environment variable shown below.

Example:

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

3. Execute the horcmstart script, and then execute the pairdisplay command to verify the configuration.

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>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)75000174 1.SMPL -----,----- ---- -
VG01  oradb1(R)    (CL1-A , 1, 2-0)75000174 2.SMPL -----,----- ---- -
```

Preparing for ShadowImage operation is now complete.

Chapter 6 Performing ShadowImage Operations

This chapter provides examples of ShadowImage commands using the Windows NT/2000/Windows Server 2003 systems. The icons of the pairs, which indicate pair status in Navigator, are also displayed.

To execute ShadowImage commands, from the host where CCI is installed, display the command prompt.

This chapter contains the following:

- Creating ShadowImage Pairs (paircreate) (see section 6.1)
- Splitting ShadowImage Pairs (pairsplit) (see section 6.2)
- Resynchronizing ShadowImage Pairs (pairresync) (see section 6.3)
- Suspending ShadowImage Pairs (pairsplit -E) (see section 6.4)
- Releasing ShadowImage Pairs (pairsplit -S) (see section 6.5)
- Confirming the Status of Pairs (see section 6.6)
- Troubleshooting (see section 6.7)

6.1 Creating ShadowImage Pairs (paircreate)

To create ShadowImage pairs:

1. For example, if the group name in the configuration definition file is VG01, follow these steps:

Execute the `pairdisplay` command to verify that the status of the ShadowImage volumes is SMPL.

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 )75000174    1.SMPL -----,----- ---- -
VG01   oradb1(R)   (CL1-A , 1, 2-0 )75000174    2.SMPL -----,----- ---- -
```

2. Execute the `paircreate` command. When executing the command, always use the `-c` option with the value 15 to specify the copying pace. Select "fast" (11 to 15) for the `-c` option, which indicates the copying pace when copying pace is preceded. Select "fast" (11 to 15) when copying pace is preceded. However, when "fast" (11 to 15) is selected, host I/O performance may be lowered. Therefore, when you need to prevent performance from being lowered, for the host I/O, select "medium" or "slow" (1 to 10). Then, execute the `pairevtwait` command to verify that the status of each volume is PAIR.

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

Figure 6.1 Paircreate Command Example

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

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 )75000174    1.P-VOL PAIR,75000174    2 -
VG01   oradb1(R)   (CL1-A , 1, 2-0 )75000174    2.S-VOL PAIR,-----    1 -
```

To assure that the data of two or more S-VOLs included in a group are of the same time, the CTG is used. The method of creating a pair using the CTG is explained below.

1. For example, if the group name in the configuration definition file is VG01, follow these steps:

Execute the `pairdisplay` command to verify that the status of the ShadowImage volumes is SMPL.

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 )75000174    1.SMPL -----,----- ---- -
VG01   oradb1(R)   (CL1-A , 1, 2-0 )75000174    2.SMPL -----,----- ---- -
VG01   oradb2(L)   (CL1-A , 1, 3-0 )75000174    3.SMPL -----,----- ---- -
VG01   oradb2(R)   (CL1-A , 1, 4-0 )75000174    4.SMPL -----,----- ---- -
VG01   oradb3(L)   (CL1-A , 1, 5-0 )75000174    5.SMPL -----,----- ---- -
VG01   oradb3(R)   (CL1-A , 1, 6-0 )75000174    6.SMPL -----,----- ---- -
```

- 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.
```

Figure 6.2 Paircreate -m grp Command Example

- Execute the `pairdisplay` command to verify the pair status and the configuration.

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 )75000174    1.P-VOL PAIR,75000174    2 -
VG01   oradb1(R)    (CL1-A , 1, 2-0 )75000174    2.S-VOL PAIR,-----    1 -
VG01   oradb2(L)    (CL1-A , 1, 3-0 )75000174    3.P-VOL PAIR,75000174    4 -
VG01   oradb2(R)    (CL1-A , 1, 4-0 )75000174    4.S-VOL PAIR,-----    3 -
VG01   oradb3(L)    (CL1-A , 1, 5-0 )75000174    5.P-VOL PAIR,75000174    6 -
VG01   oradb3(R)    (CL1-A , 1, 6-0 )75000174    6.S-VOL PAIR,-----    5 -
```

The ShadowImage pair is created. For details on the `paircreate` command and its options, refer to the *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Command Control Interface (CCI) User and Reference Guide (MK-95DF701)*.

6.2 Splitting ShadowImage Pairs (pairsplit)

To split ShadowImage pairs:

- For example, if the group name in the configuration definition file is VG01:

Execute the `pairsplit` command to split the ShadowImage pair in the PAIR status.

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

Figure 6.3 Pairsplit Command Example

- Execute the `pairdisplay` command to verify the pair status and the configuration.

```
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 )75000174    1.P-VOL PSUS,75000174    2 -
VG01   oradb1(R)    (CL1-A , 1, 2-0 )75000174    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.

The ShadowImage pair is split. For details on the `pairsplit` command and its options, refer to the *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Command Control Interface (CCI) User and Reference Guide (MK-95DF701)*.

6.3 Resynchronizing ShadowImage Pairs (pairresync)

To resynchronize ShadowImage pairs:

1. For example, if the group name in the configuration definition file is VG01:

Execute the `pairresync` command to resynchronize the ShadowImage pair. When executing the command, select "fast" (11 to 15) for `-c` option; this indicates the copying pace when the copying pace is preceded. Select "fast" (11 to 15) when the copying pace is preceded. However, when "fast" (11 to 15) is selected, the host I/O performance may be lowered. Therefore, when you need to prevent a lowered performance of the host I/O, select "medium" or "slow" (1 to 10), then execute the `pairevtwait` command to verify that the status of each volume is PAIR.

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

Figure 6.4 Pairresync Command Example

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

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 )75000174      1.P-VOL PAIR,75000174      2  -
VG01   oradb1(R)   (CL1-A , 1, 2-0 )75000174      2.S-VOL PAIR,-----      1  -
```

The ShadowImage pair is resynchronized. For details on the `pairresync` command and its options, refer to the *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Command Control Interface (CCI) User and Reference Guide (MK-95DF701)*.

6.4 Suspending ShadowImage Pairs (pairsplit -E)

To suspend the ShadowImage pairs:

1. For example, if the group name in the configuration definition file is VG01:

Execute the `pairdisplay` command to verify that the pair to be suspended is in PAIR status.

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 )75000174      1.P-VOL PAIR,75000174      2  -
VG01   oradb1(R)   (CL1-A , 1, 2-0 )75000174      2.S-VOL PAIR,-----      1  -
```

2. Execute the `pairsplit (pairsplit -E)` command to suspend under error the ShadowImage pair.

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

Figure 6.5 Pairsplit -E Command Example

- Execute the `pairdisplay` command to verify that the pair status changed to PSUE.

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 )75000174      1.P-VOL PSUE,75000174      2 -
VG01    oradb1(R)    (CL1-A , 1, 2-0 )75000174      2.S-VOL PSUE,-----      1 -
```

The ShadowImage pair is split. For details on the `pairsplit` command and the `-E` option, refer to the *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Command Control Interface (CCI) User and Reference Guide* (MK-95DF701).

6.5 Releasing ShadowImage Pairs (pairsplit -S)

To release the ShadowImage pair and change the status to SMPL:

1. For example, if the group name in the configuration definition file is VG01:

Execute the `pairdisplay` command to verify that the status of the ShadowImage pair is PAIR.

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 )75000174    1.P-VOL PAIR,75000174    2  -
VG01   oradb1(R)  (CL1-A , 1, 2-0 )75000174    2.S-VOL PAIR,-----    1  -
```

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

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

Figure 6.6 Pairsplit -S Command Example

3. Execute the `pairdisplay` command to verify that the pair status changed to SMPL.

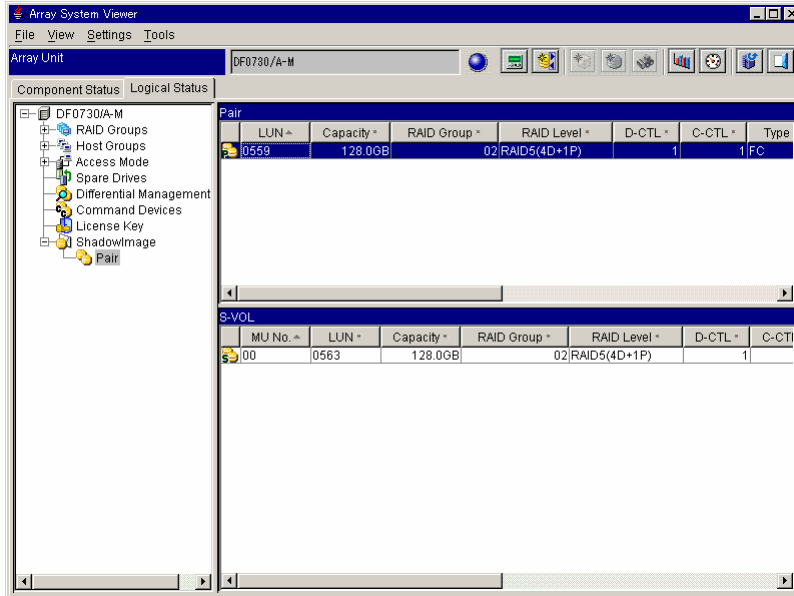
Example:

```
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 )75000174    1.SMPL -----,-----    ---- -
VG01   oradb1(R)  (CL1-A , 1, 2-0 )75000174    2.SMPL -----,-----    ---- -
```

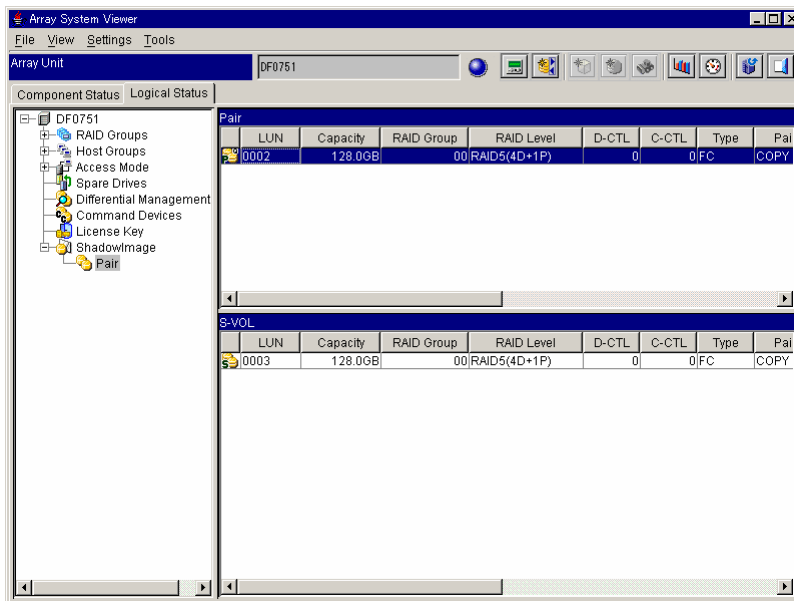
The ShadowImage pair is released. For details on the `pairsplit` command and the `-S` option, refer to the *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Command Control Interface (CCI) User and Reference Guide* (MK-95DF701).

6.6 Confirming the Status of Pairs

The status of the logical units is displayed in icon format in the **Logical Status** dialog box of the **Array System Viewer** window. **P** is used to designate a primary volume and **S** is used to designate a secondary volume.



Navigator: Version 5.00 or later



Navigator: Less than 5.00 version

Figure 6.7 Array System Viewer Window (Pair Display)

Table 6.1 Icons for Pair Status

Volume Status				Paired Status	Meaning
Formatted Yellow	Unformatted Gray	Regressed Pink	Blockaded Red		
				PAIR	Indicates the paired volume.
				COPY	Indicates a paired status, but initial copy or resync operation is not complete.
				PSUS	In paired status, but updates to the S-VOL data are suspended due to user-requested pairsplit.
				PSUE	In paired status, but updates to the S-VOL data are suspended forcibly or due to an error condition.

6.7 Troubleshooting

When a software error occurs while operating ShadowImage, user operation (CCI operation by the user), is necessary in addition to system maintenance performed by Hitachi Data Systems personnel.

For example, when formatting is necessary due to an LU error and that LU is used for ShadowImage, the pair must be released by the user (CCI operation by the user) before formatting the LU. Therefore, contact Hitachi Data Systems personnel because maintenance requires the user to issue CCI commands. Note that Hitachi Data Systems personnel can only remove the error caused by the hardware. Operations such as recovering a ShadowImage pair status (e.g. resynchronizing) must be done by the user.

If you have a problem with the ShadowImage operation, refer to the troubleshooting section in the *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Command Control Interface (CCI) User and Reference Guide* (MK-95DF701). ShadowImage error messages and procedures for recovering the PSUE errors are described in detail.

Chapter 7 ShadowImage Basic Specifications

This chapter includes the following:

- External Specifications (see section 7.1)
- Functions (see section 7.2)
- Resync Operation and the Differential Data Copy Process (see section 7.3)
- Concurrent Use of LUN Expansion (see section 7.4)
- Cascade Connection (see section 7.5)

7.1 Supported Capacities and Configurations

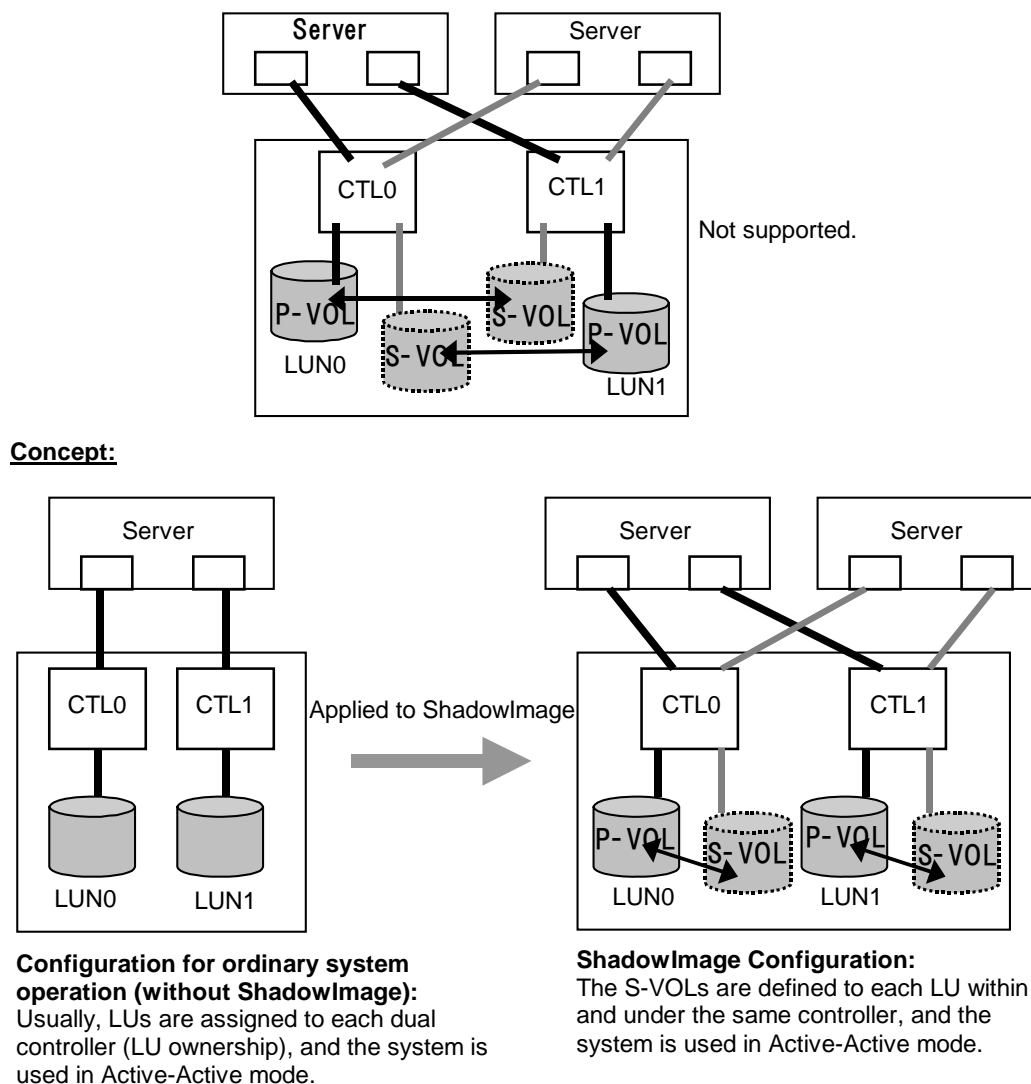


Figure 7.1 Conceptual Configuration

The configuration shown in Figure 7.1 (the P-VOL and S-VOL arranged to go across two different CTLs) is not supported.


Note: Mixing LUs (P-VOL and S-VOL) of ShadowImage and LUs of non-ShadowImage are available within the TagmaStore subsystem. However, note that there may be some effects to the performance. The performance decreases when the **resync** mode is in priority during resynchronization (even if the LUs are non-ShadowImage). For details, see section 8.4.

The maximum capacity supported by ShadowImage can be calculated from the following formula. The single maximum capacity of ShadowImage is shown in Table 7.1.

ShadowImage: Maximum supported capacity value of S-VOL (TB)

$$\begin{aligned}
 &= (\text{Maximum ShadowImage single capacity} \\
 &\quad - (\text{Total P-VOL and S-VOL capacity of TrueCopy}) \\
 &\quad - 17 \times (\text{Total P-VOL capacity of SnapShot}) \\
 &\quad - 51 \times (\text{Total P-VOL and S-VOL capacity of TCE})) \div 2
 \end{aligned}$$

Note 5: ShadowImage supported capacity is calculated not based on the P-VOL capacity but based on the S-VOL capacity only. The total sum of the P-VOL and S-VOL capacities varies depending on whether the pair configuration (correspondence between the P-VOL and S-VOL) is one-to-one or not. An example of the pair configuration, which can be constructed when the maximum S-VOL capacity that is supported is 3 TB, is shown below.



① Pair configuration is one-to-one:
The sum of the P-VOL and S-VOL capacities is 6 TB.

② Pair configuration is one-to-three:
The sum of the P-VOL and S-VOL capacities is 4 TB.

Note 6: Because TrueCopy and TCE cannot be used together, the capacity supported by TrueCopy or TCE in the above formula is zero.

Table 7.1 Single Maximum Supported Capacity of ShadowImage (TB)

Equipment Type	Mounted Memory Capacity	Single Maximum Supported Capacity (TB)
WMS100	512 MB/CTL	39
	1 GB/CTL	180
AMS200	1 GB/CTL	110
	2 GB/CTL	350
AMS500	1 GB/CTL	110
	2 GB/CTL	430
	3 GB/CTL	430
	4 GB/CTL	770
AMS1000	2 GB/CTL	750
	4 GB/CTL	1,100
	6 GB/CTL	1,340
	8 GB/CTL	1,610

7.2 Functions

Note: When ShadowImage is used in NAS, do not perform a pair creation, resynchronization, or reverse resynchronization during operations such as installation, boot, shutdown, node stop, and cluster stop of the NAS OS. Do not perform operations such as installation, boot, shutdown, node stop, and cluster stop of the NAS OS during pair creation, resynchronization, or reverse resynchronization of ShadowImage. If the pair creation, resynchronization, or reverse resynchronization of ShadowImage and one of the above operations of the NAS OS are performed at the same time, operation to the NAS OS may fail due to overload.

Table 7.2 lists and describes ShadowImage functions.

Table 7.2 ShadowImage Functions

No.	Function	Contents
1	Pair Creation	Supported For P-VOL: Read/Write For S-VOL: Read only
2	Pair Splitting	Supported For P-VOL: Read/Write For S-VOL: Read/Write
3	Resynchronization (resync) in the Direction of P-VOL to S-VOL	Supported For P-VOL: Read/Write For S-VOL: Read only
4	Resynchronization (reverse resync) in the Direction of S-VOL to P-VOL	Supported For P-VOL: Read from S-VOL/Write For S-VOL: Read only See explanation below for details. (Figure 5.3)
5	Pair Suspension (Note)	Supported For P-VOL: Read/Write For S-VOL: Read only
6	Pair Release	Supported
7	Splitting	Supported Immediate split is possible (data is also fixed immediately).

Note: Read/Write command cannot be executed for P-VOL and for S-VOL when the pair is suspended during resynchronization.

7.2.1 Pair Creation

The ShadowImage paircreate operation establishes the newly specified ShadowImage pair(s). The volumes that will become the P-VOL and S-VOL must be in the SMPL (simplex) state before being joined as a ShadowImage pair. Also, when the P-VOL that will be an operation target configures a ShadowImage pair in the PAIR status and the COPY status or the PSUE status with another S-VOL, a pair cannot be newly created using the paircreate command.

The ShadowImage initial copy operation takes place when you create a new ShadowImage pair. 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 COPY(PD) (PD = pending) while the initial copy operation is in progress. The pair status changes to PAIR when the initial copy is complete.




Figure 7.2 Example of Pair Creation

7.2.2 Pair Splitting

The ShadowImage pairsplit operation performs all pending S-VOL updates (those issued prior to the split command and recorded in the P-VOL track map) to make the S-VOL identical to the state of the P-VOL when the split command was issued, and then provides full read/write access to the split S-VOL. You can split existing pairs as needed, and you can also use the pairsplit operation to create and split pairs in one step. 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.

7.2.3 Resynchronization (resync) from P-VOL to S-VOL

The ShadowImage pairresync operation resynchronizes the split pairs (PSUS) or the suspended pairs (PSUE). When the pairresync operation starts, the pair status changes to COPY(RS) or COPY(RS-R). The pair status changes to PAIR when the pairresync operation completes.

The normal pairresync operation resynchronizes the S-VOL with the P-VOL. The copy direction for a normal pairresync operation is P-VOL to S-VOL. The pair status during a normal resync operation is COPY(RS). 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 pairresync. The normal pairresync operation can be executed for pairs with the status PSUS and PSUE. However, when the P-VOL that will be an operation target configures a ShadowImage pair in the PAIR status and the COPY status or the PSUE status with another S-VOL, the pairresync operation cannot be executed. When the pairresync operation is executed for a pair in the PSUE (S-VOL Switch) status, the reverse resynchronization operation is performed.

7.2.4 Resynchronization (reverse resync) from S-VOL to P-VOL

The reverse pairresync (pairresync -restore) operation synchronizes the P-VOL with the S-VOL. This function enables Read and Write access to P-VOL while resynchronizing in the direction of S-VOL to P-VOL (the Read access to P-VOL is performed from S-VOL). Therefore, reconstructing P-VOL data (due to a failure) from S-VOL can be performed and recovered quickly (shorten the time for stopping the host applications to P-VOL) without activating the P-VOL/S-VOL switch function. However, when the P-VOL that will be an operation target configures a ShadowImage pair in the PAIR status and the COPY status or the PSUE status with another S-VOL, the pairresync -restore operation cannot be executed.

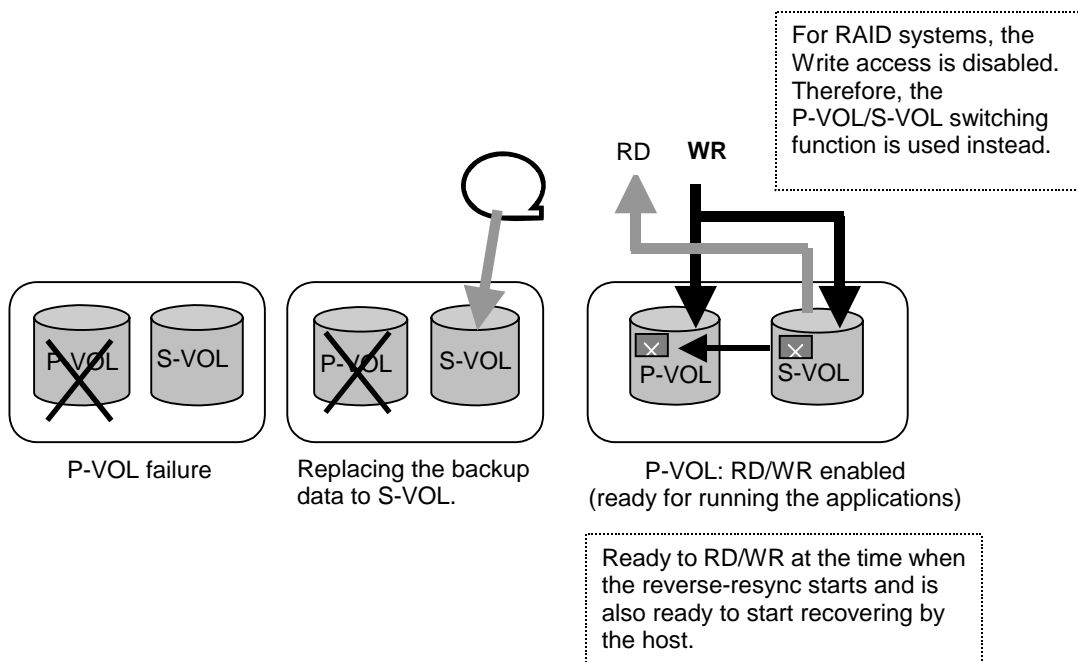


Figure 7.3 Resynchronization (reverse resync) from S-VOL to P-VOL

7.2.5 Pair Suspension

The ShadowImage pairsplit-E operation suspends the ShadowImage copy operations to the S-VOL of the pair. A ShadowImage pair can be suspended by the user at any time. When a ShadowImage pair is suspended (status = PSUE), the subsystem 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 pairresync operation is performed on a suspended pair, the entire P-VOL is copied to the S VOL (when a pairresync -restore operation is performed, the entire S-VOL is copied to the P-VOL). While the pairresync operation for a split ShadowImage pair can be very fast, the pairresync operation for a suspended pair will take as long as the initial copy operation.

The TagmaStore subsystem will automatically suspend a ShadowImage pair when it cannot keep the pair mirrored for any reason. When the TagmaStore suspends a pair, a file is output to the system log or event log to notify the host. The TagmaStore will automatically suspend a pair under the following conditions:

- When the ShadowImage volume pair has been suspended or deleted from the UNIX/PC server host using the Hitachi CCI.
- When the TagmaStore detects an error condition related to an initial copy operation. When a volume pair with COPY(PD) status is suspended, the TagmaStore aborts the initial copy operation, changes the status of the P-VOL and S-VOL to PSUE and accepts all subsequent write I/Os to the P-VOL.

7.2.6 Pair Release

The ShadowImage pairsplit-S operation stops the ShadowImage copy operations to the S-VOL of the pair and releases the volume in paired status. A ShadowImage pair can be released by the user at any time except during the pairsplit operation (i.e., any status except SMPL). Volumes status changes to SMPL for both.

7.3 Resync Operation and Differential Data Copy Process

When a pair is split in ShadowImage, the P-VOL and S-VOL can be updated individually. It is different than coupling; updates on P-VOL will not be reflected to S-VOL while it is in suspension. Therefore, the user must synchronize the data again to match the P-VOL and S-VOL data after the suspend operation.

7.3.1 Managing Differential Data

When the pair is suspended, the status of each volume (P-VOL and S-VOL) changes to be managed individually. To reunite the pairs, the user must resynchronize the volumes in order to match the data. However, regardless of the amount of updates, copying the entire volume takes time. Therefore, ShadowImage manages (records) the updated data for each volume (P-VOL and S-VOL) and copies only the differential data (called “differential data management”) to help shorten the time for resynchronization. When one P-VOL configures a pair with a maximum of three S-VOLs, the difference is managed for each pair.

7.3.2 Differential Data

The ShadowImage records both the P-VOL and S-VOL host Write positions on a bitmap and refers to these bitmaps when resynchronizing the volumes. ShadowImage copies only the differential data, so that the contents of the P-VOL and S-VOL match.

The range for a single bit within the bitmap is equivalent to 1 MB, therefore:

- A 1 KB update requires a 1 MB copy when resynchronizing the volume.
- The amount of data to copy becomes less when the access pattern from the host is high.

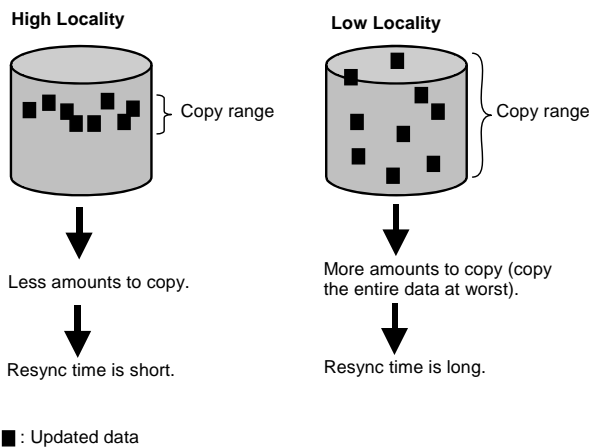


Figure 7.4 Example of Differential Data

Note: The ShadowImage performance evaluation described in Chapter 9 is measured in the status of low locality (entire copy).

7.4 Concurrent Use of LUN Expansion

In ShadowImage, a unified LU can be used as a P-VOL or an S-VOL.

Preconditions and restrictions of concurrent use with ShadowImage and LUN Expansion are as follows:

- S-VOL and P-VOL Capacities:
 - Capacities of the P-VOL and S-VOL must be the same. The P-VOL and S-VOL can be paired as long as their capacities are the same, even if numbers of LUs composing them (unified LUs) are different from each other.
 - A unified LU composed of 17 or more LUs cannot be assigned to the P-VOL.




Figure 7.5 Example: Unified LU Assigned to P-VOL or S-VOL




Figure 7.6 Example: Restriction on a Number of LUs Composing a Unified LU

- P-VOL and S-VOL RAID Levels:
 - All the LUs including the unified LU(s) assigned to the P-VOL and S-VOL must not be on the same RAID level and have the same number of data disks.




Figure 7.7 Example: Combination of RAID Levels




Figure 7.8 Example: Another Combination of RAID Levels Allowed

- Others:
 - A P-VOL or an S-VOL belonging to a ShadowImage pair cannot be used to compose a unified LU. Unify it after splitting the pair once.
 - An LU consisting of the FC drive and the SATA drive cannot be unified. Unify LUs consisting of drives of the same type.




Figure 7.9 Example: Restriction on Unification of LUs Consisting of Different Types of Drives

7.5 Cascade Connection

Because TrueCopy accepts the write operation during a copy operation, you can build a system without a cascade connection with ShadowImage. However, you can create a system with ShadowImage. Some examples of the structure are shown in this section.

- Cascade with a ShadowImage P-VOL:

When both TrueCopy and ShadowImage pairs are placed in PAIR status, host performance on the local side is lowered. Therefore, set a PSUS status for TrueCopy and ShadowImage pairs when host I/Os are frequently instructed.

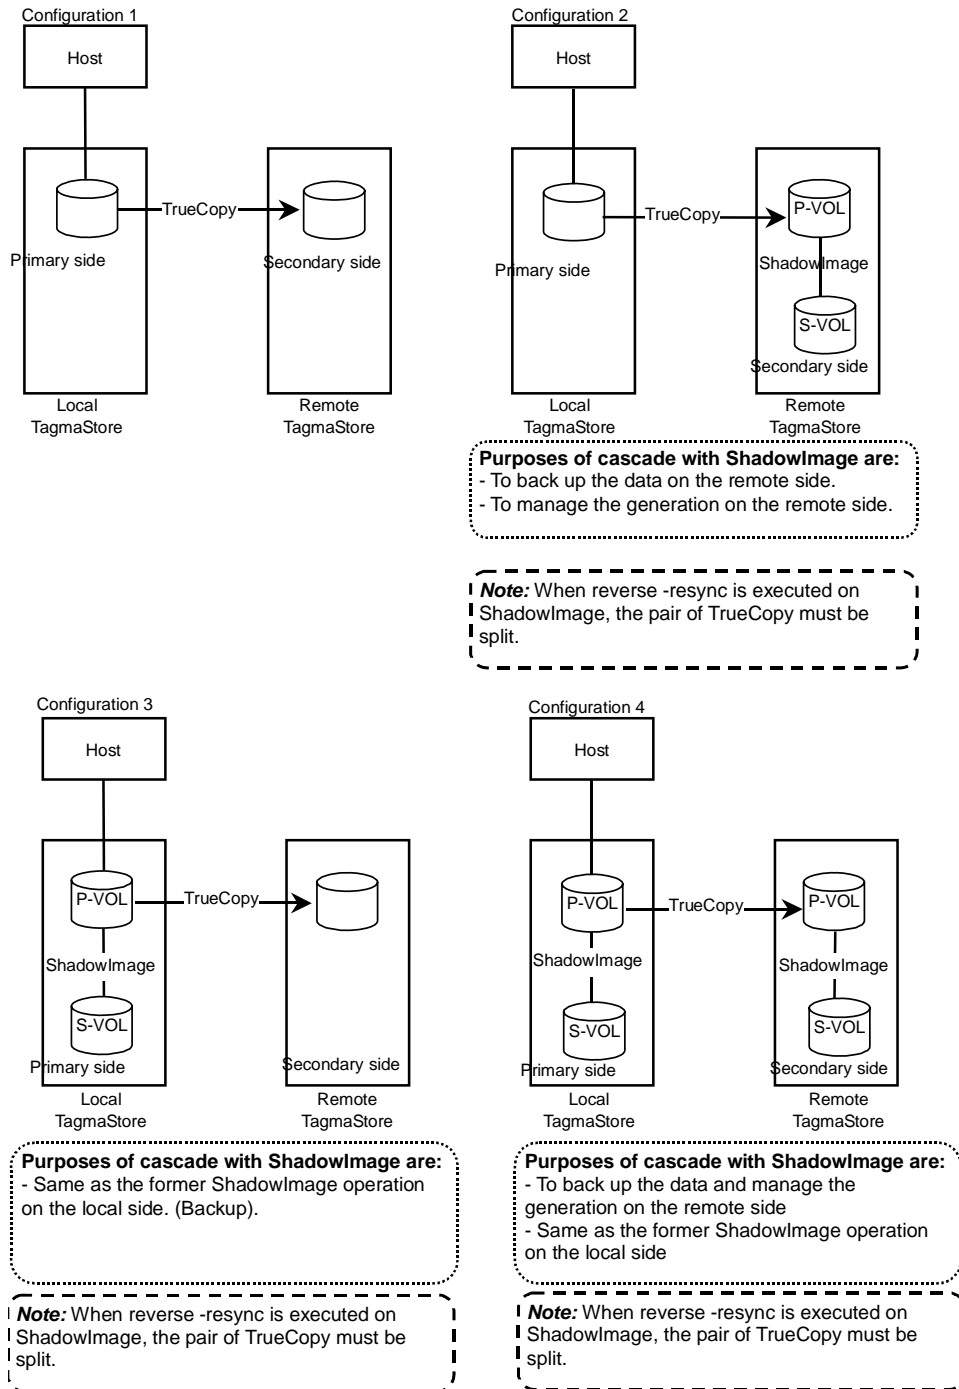


Figure 7.10 Cascade Structure: ShadowImage P-VOL (P-VOL: S-VOL=1: 1)

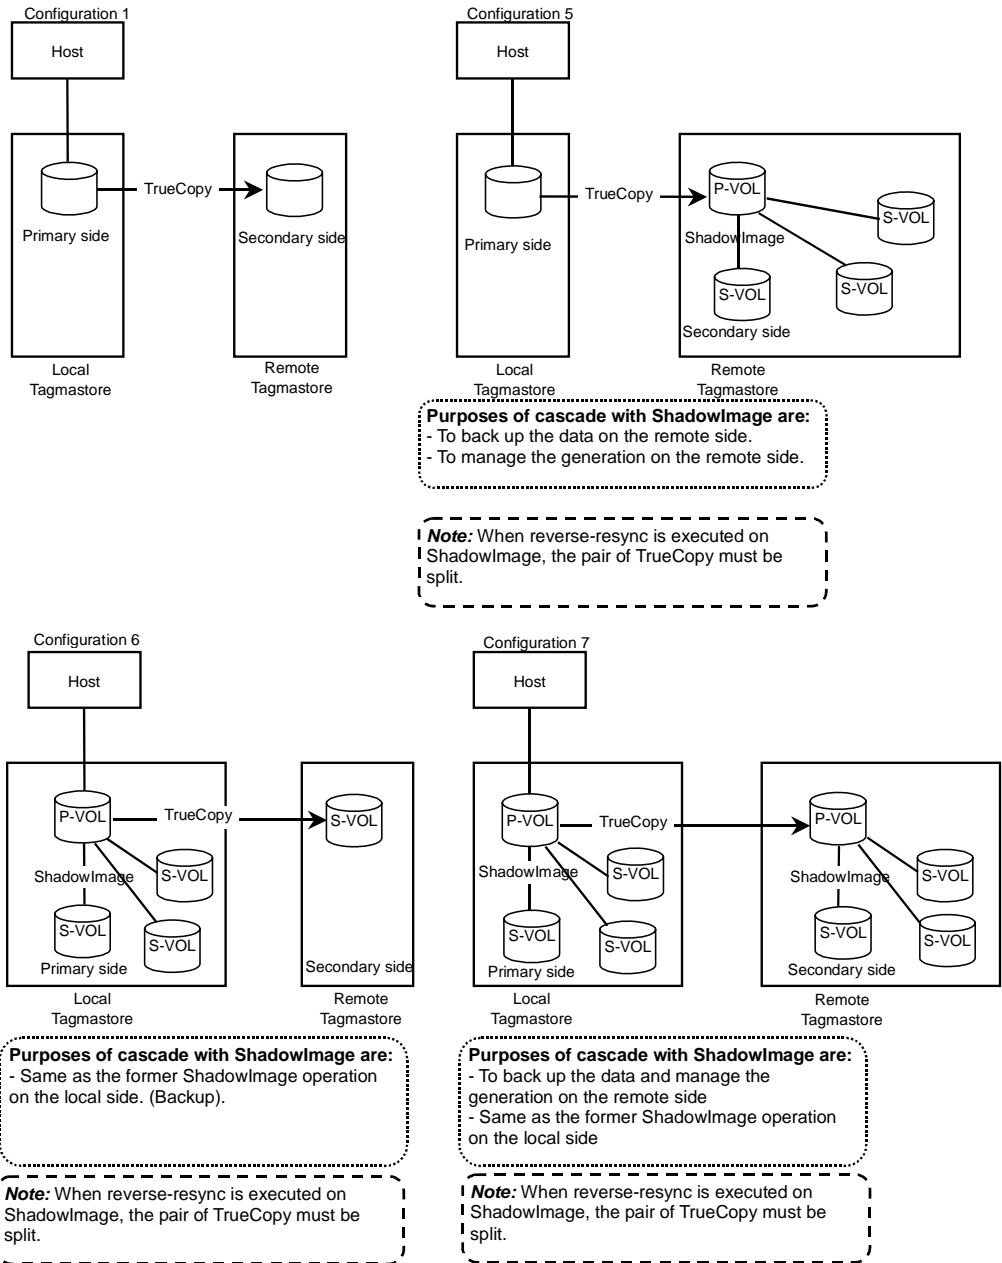


Figure 7.11 Cascade Structure: ShadowImage P-VOL (P-VOL: S-VOL=1: 3)

- Cascade with a ShadowImage S-VOL:

A cascade of a ShadowImage S-VOL is used when asynchronously performing a backup on the remote side. Because a backup is made from a ShadowImage S-VOL to the remote side in this configuration, lowered performance on the local side (a P-VOL of ShadowImage) during the backup can be minimized.

Note: A TrueCopy pair must have PSUS status and/or Hi-Copy must be placed in PSUS status when resynchronizing a volume of ShadowImage on the local side.

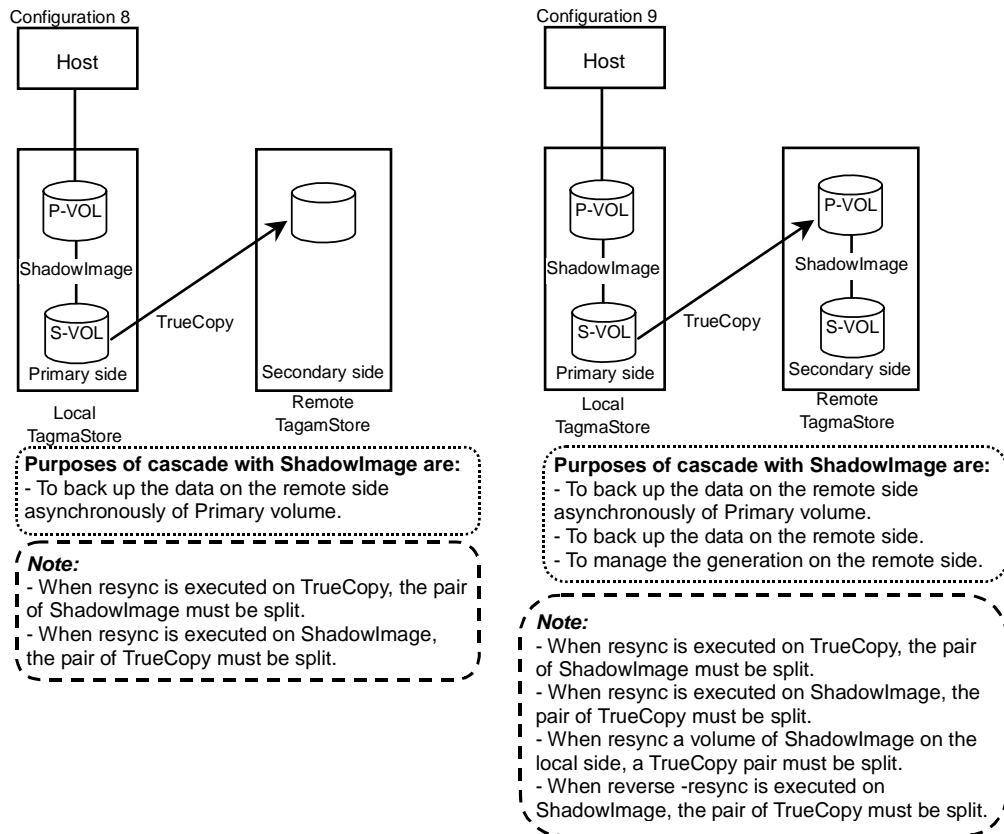


Figure 7.12 Cascade Structure: ShadowImage S-VOL (P-VOL: S-VOL=1: 1)

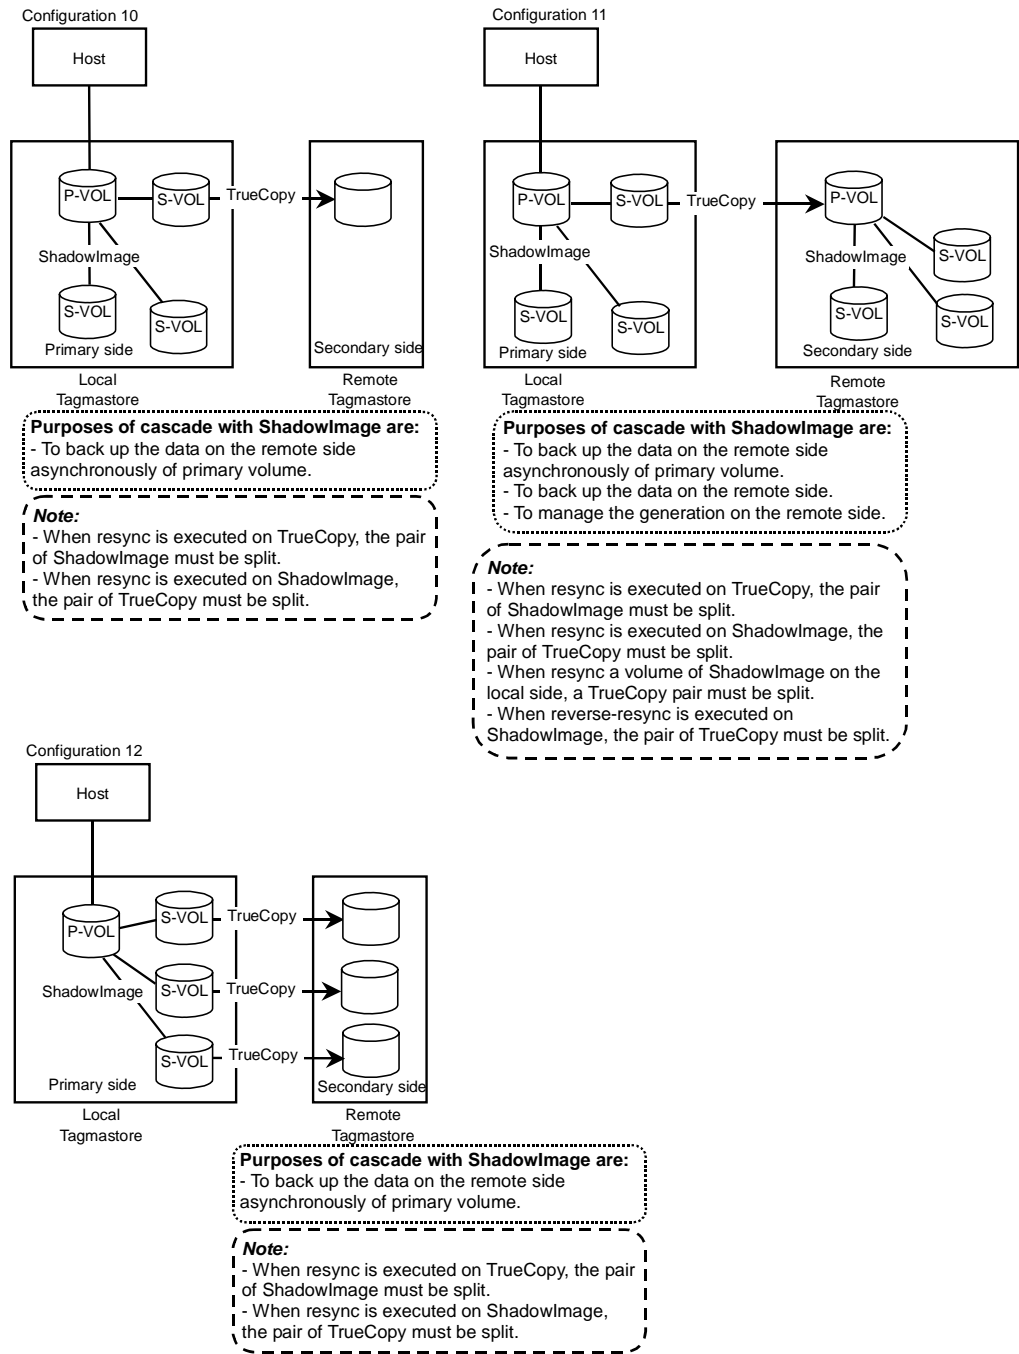
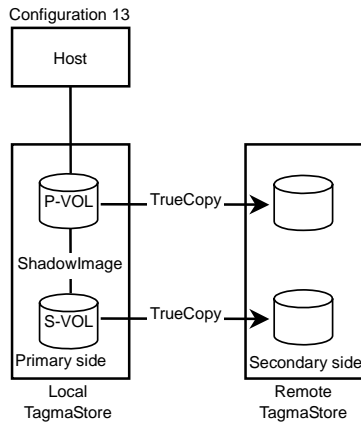
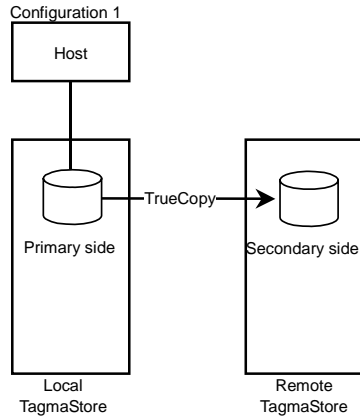


Figure 7.13 Cascade Structure: ShadowImage S-VOL (P-VOL: S-VOL=1: 3)

- Simultaneous Cascade with a ShadowImage P-VOL and S-VOL: (Simultaneous configuration is supported when the micro program version is more than 0772/A)

When you combine the ShadowImage P-VOL and S-VOL cascades, you can use TrueCopy to copy the S-VOL backup data acquired by ShadowImage to the remote side. This can be done during a TrueCopy operation. Note: Simultaneous configuration is supported by microprogram version 0772/A and later.

The operating conditions for both P-VOL and S-VOL cascade configurations must be met when using a ShadowImage pair during a TrueCopy operation. The configuration of the cascaded volume is applied during TrueCopy operations.

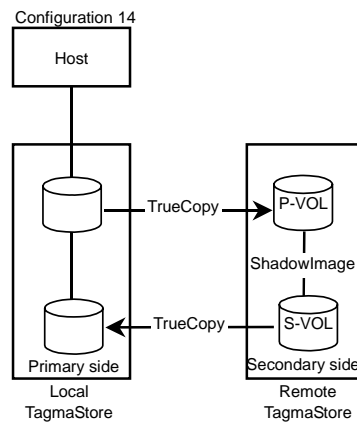


Purposes of cascade with ShadowImage are:

- Same as the former ShadowImage operation on the local side. (Backup).
- The backup data acquired on the primary side is copied to the secondary side.

Note:

- When performing the reverse resynchronization of ShadowImage, it is required to split the TrueCopy pairs cascaded with the P-VOL and the S-VOL, respectively.
- When performing the resynchronization of ShadowImage, it is required to split the TrueCopy pairs cascaded with the P-VOL and the S-VOL.
- When performing the resynchronization of TrueCopy pairs cascaded with the S-VOL, it is required to split the ShadowImage pairs.



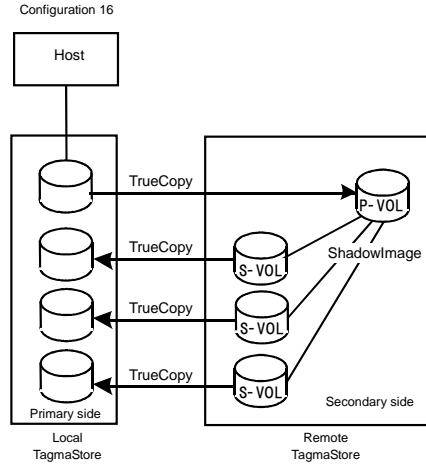
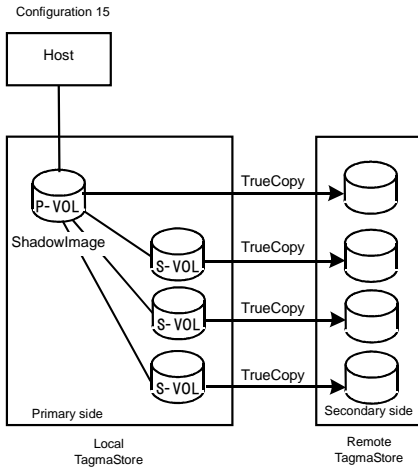
Purposes of cascade with ShadowImage are:

- A configuration resulted from the resynchronization (swapping) of the TrueCopy pair cascaded with the P-VOL in the configuration 13.

Note:

- When performing the reverse resynchronization of ShadowImage, it is required to split the TrueCopy pairs cascaded with the P-VOL and the S-VOL, respectively.
- When performing the resynchronization of ShadowImage, it is required to split the TrueCopy pairs cascaded with the P-VOL and the S-VOL.
- When performing the resynchronization of TrueCopy pairs cascaded with the S-VOL, it is required to split the ShadowImage pairs.

Figure 7.14 Simultaneous Cascading with a ShadowImage P-VOL and S-VOL (P-VOL: S-VOL=1:1)



Purposes of cascade with ShadowImage are:

- Same as the former ShadowImage operation on the local side. (Backup).
- The backup data acquired on the primary side is copied to the secondary side.

Note:

- When performing the reverse resynchronization of ShadowImage, it is required to split the TrueCopy pairs cascaded with the P-VOL and the S-VOL, respectively.
- When performing the resynchronization of ShadowImage, it is required to split the TrueCopy pairs cascaded with the P-VOL and the S-VOL.
- When performing the resynchronization of TrueCopy pairs cascaded with the S-VOL, it is required to split the ShadowImage pairs.

Purposes of cascade with ShadowImage are:

- A configuration resulted from the resynchronization (swapping) of the TrueCopy pair cascaded with the P-VOL in the configuration 15.

Note:

- When performing the reverse resynchronization of ShadowImage, it is required to split the TrueCopy pairs cascaded with the P-VOL and the S-VOL, respectively.
- When performing the resynchronization of ShadowImage, it is required to split the TrueCopy pairs cascaded with the P-VOL and the S-VOL.
- When performing the resynchronization of TrueCopy pairs cascaded with the S-VOL, it is required to split the ShadowImage pairs.

Figure 7.15 Simultaneous Cascading with a ShadowImage P-VOL and S-VOL (P-VOL: S-VOL=1:3)

7.5.1 Cascade Restrictions

- Cascade Restrictions with ShadowImage:

Cascade connection with ShadowImage is not available. Three S-VOLs can be made for a P-VOL (micro program less than 0750/A version: only one S-VOL can be made for a P-VOL). When one P-VOL configures a pair with the maximum of three S-VOLs, only one S-VOL can be in the PAIR status or the COPY status.

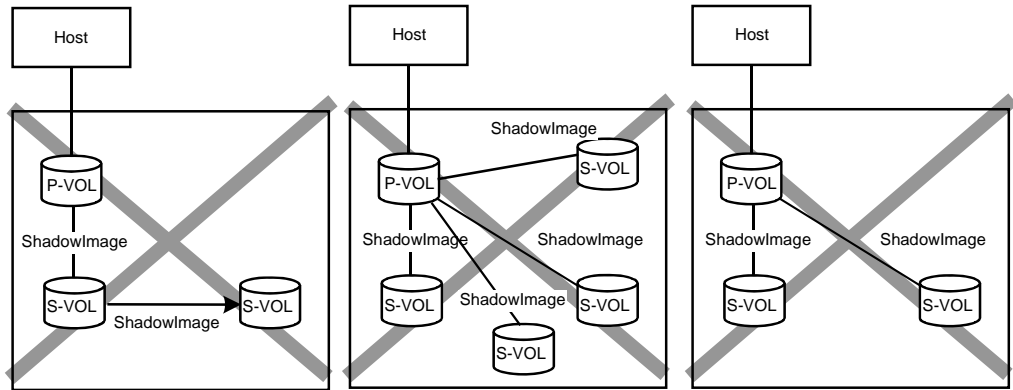


Figure 7.16 Restrictions in Cascade Connection-1

- Cascade Restrictions with SnapShot:

ShadowImage cannot be cascaded with SnapShot.

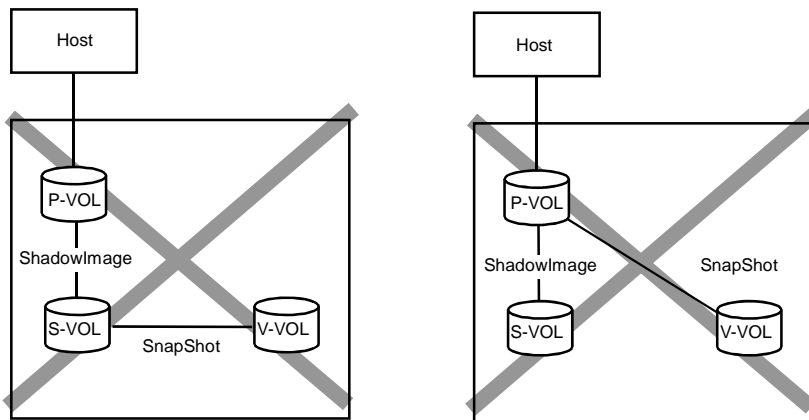


Figure 7.17 Restrictions in Cascade Connection-2

- Cascade Restrictions with POOL of SnapShot:

A ShadowImage pair cannot be created using a POOL.

- Cascade Restrictions with TrueCopy:

The following is an example of a configuration in which restrictions are placed on cascading TrueCopy with ShadowImage.

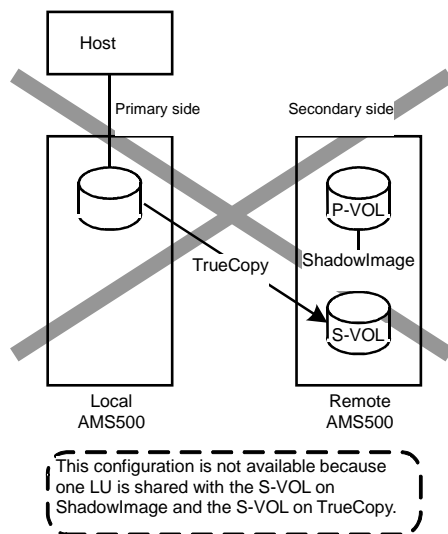


Figure 7.18 Restrictions in Cascade Connection-3

When a ShadowImage P-VOL connects a TrueCopy LU and the reverseresync operation is executed on ShadowImage, a TrueCopy pair must be split (PSUS).

Chapter 8 ShadowImage System Configuration

This section includes the following:

- Path Connection (see section 8.1)
- Locating P-VOLs and S-VOLs (see section 8.2)
- P-VOLs and S-VOLs in a RAID Configuration (see section 8.3)
- Others (Example: Mode Setting) (see section 8.4)
- Recommendations (see section 8.5)

8.1 Path Connection

8.1.1 AMS500

Configuration A is not recommended (the S-VOL is run by the other system's CTL). For the host side, use two ports per CTL (a single port per CTL causes an FC bottleneck on the host side and disables the process to perform).

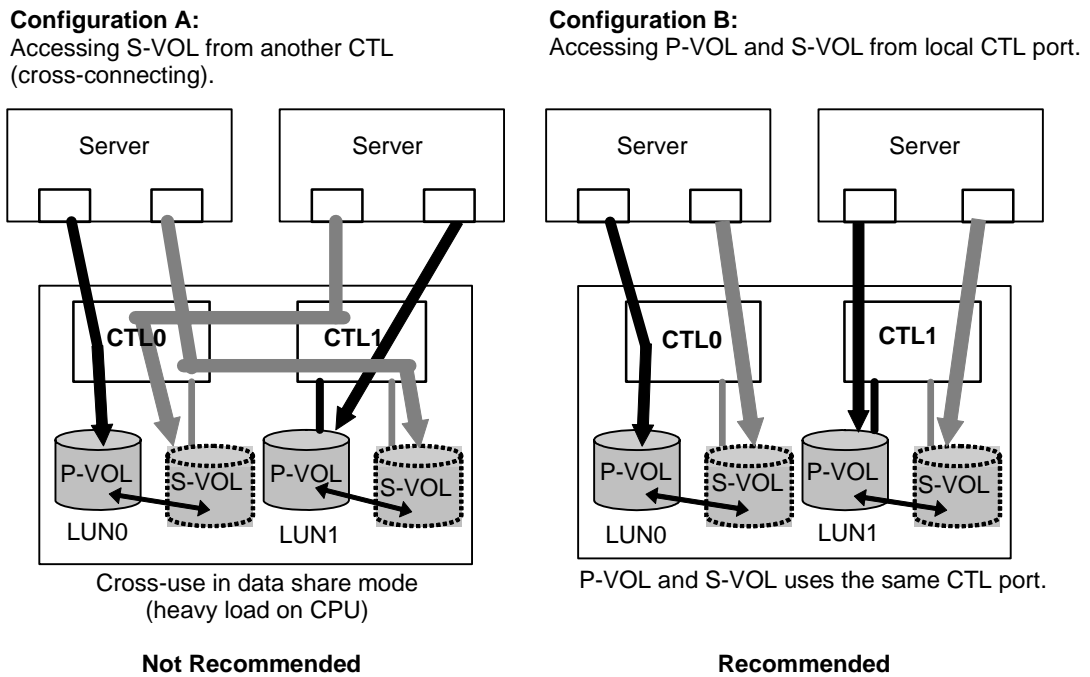


Figure 8.1 Path Connection (AMS500)

8.2 Locating P-VOLs and S-VOLs

This section includes the following:

- Location recommendations for P-VOLs and S-VOLs
- Locating multiple LUs within the same drive column
- Pair status differences when setting multiple pairs
- RAID group drive types on which a P-VOL/SVOL is located
- Setting multiple command devices

8.2.1 Location Recommendations for P-VOLs and S-VOLs

DO NOT locate P-VOLs and S-VOLs within the same ECC group of the same RAID group:

- 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.

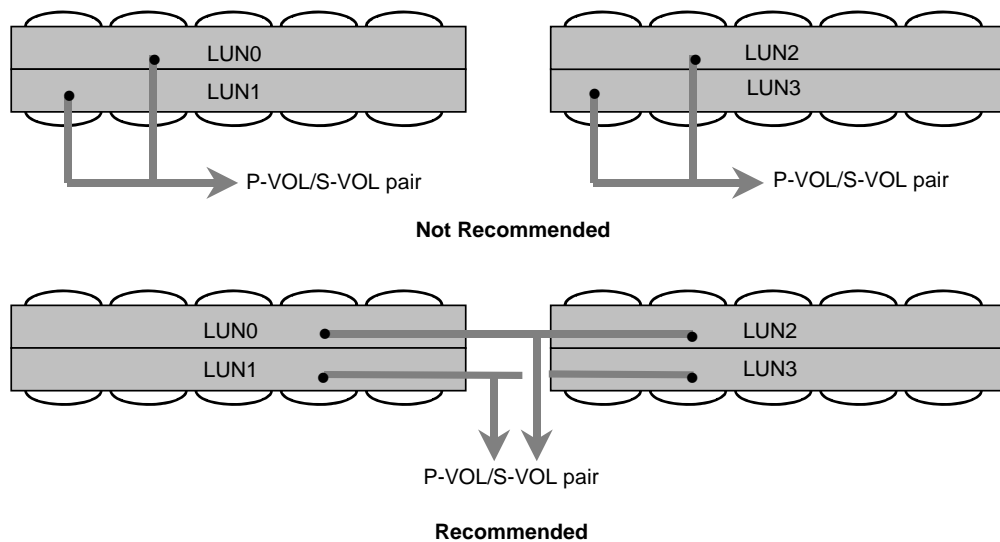


Figure 8.2 Locations for P-VOLs and S-VOLs (Not Recommended and Recommended)

8.2.2 Locating Multiple LUs within Same Drive Column

If multiple LUs are set within the same parity group and if each pair state differs, it is difficult to estimate the performance and to design the system operational settings. An example is as follows: the LU0 and LU2 are both P-VOLs and exist within the same group in the same drive (S-VOLs are located in a different parity group) and when the LU0 is in a coupling process and the LU2 is in resync process.




Figure 8.3 Locating Multiple LUs within Same Drive Column

8.2.3 Pair Status Differences when Setting Multiple Pairs

Even with a single LU per parity group, it is recommended that you keep the status of pairs the same (such as SMPL, coupling, and split) when setting multiple ShadowImage pairs. If each ShadowImage pair status differs, it is difficult to estimate the performance when designing the system operational settings. For details on ShadowImage status, see the *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Command Control Interface (CCI) User and Reference Guide (MK-95DF701)*.

8.2.4 P-VOLs and S-VOLs in a RAID Configuration

FC drive performance exceeds SATA drive performance; therefore, when a P-VOL or S-VOL is located in a RAID group consisting of SATA drives performance is lower than when a P-VOL or S-VOL is located in a RAID group consisting of FC drives. The following are recommended:

- Locate a P-VOL in a RAID group consisting of FC drives.
- When locating an S-VOL in a RAID group consisting of the SATA drives, conduct a thorough investigation beforehand.

8.2.5 Setting Multiple Command Devices

When two command devices are set within the one disk subsystem, assign them to the respective RAID groups. If they are assigned to the same RAID group, both command devices become unavailable due to difficulties, such as a drive failure.

8.3 P-VOLS and S-VOLs in a RAID Configuration

Note the following about P-VOLs and S-VOLs in RAID configuration:

- For P-VOL: RAID level with redundancy is recommended.
- For S-VOL: RAID level with redundancy is recommended.
- P-VOL redundancy => S-VOL redundancy

Table 8.1 P-VOLs and S-VOLs in a RAID Configuration

No.	P-VOL	S-VOL	Amount of User Data	Total Amount of ShadowImage	Share of User Data
1	RAID 1+0 ($N = 1$ to 8)	RAID 1+0 ($N = 1$ to 8)	1	4	1/4
2	RAID 1+0 ($N = 1$ to 8)	RAID 5 ($N = 1$ to 8)	1	$2 + 1.25 = 3.25$	1/3.25
3	RAID 1+0 ($N = 2$ to 8)	RAID 0 ($N = 2$ to 8)	1	$2+1=3$	1/3
4	RAID 5 (see <i>Note 1</i>)	RAID 5 (When $N = 4$)	1	$1.25+1.25 = 2.5$	1/2.5
5	RAID 5 (see <i>Note 2</i>)	RAID 0 (When $N = 4$)	1	$1.25+1 = 2.25$	1/2.25

Note 1: Total amount of ShadowImage = $(1+1/N) \times 2$ (N = Number of data drives in RAID.)

Note 2: Total amount of ShadowImage = $(1+1/N)+1$

The recommended RAID configuration is RAID 5 (4D+1)/RAID 5 (4D+1):

- 4D+1 has an outstanding performance in RAID.
- Large number of 4D+1 has been adopted based on the past sales record.
- The cost can be cut down due to the large share of user data.

8.4 Others (Example: Mode Setting)

Set the copy pace for initial copy and restoring. The criteria are as follows:

1. Prioritize the resync operation (recommended).
Enter a value 15 for <size> to set the copy pace.
2. Select **Normal** only when placing importance on the performance to the P-VOL during a resync operation.

This can be specified when there is no update processes during the resync operation. The copy time for the resync operation is not guaranteed. (The copy time for resync operation delay depends on the I/O performance.)

3. For estimating the resync time, see section 7.3. **DO NOT SELECT** "Slow".

Table 8.2 Relationship between Copy Pace Parameters and Resync Process

No.	Parameters	Resync Pace	Resync Process	Host's I/O Performance
1	1 to 5	Slow <i>(do not select)</i>	Single copy unit (stripe size) is processed per copy every 10 ms. The process is performed in longer intervals if the I/O performance from the host is heavy. The time of resync completion cannot be guaranteed because the pace differs depending on the I/O performance from the host.	Read: Executing approximately 90 of host limitation is possible. Write: Executing approximately 90 of host limitation is possible.
2	6 to 10	Normal	The resync process is performed continuously, but the time of resync completion cannot be guaranteed. The pace differs depending on the I/O performance from the host.	Same as above.
3	11 to 15	Prior	The resync process is performed continuously and takes prior. Therefore, the I/O performance from the host is restricted. The time of resync completion is guaranteed.	Random Read: Random Write: Sequential Read: Sequential Write:

8.5 Recommendations

8.5.1 Pair Assignment

- Do not assign an LU (required for a quick response to a host) to a pair.
When LUs are paired, data written to a primary volume is also written to a secondary volume. This occurs particularly when the writing load become heavier due to a large number of write operations, writing data with a large block size, frequent write I/O operations, and continuous writing. Select the ShadowImage pair carefully. When applying ShadowImage to an LU with a heavy writing load, make the loads on other LUs lighter.
- Assign two different RAID groups to each of the primary and secondary volumes.
When a secondary volume is assigned to a RAID group in which the primary volume has been assigned, the reliability of data is lowered because a failure that occurs in a single drive affects both of the primary and secondary LUs. The performance becomes limited because the load of writing applied on a drive is doubled. Therefore, it is recommended to assign primary and secondary volumes to respective RAID groups (refer to section 8.2).
- Assign a small number of volumes within the same RAID group.
When volumes are assigned to the same RAID group and used as pair volumes, there may be a case where a pair creation or resynchronization for one of the volumes causes a restriction to be placed on performance of a host I/O, pair creation, resynchronization, etc., for the other volume(s) because of contention between drives. It is recommended that you assign a small number of (one or two) volumes to be paired to the same RAID group (refer to section 8.2). When creating two or more pairs within the same RAID group, standardize the controllers that control LUs in the same RAID group and pay attention to make the pair creation or resynchronization timely.
- For a primary volume, use the FC drives.
When a P-VOL is located in a RAID group consisting of the SATA drives, performance of a host I/O, pair creation, and pair resynchronization, etc., is lowered because of the lower performance of the SATA drive. Therefore, it is recommended to assign a primary volume to a RAID group consisting of the FC drives (refer to section 8.2).
- Assign four or more disks to the data disks.
- When the data disks that compose a RAID group are not sufficient, it affects the host performance and/or copying performance adversely because reading/writing from/to the drives is restricted. Therefore, when operating pairs with ShadowImage, it is recommended that you use an LU consisting of four or more data disks.

8.5.2 Pair Creation

- Create a pair when the I/O load is light.
Pair creation initiates data copying (initial copying) from a primary volume to a secondary volume and therefore affects the performance of the host I/O. It affects the same parity group to which the paired LUs belong. Therefore, it is best to create a pair when the I/O load is light. Perform an operation (such as a batch processing) at a time different than when a pair is being created.
- Restrict the creation of multiple pairs simultaneously.
If many pairs are created at the same time, pair creation not only affects the performance of the other processes, but it also takes longer to create the pair. Therefore, restrict creating multiple pairs simultaneously. Create pairs in order of necessity.

8.5.3 Resynchronization

- Resynchronize a pair at a time when the I/O load is light.
Resynchronization initiates data copying between the primary and secondary volumes and therefore affects the performance of the host I/O. It affects the same parity group to which the paired LUs belong. Therefore, it is best to resynchronize a pair when the I/O load is light. Perform an operation (such as a batch processing) at a time different from than when a pair is resynchronizing.
- Give priority to resynchronization.
It is possible to specify a priority to the resynchronization process and the host process. Give priority to the resynchronization process because the time required for it can be estimated. When priority is given to the host process, the influence of resynchronization on the host I/O performance can be reduced; it is difficult to estimate when the resynchronization process is completed because the progress of the resynchronization depends on the load of the host I/O (refer to section 8.4).

Chapter 9 System Operation Example

This section includes the following:

- Full-Time Operation (Backup and Resync at Night) (see section 9.1)
- Full-Time Split Operation (see section 9.2)

9.1 Full-Time Operation (Backup and Resync at Night)

- Backup using tape:
 - Backup must be performed from two ports at the same time.
 - The total of P-VOL and S-VOL must be less than or equal to 2 TBs.
E.g., if P-VOL is 1 TB, backup takes 2 hours and 40 minutes for 108 MB/sec.)
 - During resynchronization, backup is performed when the host load is at minimum (less than 1,000 IOPS).
 - Resynchronization is executable at a high speed of 80 MB/sec when the load of the updating process (Random) is low. When the load of the updating process is high, the processing speed should be approximately 50 MB/sec.
 - It takes 3 hours 40 minutes to resynchronize the entire 1 TB.

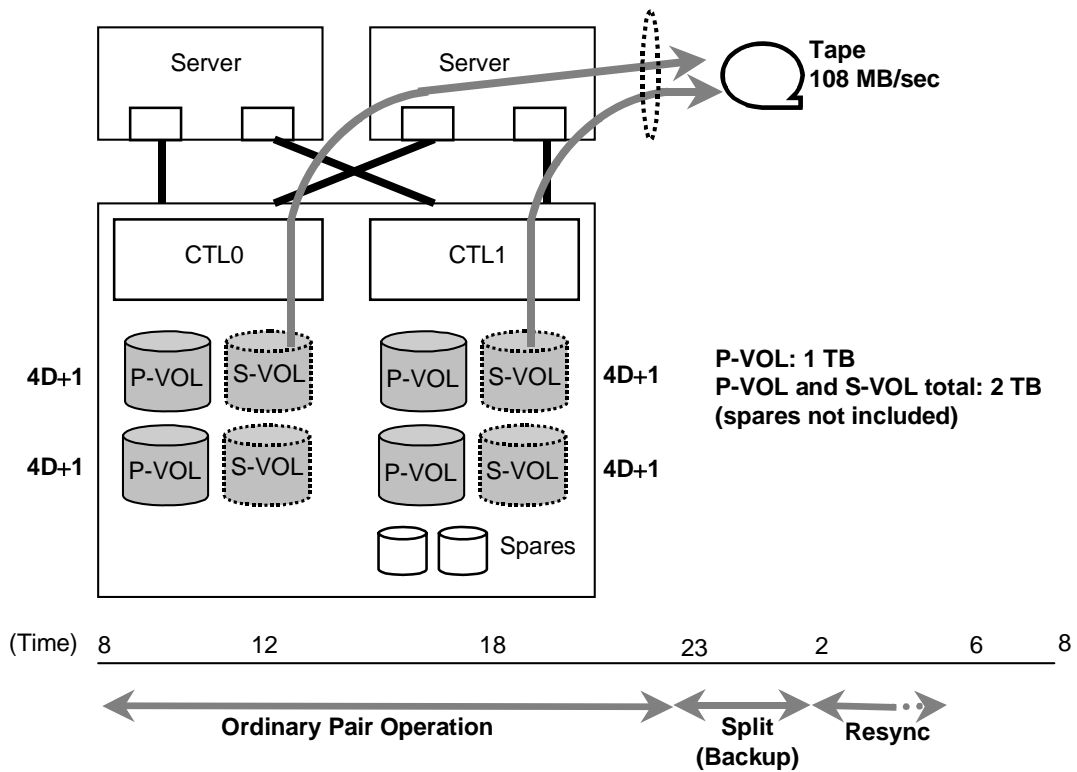


Figure 9.1 Backup Using Tape

Note: The figure above illustrates when six hours are allocated for the backup and the resynchronization time process. If more than six hours can be used for this process, a larger capacity of backup data is obtainable.

Chapter 10 Supported Platforms and Restrictions

This section includes the following:

- Supported Platforms (see section 10.1)
- Restrictions (see section 10.2)

10.1 Supported Platforms

The following table describes the supported platforms.

Table 10.1 External Specifications

No.	Platforms	Operating System Version
1	SUN™	Solaris™ 8 (SPARC)
		Solaris™ 9 (SPARC)
		Solaris™ 10 (SPARC)
		Solaris™ 10 (x86)
		Solaris™ 10 (x64)
2	PC Server (Microsoft®)	Windows® 2000
		Windows Server™ 2003 (IA32)
		Windows Server™ 2003 (x64)
		Windows Server™ 2003 (IA64)
3	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
4	IBM®	AIX® 5.1
		AIX® 5.2
		AIX® 5.3
5	Red Hat®	Red Hat® Linux® AS2.1 (IA32)
		Red Hat® Linux® AS3.0 (IA32)
		Red Hat® Linux® AS4.0 (IA32)
		Red Hat® Linux® AS3.0 (EM64T)
		Red Hat® Linux® AS4.0 (EM64T)
		Red Hat® Linux® AS3.0 (IA64)
		Red Hat® Linux® AS4.0 (IA64)
6	SGI®	IRIX® 6.5.x

10.2 Restrictions

This section describes the restrictions for configuration.

10.2.1 LU Mapping and ShadowImage Configuration

- When the Mapping mode is valid:
The P-VOL and S-VOL, whose mapping information is not set for the port that has been set in the configuration definition file (not recognizable from a host), cannot be paired. When you do not want to have them recognized by a host, use the LUN Security or LUN Management.

10.2.2 For VxVM and ShadowImage Configuration

- Volumes recognized by the host:
If the P-VOL and the S-VOL are recognized by the same host, the VxVM will not operate properly. Allow only the P-VOL to be recognized by the host and let another host recognize the S-VOL.

10.2.3 MSCS and ShadowImage Configuration

- When setting the S-VOL to be recognized by the host:
 - Use the CCI `mount` command instead of using the Disk Administrator. For details on the `mount` command, see the *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Command Control Interface (CCI) User and Reference Guide* (MK-95DF701).
 - The same host cannot recognize both the P-VOL and the S-VOL. Allow one host to recognize only the P-VOL and another host to recognize only the S-VOL.
 - Do not place the MSCS Quorum Disk in RAID Manager (CCI).
 - Shutdown MSCS before executing the CCI `sync` command.

10.2.4 AIX® and ShadowImage Configuration

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

10.2.5 Windows Server 2003 and ShadowImage TagmaStore Configuration

- When mounting a volume, use Volume {GUID} as an argument of the mount command of CCI. The Volume {GUID} can be used by CCI 01-13-03/00 or later. (For more detail, see the *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Command Control Interface (CCI) User and Reference Guide* (MK-95DF701).)
- When describing a command device in the configuration definition file, specify it as Volume{GUID}. (For more detail, see the *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Command Control Interface (CCI) User and Reference Guide* (MK-95DF701).)
- 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.

10.2.6 Linux and LVM Configuration

- Volumes recognized by the host:
If the P-VOL and the S-VOL are recognized by the same host, the LVM will not operate properly. Allow only the P-VOL to be recognized by the host and let another host recognize the S-VOL.

Chapter 11 Notes on System Configuration

This section includes the following:

- Previously Designed Procedure for Recovering Data during Malfunction (see section 11.1)
- Applying ShadowImage to an Existing System (see section 11.2)

11.1 Previously Designed Procedure for Recovering Data during Malfunction

In ShadowImage, the P-VOL and S-VOL synchronous may become corrupt due to a hardware malfunction. Even if the user has a hardware maintenance contract (including free warranty period) with Hitachi, recovering ShadowImage status (such as resynchronization) is beyond the maintenance personnel's task. Maintenance personnel's include removing the cause of malfunction and recovering the hardware only. Therefore, the procedure for recovering the ShadowImage status must be specified by the user when designing the system.

11.2 Applying ShadowImage to an Existing System

When newly applying ShadowImage to an already existing system, follow the description in Chapter 8 and determine the optimal configuration.

Appendix A Operations Using CLI

This section describes the following operation procedure for ShadowImage using the CLI of the Navigator. The following sections are included:

- Installing ShadowImage (see section A.1)
- Uninstalling ShadowImage (see section A.2)
- Enabling or Disabling ShadowImage (see section A.3)
- Setting the Command Device (see section A.4)
- Setting the Differential Management LU (see section A.5)
- Setting the Target ID (LU Mapping) (see section A.6)
- Setting the ShadowImage I/O Switching Mode (see section A.7)
- Confirming the Status of the Pairs (see section A.8)

For details on Navigator, refer to the *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Command Control Interface (CCI) User and Reference Guide* (MK-95DF701).

A.1 Installing ShadowImage

Since ShadowImage is an extra-cost option, ShadowImage cannot usually be selected (locked) before it can be used. To make ShadowImage available, you must install ShadowImage and set its function selectable (unlock).

Note: Before installing/uninstalling ShadowImage, verify that the array unit to be operated is functioning normally. If a failure such as a controller blockage has occurred, installation/uninstallation cannot be performed.

To install ShadowImage, a key code or key file provided with the optional feature is required. The following describes the installation procedure.

To install ShadowImage:

1. From the command prompt, register the subsystem (array unit) in which the ShadowImage is to be installed, then connect to the subsystem.
2. Execute the `auopt` command to install ShadowImage. The example is shown below.

Example:

Navigator version is 5.00 or later and Cache Partition Manager is enabled

```
% auopt -unit subsystem-name -lock off -keycode manual-attached-keycode
Password: manager-password
Are you sure you want to install the option? (y/n [n]): y
When Cache Partition Manager is enabled, if the option using data pool will be
enabled the default cache partition information will be restored.
Do you want to continue processing? (y/n [n]): y
The option is installed successfully.
%
```

Navigator version is less than 5.00 and Cache Partition Manager is enabled

```
% auopt -unit subsystem-name -lock off -keycode manual-attached-keycode
Password: manager-password
Are you sure you want to unlock the option? (y/n [n]): y
When Cache Partition Manager is enabled, if the option using data pool will be
enabled the default cache partition information will be restored.
Do you want to continue processing? (y/n [n]): y
The option is unlocked.
%
```

Navigator version is less than 3.00 and Cache Partition Manager is enabled

```
% auopt -unit subsystem-name -lock off -keycode manual-attached-keycode
Password: manager-password
Are you sure you want to unlock the option? (y/n [n]): y
The option is unlocked.
%
```

3. Execute the `auopt` command to confirm whether ShadowImage has been installed. The example is shown below.

Example:

```
% auopt -unit subsystem-name -refer
Password: manager-password
Option Name      Type      Term      Status
SHADOWIMAGE     Permanent ---      Enable
%
```

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

A.2 Uninstalling ShadowImage

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

Note: To uninstall ShadowImage, all ShadowImage pairs must be released (that is, the status of all LUs are SMPL).

The following describes the uninstallation procedure.

To uninstall ShadowImage:

1. From the command prompt, register the subsystem (array unit) in which the ShadowImage is to be uninstalled, then connect to the subsystem.
2. Execute the `auopt` command to uninstall ShadowImage. The example is shown below.

Example:

Navigator: Version 5.00 or later

```
% auopt -unit subsystem-name -lock on -keycode manual-attached-keycode
Password: manager-password
Are you sure you want to de-install the option? (y/n [n]): y
The option is de-installed successfully.
%
```

Navigator: Less than 5.00 version

```
% auopt -unit subsystem-name -lock on -keycode manual-attached-keycode
Password: manager-password
Are you sure you want to lock the option? (y/n [n]): y
The option is locked.
%
```

3. Execute the `auopt` command to confirm whether ShadowImage has been uninstalled. The example is shown below.

Example:

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

Uninstalling ShadowImage is now complete.

A.3 Enabling or Disabling ShadowImage

Once installed, ShadowImage can be enabled or disabled.

Note: To disable ShadowImage, all ShadowImage pairs must be released (that is, the status of all LUs are SMPL).

The following describes the enabling/disabling procedure.

1. From the command prompt, register the subsystem (array unit) in which the status of the feature is to be changed, then connect to the subsystem.
2. 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.

Example:

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

3. Execute the `auopt` command to confirm whether the status has been changed. The example is shown below.

Example:

```
% auopt -unit subsystem-name -refer
Password: manager-password
Option Name      Type      Term      Status
SHADOWIMAGE     Permanent ---      Disable
%
```

Enabling or disabling ShadowImage is now complete.

A.4 Setting the Command Device

The command device is a user-selected, dedicated logical volume on the TagmaStore subsystem that functions as the interface to the CCI software. The ShadowImage commands are issued by the CCI (HORCM) to the TagmaStore command device.

In order to accept read and write commands that are executed by the TagmaStore subsystem and return read requests to the UNIX/PC host, the command device must be designated. The command device must be defined in the HORCM_CMD section of the configuration definition file for the CCI instance on the attached host. Two command devices can be designated for the TagmaStore subsystem. You can designate command devices using the Navigator.

Note 1: LUs set for command devices must be recognized by the host. The Command Device LU size must be greater than or equal to 33 MB.

Note 2: The following restrictions apply when either a pair of ShadowImage, SnapShot, TrueCopy or TCE exists, or the path of True Copy or TCE is defined.

- When two command devices are set, only one command device can be released.
- When only one command device is set, the command device cannot be released.

To designate command device(s):

1. From the command prompt, register the subsystem (array unit) to which you want to create the command device. Connect to the subsystem.
2. Execute the `aucmddev` command to create a command device. The following is an example of specifying LU 2 for command device 1.

First, display LUs to be assignable command device, and later create a command device.

When you want to use the protection function of CCI, enter "enable" following the `-dev` option.

Example:

```
% aucmddev -unit subsystem-name -availablelist
Password: manager-password
Available Logical Units
  LUN  Capacity  RAID Group RAID Level  D-CTL C-CTL Type Status
    2   35.0 Mbyte      0  5( 4D+1P)      0    0 FC  Normal
    3   35.0 Mbyte      0  5( 4D+1P)      1    1 FC  Normal
%
% aucmddev -unit subsystem-name -set -dev 1 2
Password: manager-password
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 created. The following shows the example.

Note: To set the alternate command device function or to avoid data loss and subsystem downtime, designate two command devices. For details on alternate command device function, refer to the *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Command Control Interface (CCI) User and Reference Guide* (MK-95DF701).

Example:

```
% aucmddev -unit subsystem-name -refer
Password: manager-password
Command Device LUN RAID Manager Protect
1          2  Disable
%
```

4. To release an already set command device, specify as follows.

The following is an example of releasing command device 1.

Example:

```
% aucmddev -unit subsystem-name -rm -dev 1
Password: manager-password
Are you sure you want to release the command devices? (y/n [n]): y
The command devices have been released successfully.
%
```

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

Example:

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

A.5 Setting Differential Management LU

When the Differential Management LU is not set, create it. The Differential Management LU is an exclusive logical unit for storing the differential data at the time when the volume is copied. The Differential Management LU in the disk subsystem is treated in the same way as the other logical units. However, a logical unit that is set as the Differential Management LU is not recognized by a host (it is hidden).

It is required to set a logical unit with a size of 5 GB at the least as the Differential Management LU. Up to two Differential Management LUs can be set. The second one is used for the mirroring. It is recommended that two Differential Management LUs be set.

To designate Differential Management LUs:

1. From the command prompt, register the subsystem on which you want to create the differential management LU. Connect to the subsystem.
2. Execute the `audmlu` command to create a differential management LU.

First, display LUs to be assignable differential management LU, and later create a differential management LU.

Example:

```
% audmlu -unit subsystem-name -availablelist
Password: manager-password
Available Logical Units
  LUN Capacity   RAID Group  RAID Level  D-CTL  C-CTL Type Status
    0    5.0 GByte        0    5( 4D+1P)    0      0 FC  Normal
%
% audmlu -unit subsystem-name -set -lu 0
Password: manager-password
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 differential management LU, specify as follows.

Example:

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

Note: The following restrictions apply when a pair of ShadowImage, SnapShot, TrueCopy or TCE exists, the path of True Copy or TCE is defined, or POOL of SnapShot or TCE is defined.

- When two differential management LUs are set, only one differential management LU can be released.
- When only one differential management LU is set, the differential management LU cannot be released.

A.6 Setting the Target ID (LU Mapping)

The following is the procedure for specifying the target ID. The target ID is specified using the Navigator:

Note: For iSCSI model, use an autargetmap command instead of an auhgmap command.

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

Example:

```
% auhgmap -unit subsystem-name -MappingMode on
Password: manager-password
Are you sure you want to set the mapping mode? (y/n [n]): y
when setting starts, the subsystem stops accepting access to the controller from
the host.
Before setting, stop access to the controller from the host.
Do you want to continue processing? (y/n [n]): y
The mapping mode has been set successfully.
%

% auhgmap -unit subsystem-name -add 0 A 0 6 0
Password: manager-password
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 Target ID (LU Mapping) has been set.

Example:

```
% auhgmap -unit subsystem-name -refer
Password: manager-password
Mapping mode = ON
Port  Group  H-LUN  LUN
  0A      0      6      0
%
```

A.7 Setting the ShadowImage I/O Switching Mode

When using ShadowImage I/O Switching Function, set the ShadowImage I/O Switching mode with Navigator.

To set the ShadowImage I/O Switching Mode:

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

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

Example:

```
% ausystemparam -unit subsystem-name -set -ShadowImageIOSwitch enable
Password: manager-password
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.

Example:

```
% ausystemparam -unit subsystem-name -refer
Password: manager-password
---- Common Parameter ----
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 ShadowImage I/O Switching mode, it is necessary to make pair statuses of all ShadowImage pairs to those other than PSUS (S-VOL Switch) and COPY (S-VOL Switch).

A.8 Confirming the Status of Pairs

You can display the ShadowImage paired logical units using CLI. The following steps describe the procedure:

1. From the command prompt, register the subsystem (array unit) to which you want to display the status of paired logical volumes. Connect to the subsystem.
2. Execute the `aumrcfluc` command to display the pair status.

Example:

```
% aumrcfluc -unit subsystem-name
Password: manager-password
  RAID   RAID   Capacity
LU Group Level  [block] D-CTL C-CTL Status
2/ 3 0/1   5 / 5 100000000 0    0 PAIR
7/ 8 2/3   1 / 1 100000000 1    1 PAIR
%
```

The pair status is displayed.

Acronyms

Acronym	Expansion
AMS	Adaptable Modular Storage
CCI	command control interface
CLI	command line interface
CTG	consistency group
CTL	controller
CU	controller unit
dB(A)	decibel (A-weighted)
DM-LU	differential management logical unit
FC	fibre channel
GUI	raphical user interface
HDD	hard disk drive
HDLM	Hitachi Dynamic Link Manager
HIPAA	Health Insurance Portability and Accountability Act
H-LUN	host logical unit
IAEA	International Atomic Energy Agency
IOPS	input output operations per second
IOS	internet work operating system
iSCSI	internet small computer system interface
LDEV	logical device

LU	logical unit
LVM	logical volume manager
MU	mirror unit
NAS	network attached storage
NSC	network storage controller
PSUS	pair suspended-split
P-VOL	primary volume
RAID	edundant array of independent disks
RCU	remote control unit
RPO	recovery point objective
RTO	recovery time objective
SAN	storage-area networkRTC
SNIA	Storage Networking Industry Association
SNMP	simple network management protocol
S-VOL	secondary volume
TCE	TrueCopy Extended Distance
TID	target identifier
T-VOL	target volume
USP	Universal Storage Platform
V-VOL	virtual volume (Snapshot Image)
VxVM	Veritas Volume Manager

