

Hitachi AMS 2000 Family TrueCopy Remote Replication User Guide

FASTFIND LINKS

Document organization

Release notes and readme

Getting help

Table of Contents

@Hitachi Data Systems

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

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

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

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

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

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

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

Table of Contents

	Prefaceix
	Intended audience Product version Release notes and readme Document revision level Changes in this revision Document organization Conventions Convention for storage capacity values Accessing product documentation Setting help
1	Overview1-1
	TrueCopy synchronous remote replication 1-2 How TrueCopy works 1-2 Typical environment 1-2 Volume pairs 1-3 Remote Path 1-4 Differential Management LUs (DMLU) 1-4 Consistency Group (CTG) 1-4 TrueCopy interfaces 1-4 Typical workflow 1-5 Operations overview 1-5
2	Remote path design2-1
	Determining remote path bandwidth

	Remote path requirements, supported configurations	2-8
	Management LAN requirements	2-9
	Remote path requirements	
	WAN Optimization Controller (WOC) requirements	
	Remote path configurations	
	Remote path configurations for Fibre channel	
	Direct connection	
	Switch connections	
	Fibre channel extender	
	Path and switch performance	
	Port transfer rate for Fibre channel	
	Remote path configurations for iSCSI	
	Direct iSCSI connection	
	Single LAN switch, WAN connection	
	Multiple LAN switch, WAN connection	
	Single LAN switch, WOC, WAN connection	
	Multiple LAN switch, WOC, WAN connection	2-19
	Multiple array, LAN switch, WOC connection with single WAN	2-20
	Multiple array, LAN switch, WOC connection with two WANs	2-21
	Using the remote path — best practices	2-22
	Remote processing	2-22
3	Planning for TrueCopy	3-1
	The planning workflow	
	Planning arrays	3-2
	Planning Logical Units	
	Volume pair recommendations	
	LUN expansion	
	Operating system recommendations and restrictions	
	Host time-out	
	P-VOL, S-VOL recognition by same host on VxVM, AIX®, LVM	3-6
	Setting the Host Group options	
	Windows 2000 Servers	
	Windows Server 2003/2008	
	Identifying P-VOL and S-VOL LUs on Windows	
	VMWare and TrueCopy configuration	
	Concurrent use of Dynamic Provisioning	
	Calculating supported capacity	
	On cascading TrueCopy with local replication software	
	Cascade connection restrictions	
4	Requirements	<u> 4-1</u>
•	•	
	System requirements	4-2

	Displaying the hardware revision	4-3
5	Installation and setup	. 5-1
	Installation procedures	5-2
	Installation	
	Enabling, disabling	
	Uninstalling TrueCopy	
	Setup procedures	
	Defining the DMLU	5-4
	Adding or changing the remote port CHAP secret	5-4
	Defining the remote path	5-5
6	Using TrueCopy	.6-1
	TrueCopy operations	
	Checking pair status	
	Creating pairs	
	Prerequisite information and best practices	
	Copy pace	
	Fence level	
	Create pair procedure	
	Splitting pairs	
	Resynchronizing pairs	
	Swapping pairs	
	Editing pairs	
	Deleting pairs	
	Deleting the remote path	
7	Disaster recovery	.7-1
	Disaster recovery	
	Resynchronizing the pair	
	Data path failure and recovery	
	Host server failure and recovery	
	Production site failure and recovery	
	Special problems and recommendations	
8	Monitoring, maintenance, troubleshooting	.8-1
_		
	Monitoring and maintenance	
	Monitoring the remote path	
	Troubleshooting	
	Pair failure	

A	Specifications	\-1
	TrueCopy specifications	4-2
В	Operations using CLI E	3-1
	Installation and setup	B-2
	Installing	
	Enabling or disabling	
	Uninstalling	
	Setting the Differential Management-Logical Unit	
	Release a DMLU	
	Setting the remote port CHAP secret	
	Setting the remote path	
	Deleting the remote path	
	Pair operations	
	Displaying status for all pairs	
	Displaying detail for a specific pair	
	Creating a pair	
	Creating pairs belonging to a group	
	Splitting a pair	
	Swapping a pair	
	Deleting a pair	
	Changing pair information	
	Sample scripts	
	Backup script	
	Pair-monitoring script	
С	Operations using CCI	:-1
	Preparing for CCI operations	
	Setting the command device	
	Setting LU mapping	
	Defining the configuration definition file	C
	Setting the environment variable	
	Pair operations	
	Multiple CCI requests and order of execution	
	Operations and pair status	
	Confirming pair status	
	Creating pairs (paircreate)	
	Splitting pairs (pairsplit)	
	Resynchronizing pairs (pairresync)	
	Suspending pairs (pairsplit -R)	
	Releasing pairs (pairsplit -S)	
	Mounting and unmounting a volume	
	mounting and unmounting a volume	12

D	Cascading with ShadowImage				. D) -1
	Cascade overview				. [)-2
	Cascade configurations					
	Configurations with ShadowImage P-VOLs				. [)-3
	Configurations with ShadowImage S-VOLs					
	Configurations with ShadowImage P-VOLs and S-VOLs	٠.			[)-6
	Replication operations supported		٠.		[)-7
	TrueCopy operations supported					
	ShadowImage operations supported					
	Read/write operations, combined status supported					
	Guidelines and restrictions					
	Configurations not supported					
	Creating a backup with ShadowImage		٠.		D-	-14
Е	Cascading with SnapShot				. E	:-1
	Cascade overview					=-2
	Cascade configurations					
	Configurations with SnapShot P-VOLs					
	Configurations with SnapShot V-VOLs					
	Replication operations supported					
	TrueCopy operations supported					
	SnapShot operations supported					
	Read/write, status combinations supported					
	Guidelines and restrictions				F	E-7
	Configurations not supported				F	E-8
	Creating a backup with SnapShot	٠.			F	E-9
	When to create a backup	٠.			. E-	-10
_					_	
F	Wavelength Division Multiplexing (WDM) and dark fibre	•	•	• •	. Р	·-1
	WDM and dark fibre	٠.			!	- -2
	Glossary					

Index

Preface

This manual provides instructions and information for planning, installing, configuring, using, monitoring, and troubleshooting TrueCopy Remote Replication software. See a full description in Chapter 1, Overview.

This preface includes the following information:

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

Preface ix

Intended audience

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

Product version

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

Release notes and readme

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

Document revision level

Revision	Date	Description
MK-97DF8052-00	October 2008	Initial Release
MK-97DF8052-01	December 2008	Supersedes and replaces MK-97DF8052-00
MK-97DF8052-02	February 2009	Supersedes and replaces MK-97DF8052-01
MK-97DF8052-03	May 2009	Supersedes and replaces MK-97DF8052-02
MK-97DF8052-04	August 2009	Supersedes and replaces MK-97DF8052-03
MK-97DF8052-05	November 2009	Supersedes and replaces MK-97DF8052-04
MK-97DF8052-06	January 2010	Supersedes and replaces MK-97DF8052-05
MK-97DF8052-07	April 2010	Supersedes and replaces MK-97DF8052-06
MK-97DF8052-08	August 2010	Supersedes and replaces MK-97DF8052-07

Changes in this revision

This release includes updates or changes to the following:

- Added the section Remote processing (page 2-22).
- Added the section VMWare and TrueCopy configuration (page 3-11).
- Added the section Concurrent use of Dynamic Provisioning (page 3-11).
- Added the section Cascade connection restrictions (page 3-17).
- Added drives in the section Volume pair recommendations (page 3-3).
- Updated a parameter in Table A-1 (page A-2).

X Preface

Document organization

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

Chapter/Appendix Title	Description
Chapter 1, Overview	Provides a description of TrueCopy, how it works, and its components.
Chapter 2, Remote path design	Provides information and procedures for designing the data path between local and remote arrays.
Chapter 3, Planning for TrueCopy	Provides information to help you plan a TrueCopy system.
Chapter 4, Requirements	Provides minimum operational requirements.
Chapter 5, Installation and setup	Provides instructions for installing and performing setup procedures.
Chapter 6, Using TrueCopy	Provides instructions and information for performing TrueCopy operations.
Chapter 7, Disaster recovery	Provides procedures and recommendations for switching host operations, pair swapping, and restoring the production volume.
Chapter 8, Monitoring, maintenance, troubleshooting	Provides information on monitoring, maintaining and troubleshooting.
Appendix A, Specifications	Provides TrueCopy specifications.
Appendix B, Operations using CLI	Provides detailed Navigator 2 Command Line Interface instructions for configuring and using TrueCopy.
Chapter C, Operations using CCI	Provides detailed Command Control Interface instructions for configuring and using TrueCopy.
Appendix D, Cascading with ShadowImage	Provides ShadowImage cascade information including supported configurations and operations.
Appendix E, Cascading with SnapShot	Provides SnapShot cascade information including supported configurations and operations.
Appendix F, Wavelength Division Multiplexing (WDM) and dark fibre	Provides information for using WDM and dark fibre in the data path network.
Glossary	Defines special terms and acronyms used in this document.
Index	Provides links to specific information in this manual.

Preface **xi**

Document conventions

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

Symbol	Meaning	Description
Tip Tips provide helpful information, guidelines, for performing tasks more effectively.		Tips provide helpful information, guidelines, or suggestions for performing tasks more effectively.
\triangle	Note	Notes emphasize or supplement important points of the main text.
<u>^</u>	Caution	Cautions indicate that failure to take a specified action could result in damage to the software or hardware.
	WARNING	Warnings indicate that failure to take a specified action could result in loss of data or serious damage to hardware.

The following typographic conventions are used in this document.

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

xii Preface

Convention for storage capacity values

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

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

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

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

Accessing product documentation

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

This documentation set consists of the following documents.

Release notes

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



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

Preface **xiii**

Installation and getting started

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

AMS2100/2300 Getting Started Guide, MK-98DF8152

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

AMS2500 Getting Started Guide, MK-97DF8032

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

AMS 2000 Family Site Preparation Guide, MK-98DF8149

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

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

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

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

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

Storage and replication features

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

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

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

xiv Preface

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Storage Navigator Modular 2 online help

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

Hardware maintenance and operation

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

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

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

AMS 2500 Storage System Hardware Guide, MK-97DF8007

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

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

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

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

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

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

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

xvi Preface

Command and Control (CCI)

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

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

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

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

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

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

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

Command Line Interface (CLI)

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

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

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

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

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

Preface **xvii**

Dynamic Replicator documentation

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

Dynamic Replicator - Scout Release Notes, RN-99DF8211

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

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

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

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

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

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

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

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

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

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

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

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

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

Dynamic Replicator - Scout File Server Solution, MK-98DF8225

Dynamic Replicator - Scout ESX - Protecting ESX Server (RCLI), MK-99DF8226

xviii Preface

Getting help

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

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

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

Comments

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

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

Preface **xix**

Preface

Hitachi AMS 2000 Family TrueCopy Remote Replication User Guide



Overview

A broken link, an accidentally erased file, the force of nature: negative occurrences cause problems to storage systems. When access to critical data is interrupted, a business can suffer irreparable harm.

TrueCopy helps you keep critical data backed up in a remote location, so that negative incidents do not have a lasting negative impact.

This manual provides information and instructions for planning and designing a TrueCopy system, implementing the design, creating and maintaining remote copies, and monitoring and troubleshooting the system.

This overview consists of the following sections:

- TrueCopy synchronous remote replication
 How TrueCopy works
 Typical environment
 TrueCopy interfaces
 Typical workflow
- Operations overview

TrueCopy synchronous remote replication

TrueCopy creates a duplicate of a production volume to a secondary volume located at a remote site. Data in a TrueCopy backup stays synchronized with the data in the local array. This happens when data is written from the host to the local array then to the remote system, via fibre channel or iSCSI link. The host holds subsequent output until acknowledgement is received from the remote array for the previous output.

When a synchronized pair is split, writes to the primary volume are no longer copied to the secondary side. Doing this means that the pair is no longer synchronous. Output to the local array is cached until the primary and secondary volumes are re-synchronized. When resynchronization takes place, only the changed data is transferred, rather than the entire primary volume. This reduces copy time.

TrueCopy can be teamed with ShadowImage or SnapShot, on either or both local and remote sites. These in-system copy tools allow restoration from one or more additional copies of critical data.

Besides disaster recovery, TrueCopy backup copies can be used for development, data warehousing and mining, or migration applications.

How TrueCopy works

A TrueCopy "pair" is created when you:

- Select a volume on the local array that you want to copy.
- Create or identify the volume on the remote array that will contain the copy.
- Connect the local and remote arrays with a fibre channel or iSCSI link
- Copy all primary volume data to the secondary volume.

Under normal TrueCopy operations, all data written to the primary volume is copied to the secondary volume, insuring that the secondary copy is a complete and consistent backup.

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

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

To plan a TrueCopy system an understanding is needed of its components.

Typical environment

A typical configuration consists of the following elements. Many but not all require user set up.

1–2 Overview

- Two AMS arrays—one on the local side connected to a host, and one on the remote side connected to the local array. Connections are made via fibre channel or iSCSI.
- A primary volume (P-VOL) on the local array that is to be copied to the secondary volume (S-VOL) on the remote side. Primary and secondary volumes may be composed of several LUs.
- A differential management LU on local and remote arrays, which hold TrueCopy information when the array is powered down
- Interface and command software, used to perform TrueCopy operations. Command software uses a command device (volume) to communicate with the arrays.

Figure 1-1 shows a typical TrueCopy environment.

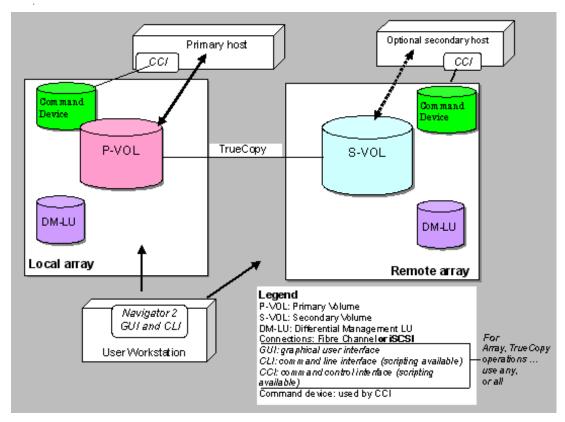


Figure 1-1: TrueCopy Components

Volume pairs

As described above, original data is stored in the P-VOL and the remote copy is stored in the S-VOL. The pair can be paired, split, re-synchronized, and returned to the simplex state. When synchronized, the volumes are paired; when split, new data sent is to the P-VOL but held from the S-VOL. When re-synchronized, changed data is copied to the S-VOL. When necessary, data in the S-VOL can be copied to the P-VOL (P-VOL restoration). Logical units on the local and remote arrays must be defined and formatted prior to pairing.

Overview 1–3

Remote Path

TrueCopy operations are carried out between local and remote arrays connected by a fibre channel or iSCSI interface. A data path, referred to as the remote path, connects the port from the local array that executes the volume replication to the port on the remote array. User setup is required on the local array.

Differential Management LUs (DMLU)

The DMLU is an exclusive volume used for storing TrueCopy information when the local or remote array is powered down. The DMLU is treated the same as other volumes in the storage subsystem, but is hidden from a host. User setup is required on the local and remote arrays. See Defining the DMLU on page 5-4.

Consistency Group (CTG)

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

Managing primary volumes as a consistency group allows TrueCopy operations to be performed on all volumes in the group concurrently. Write order in secondary volumes is guaranteed across application logical volumes. User setup is required.

TrueCopy interfaces

TrueCopy can be operated using of the following interfaces:

- The GUI (Hitachi Storage Navigator Modular 2 Graphical User Interface) is a browser-based interface from which TrueCopy can be setup, operated, and monitored. The GUI provides the simplest method for performing operations, requiring no previous experience. Scripting is not available.
- CLI (Hitachi Storage Navigator Modular 2 Command Line Interface), from which TrueCopy can be setup and all basic pair operations can be performed—create, split, resynchronize, restore, swap, and delete. The GUI also provides these functionalities. CLI also has scripting capability.
- CCI (Hitachi Command Control Interface (CCI), used to display volume information and perform all copying and pair-managing operations. CCI provides a full scripting capability which can be used to automate replication operations. CCI requires more experience than the GUI or CLI. CCI is required for performing failover and fall back operations. It is also required on Windows 2000 Server and Windows Server 2008 for mount/unmount operations.

1–4 Overview

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

 \triangle

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

Typical workflow

Designing, creating, and using a TrueCopy system consists of the following tasks:

- Planning in which you assemble the necessary components of a TrueCopy system. This includes establishing path connections between the local and remote arrays, LUN sizing and RAID configurations, understanding how to use TrueCopy concurrently with ShadowImage and/or Copy-on-Write, and other necessary pre-requisite information and tasks.
- Design in which you gather business requirements and write workload data to size TrueCopy remote path bandwidth to fit your organization's requirements.
- Configuration during which you implement the system and create an initial pair.
- Operations in which you perform copy and maintenance operations.
- Monitoring the system
- Troubleshooting

Operations overview

The basic TrueCopy operations are shown in Figure 1-2. They consist of creating, splitting, resynchronizing, swapping, deleting a pair.

- Create Pair. This establishes the initial copy using two logical units that you specify. Data is copied from the P-VOL to the S-VOL. The P-VOL remains available to the host for read and write throughout the operation. Writes to the P-VOL are duplicated to the S-VOL. The pair status changes to Paired when the initial copy is complete.
- Split. The S-VOL is made identical to the P-VOL and then copying from the P-VOL stops. Read/write access becomes available to and from the S-VOL. While the pair is split, the array keeps track of changes to the P-VOL and S-VOL in track maps. The P-VOL remains fully accessible in Split status.
- Resynchronize pair. When a pair is re-synchronized, changes in the P-VOL since the split is copied to the S-VOL, making the S-VOL identical to the P-VOL again. During a resync operation, the S-VOL is inaccessible to hosts for write operations; the P-VOL remains accessible

Overview **1–5**

for read/write. If a pair was suspended by the system because of a pair failure, the entire P-VOL is copied to the S-VOL during a resync.

- Swap pair. The pair roles are reversed.*
- Delete pair. The pair is deleted and the volumes return to Simplex status.

Simplex

Synchronizing S-VOL

Paired S-VOL

Split

S-VOL

Resync S-VOL

Swap

P-VOL

Delete (Simplex)

Figure 1-2: TrueCopy Pair Operations

See the individual procedures and more detailed information in Chapter 6, Using TrueCopy.

1–6 Overview

Remote path design

A remote path must be designed to adequately manage your organization's data throughput. This chapter provides instructions for analyzing business requirements, measuring write-workload, and calculating your system's changed data over a given period. Remote path configurations are also provided.

- Determining remote path bandwidth
- ☐ Remote path requirements, supported configurations

Determining remote path bandwidth

Bandwidth for the TrueCopy remote path is based on the amount of production output to the primary volume. Sufficient bandwidth must be present to handle the transfer of all workload levels, from average MB/sec to peak MB/sec. Planning bandwidth also accounts for growth and a safety factor.

Measuring write-workload

To determine the bandwidth necessary to support a TrueCopy system, the peak workload must be identified and understood. Workload data is collected using performance monitoring software on your operating system. This data is best collected over a normal monthly cycle. It should include end-of month, quarter, or year, or times when workload is heaviest.

To collect workload data

- 1. Using your operating system's performance monitoring software, collect the following:
 - I/O per second (IOPS).
 - Disk-write bytes-per-second for every physical volume that will be replicated.

The data should be collected at 5-10 minute intervals, over a 4-6 week period. The period should include periods when demand on the system is greatest.

2. At the end of the period, convert the data to MB-per-second, if it is not already so. Import the data into a spreadsheet tool. Figure 2-1 shows graphed data throughput in MB per second.

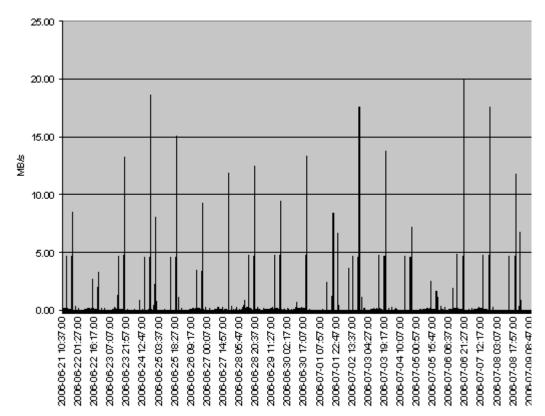


Figure 2-1: Data Throughput in MB/sec

Figure 2-2 shows IOPS throughput over the same period.

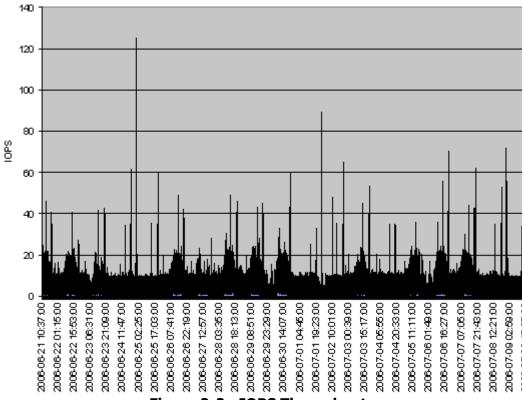


Figure 2-2: IOPS Throughput

- 3. Locate the highest peak to determine the greatest MB-per-second workload.
- 4. Be aware of extremely high peaks. In some cases, a batch job, defragmentation, or other process could be driving workload to abnormally high levels. It is sometimes worthwhile to review the processes that are running. After careful analysis, it may be possible to lower or even eliminate some spikes by optimizing or streamlining highworkload processes. Changing the timing of a process may lower workload; another option may be to schedule a suspension of the TrueCopy pair (split the pair) when a spiking process is active.
- 5. With peak workload established, take into consideration the following:
 - Channel extension. Extending fibre channel over IP telecommunication links changes workloads.
 - The addition of the IP headers and conversion from fibre channel's 2112-byte frames to the 1500-byte Maximum Transfer Unit of Ethernet add approximately 10% to the amount of data transferred.
 - Compression is also a factor. The exact compression ratio is dependent on the compressibility of the data and speed of the telecommunications link. Hitachi Data Systems uses 1.8:1 as a compression rule-of-thumb, though real-life ratios are typically higher.
 - Projected growth rate accounts for the increase expected in write workload over a 1, 2, or 3 year period.
 - Safety factor adds extra bandwidth for unusually high spikes that might occur.
- 6. The bandwidth must be at least as large as the peak MB/sec, including channel extension overhead and compression ratios.

Optimal I/O performance versus data recovery

An organization's business requirements affect I/O performance and data recovery. Business demands indicate whether a pair must be maintained in the synchronized state or can be split. Understanding the following comparisons is useful in refining bandwidth requirements.

When a pair is synchronizing, the host holds any new writes until confirmation that the prior write is copied to both the P-VOL and the S-VOL. This assures a synchronous backup, but the resulting latency impacts I/O performance.

On the other hand, a pair in Split status has no impact on host I/O. This advantage is offset if a failure occurs. In a pair failure situation, data saved to the P-VOL while the pair is split is not copied to the S-VOL, and may therefore be lost. When performance is the priority, data recovery is impacted.

To deal with these competing demands, an organization must determine its recovery point objective (RPO), which is how much data loss is tolerable before the business suffers significant impact. RPO shows the point back in

time from the disaster point that a business must recover to. A simple illustration is shown in Figure 2-3. The RPO value shows you whether a pair can be split, and how long it can stay split.

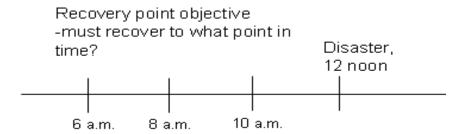


Figure 2-3: Determining Recovery Point

- If no data loss can be tolerated, the pair is never split and remains in the synchronous, paired state.
- If one hour of data loss can be tolerated, a pair may be split for one hour.
- If 8 hours of data loss can be tolerated, a pair may be split for 8 hours.

Finding the recovery point means exploring the number of lost business transactions business can survive, determining the number of hours that may be required to key-in or otherwise recover lost data.

Performance and data recovery are also affected when the TrueCopy system is cascaded with ShadowImage. The following sections describe the impact of cascading.

Performance, data recovery with ShadowImage

I/O performance and data recovery can be enhanced by cascading the TrueCopy system with ShadowImage. Your priority, performance or data recovery, is achieved by where ShadowImage is used—on the local side or remote side. The following sections discuss these options.

Impact with ShadowImage on the local side

Minimal impact to host I/O performance occurs when ShadowImage on the local side is used to transfer data to the remote side. This is done by backing up the primary volume with ShadowImage, and then using the backup copy to transfer the data to the remote side via TrueCopy, as shown in Figure 2-4.

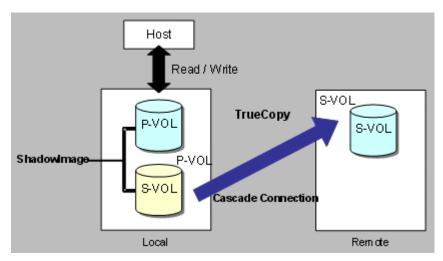


Figure 2-4: TrueCopy to Remote Side from Local ShadowImage

Distance between local and remote arrays, and low bandwidth do not factor into I/O performance when ShadowImage is used in this way, as illustrated in Figure 2-5.

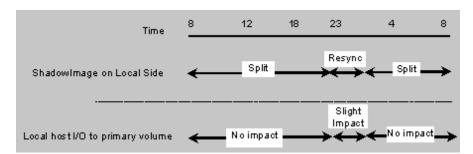


Figure 2-5: I/O Performance with Local ShadowImage

Copying from the local ShadowImage to the remote side leaves the primary volume free to be more quickly written to by the host. In addition, copying and especially resynchronizing the P-VOL to ShadowImage is far quicker than an equivalent TrueCopy operation to the remote side. Using ShadowImage in this way results in an asynchronous backup of the primary volume. Complete data recovery is impacted.

Impact with ShadowImage on the remote side

When ShadowImage is on the remote side, the TrueCopy operation occurs normally, that is, between the local P-VOL and the remote S-VOL. This is shown in Figure 2-6.

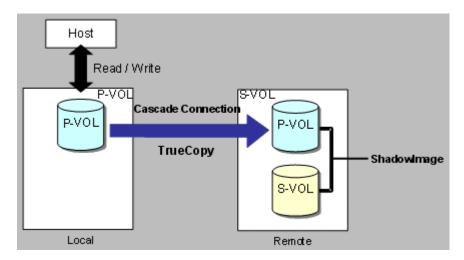


Figure 2-6: TrueCopy Remote Side to Remote ShadowImage

Backup data in the ShadowImage S-VOL is available in the event of failure during a TrueCopy resynchronization. Data recovery is assured.

Including ShadowImage in the TrueCopy system results in asynchronous remote backup. See Appendix D, Cascading with ShadowImage for supported configurations, operations, combinations of TrueCopy/ ShadowImage status, and recommendations.

Remote path requirements, supported configurations

The remote path is the connection used to transfer data between the local array and remote array. TrueCopy supports fibre channel and iSCSI port connectors and connections.

The following kinds of networks are used with TrueCopy:

- Local Area Network (LAN), for system management. Fast Ethernet is required for the LAN.
- Wide Area Network (WAN) for the remote path. For best performance:
 - A fibre channel extender is required.
 - iSCSI connections may require a WAN Optimization Controller (WOC).

Figure 2-7 shows the basic TrueCopy configuration with a LAN and WAN. More specific configurations are shown in Remote path configurations on page 2-10.

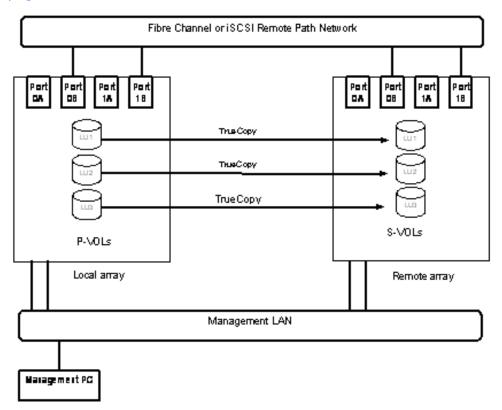


Figure 2-7: Remote Path Configuration

Requirements are provided in the following:

- Management LAN requirements on page 2-9
- Remote path requirements on page 2-9
- WAN Optimization Controller (WOC) requirements on page 2-10
- Fibre channel extender on page 2-13.

Management LAN requirements

Fast Ethernet is required for an IP LAN.

Remote path requirements

This section discusses the TrueCopy remote path requirements for a WAN connection. This includes the following:

- Types of lines
- Bandwidth
- Distance between local and remote sites
- WAN Optimization Controllers (WOC) (optional)

For instructions on assessing your system's I/O and bandwidth requirements, see:

- Measuring write-workload on page 2-2
- Determining remote path bandwidth on page 2-2

Table 2-1 provides remote path requirements for TrueCopy. A WOC may also be required, depending on the distance between the local and remote sites and other factors listed in Table 2-3.

Table 2-1: Remote Path Requirements

Item	Requirements
Bandwidth	 Bandwidth must be guaranteed. Bandwidth must be 1.5 Mb/s or more for each pair. 100 Mb/s recommended. Requirements for bandwidth depend on an average inflow from the host into the array.
Remote Path Sharing	 The remote path must be dedicated for TrueCopy pairs. When two or more pairs share the same path, a WOC is recommended for each pair.

Table 2-2 shows types of WAN cabling and protocols supported by TrueCopy and those not supported.

Table 2-2: Supported, Un-supported WAN Types

	WAN Types
Supported	Dedicated Line (T1, T2, T3 etc)
Not-supported	ADSL, CATV, FTTH, ISDN

WAN Optimization Controller (WOC) requirements

WAN Optimization Controller (WOC) is a network appliance that enhances WAN performance by accelerating long-distance TCP/IP communications. With a WOC, TrueCopy copy performance across long distance increases more than when only a WAN is used. A WOC guarantees bandwidth for each line.

- Use Table 2-3 to determine whether the TrueCopy system requires the addition of a WOC.
- Table 2-4 shows the requirements for WOC.

Table 2-3: Conditions Requiring a WOC

Item	Condition
Latency, Distance	• If round trip time is 5 ms or more, or distance between the local site and the remote site is 100 miles (160 km) or further, WOC is highly recommended.
WAN Sharing	If two or more pairs share the same WAN, WOC is recommended for each pair.

Table 2-4: WOC Requirements

Item	Requirements
LAN Interface	Gigabit Ethernet or fast Ethernet must be supported.
Performance	Data transfer capability must be equal to or more than bandwidth of WAN.
Functions	 Traffic shaping, bandwidth throttling, or rate limiting must be supported. These functions reduce data transfer rates to a value input by the user. Data compression must be supported. TCP acceleration must be supported.

Remote path configurations

One remote path must be set up per controller, two paths per array. With two paths, an alternate path is available in the event of link failure during copy operations.

Paths can be constructed from:

- Local controller 0 to remote controller 0 or 1
- Local controller 1 to remote controller 0 or 1

Paths can connect a port A with a port B, and so on.

The following sections describe Fibre channel and iSCSI path configurations. Recommendations and restrictions are included.

Remote path configurations for Fibre channel

The AMS array supports direct and switch fibre channel connections only. Hub connections are not supported.

- Direct connection (loop only) is a direct link between the local and remote arrays.
- Switch connections push data from the local array through a fibre channel link across a WAN to the remote switch and fibre channel to the remote array. Switch connections increase throughput between the arrays. F-Port (Point-to-Point) and FL-Port (Loop) switch connections are supported.
- iSCSI and FC support point-to-point direct connections

Direct connection

A direct connection is a standard point-to-point fibre channel connection between ports, as shown in Figure 2-8. Direct connections are typically used for systems 500 meters to 10 km apart.

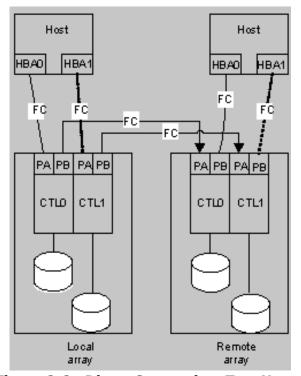


Figure 2-8: Direct Connection, Two Hosts

- Optimal performance occurs when the paths are connected to parallel controllers, that is, local controller 0 is connected with remote controller 0, and so on.
- Between a host and array, only one remote path is required. However, two are recommended, with one available for a backup.

Figure 2-9 shows path configurations using a single switch per array.

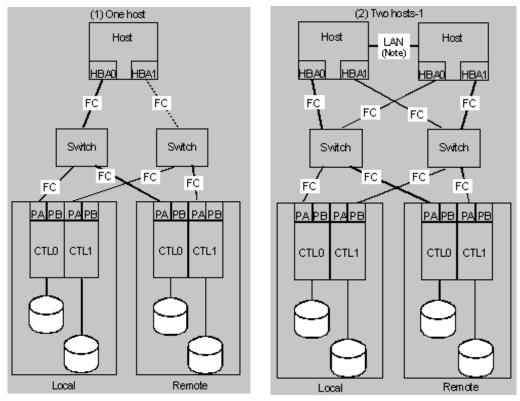


Figure 2-9: Switch Connection, 1 Host, 2 Hosts

Figure 2-10 shows multiple switches per array.

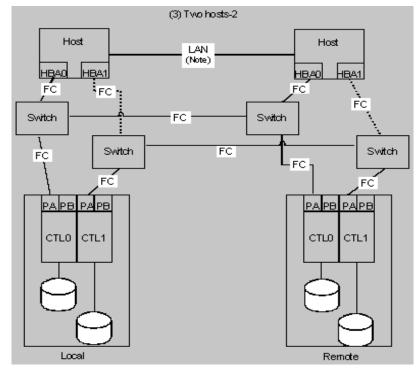


Figure 2-10: Multiple Switches with Two Hosts

Recommendations

When two hosts exist, a LAN is required to provide communication between local and remote CCIs, when used. In this case, the local host activates the CCIs on the local and remote side.

Fibre channel extender

Distance limitations can be reduced when channel extenders are used with fibre channel. This section provides configurations and recommendations for using channel extenders, and provides information on WDM and dark fibre.

Figure 2-11 illustrates two remote paths using two FC switches, Wavelength Division Multiplexor (WDM) extender, and dark fibre to make the connection to the remote site.

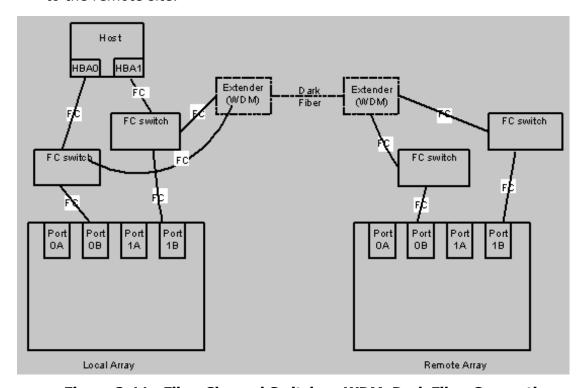


Figure 2-11: Fibre Channel Switches, WDM, Dark Fibre Connection

Recommendations

- Two remote paths are recommended between local and remote arrays.
 In the event of path failure, data copying is automatically shifted to the alternate path.
- WDM has the same speed as fibre channel; however, response time increases when distance between sites increases.

For more information on WDM, see Appendix F, Wavelength Division Multiplexing (WDM) and dark fibre.

Path and switch performance

Performance guidelines for two or more paths between hosts, switches, and arrays are shown in Table 2-5.

Table 2-5: Performance Guidelines for Paths and Switches

Array Mode	Switch				
Array Mode		4 G bps	2 G bps	1 G bps	
Auto Mode	One path per controller between the array and a switch is sufficient for both host I/O and TrueCopy operations. (Shown in Figure 2-10.)	Same as 8 G bps/Auto Mode.	Same as 8 G bps/Auto Mode.	Two ports between the array and a switch must be provided. Or, one switch/port must be provided for host I/O, and one switch/port for TrueCopy.	
8 G bps Mode			Not available	Not available	
4 G bps Mode			Not available	Not available	
2 G bps Mode			Same as 8 G bps/Auto Mode	Not available	
1 G bps Mode	Two ports between the array and a switch must be provided. Or, one switch/port must be provided for host I/O, and one switch/port for TrueCopy. For the configuration shown in Figure 2-10 two paths per switch are recommended.	Same as 8G bps/1 G bps.	Same as 8G bps/1 G bps.	Same as 8G bps/1 G bps.	

Port transfer rate for Fibre channel

The communication speed of the fibre channel port on the array must match the speed specified on the host port. These two ports—fibre channel port on the array and host port—are connected via fibre channel cables. Each port on the array must be set separately.

Table 2-6: Setting Port Transfer Rates

	If the host port is set to	Set the array port to
	1 G bps	1 G bps
Manual mode	2 G bps	2 G bps
Manual Inoue	4 G bps	4 G bps
	8 G bps	8 G bps
	2 G bps	Auto, with max of 2 G bps
Auto mode	4 G bps	Auto, with max of 4 G bps
	8 G bps	Auto, with max of 8 G bps

 When using a direct connection, Auto mode may cause blockage of the data path. In this case, change the transfer rate of Manual mode. Maximum speed is ensured using the manual settings.

Specify port transfer rate in Navigator 2 GUI, on the **Edit FC Port** screen (Settings/FC Settings/port/Edit Port button).



NOTE: Do not modify transfer rate while the array power is on, if your remote path is a direct connection. This causes remote path blockage.

Find details on communication settings in the *Hitachi AMS 2100/2300 Storage System Hardware Guide*.

Remote path configurations for iSCSI

The iSCSI remote path can be set up in the following configurations:

- Direct connection
- Local Area Network (LAN) switch connections
- Wide Area Network (WAN) connections
- WAN Optimization Controller (WOC) connections

Direct iSCSI connection

Figure 2-12, illustrates two remote paths directly connecting the local and remote arrays. Direct connections are used when the local and remote arrays are set up at the same site. In this case, the arrays can be linked with category 5e9 or 6 copper LAN cable.

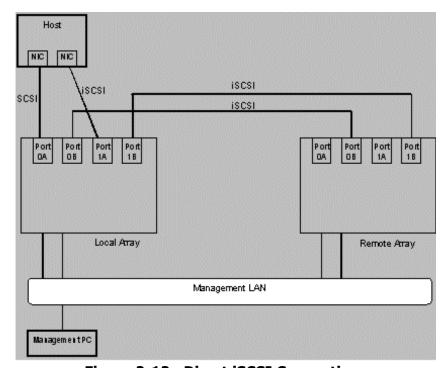


Figure 2-12: Direct iSCSI Connection

One path between arrays is required.

One path between the host and local array is required.

Recommendations

- Though one path is required, two paths are recommended from host to array and between local and remote arrays. This provides a backup path in the event of path failure.
- When a large amount of data is to be copied to the remote site, the initial copy between local side and remote systems may be performed at the same location. In this case, category 5e9 or 6 copper LAN cable is recommended.

Single LAN switch, WAN connection

Figure 2-13, illustrates two remote paths using one LAN switch and network to the remote array.

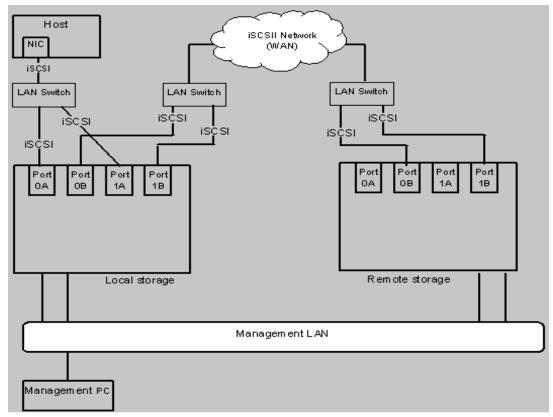


Figure 2-13: Single-Switch Connection

- This configuration is not recommended because a failure in a LAN switch or WAN would halt operations.
- Separate LAN switches and paths should be used for host-to-array and array-to-array, for improved performance.

Multiple LAN switch, WAN connection

Figure 2-14, illustrates two remote paths using multiple LAN switches and WANs to make the connection to the remote site.

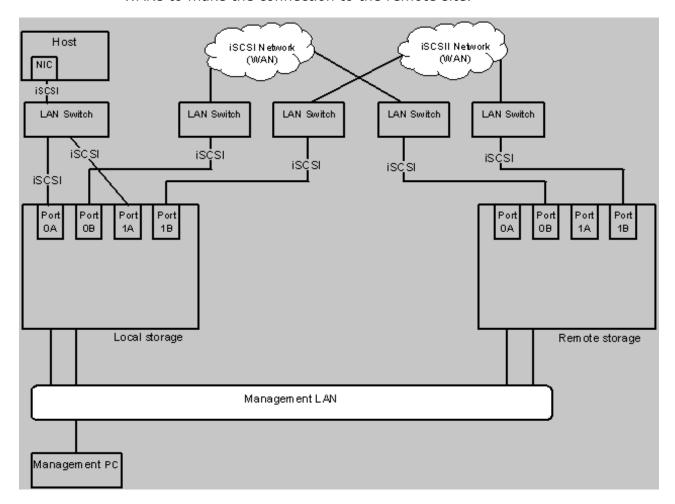


Figure 2-14: Multiple-Switch and WAN Connection

Recommendations

• Separate LAN switches and paths should be used for host-to-array and array-to-array for better performance, and to provide a backup.

Single LAN switch, WOC, WAN connection

WOCs may be required for TrueCopy, depending on your system's bandwidth, latency, and so on. Use of a WOC improves performance. See WAN Optimization Controller (WOC) requirements on page 2-10 for more information.

Figure 2-15, illustrates two remote paths using a single LAN switch, WOC, and WAN to make the connection to the remote site.

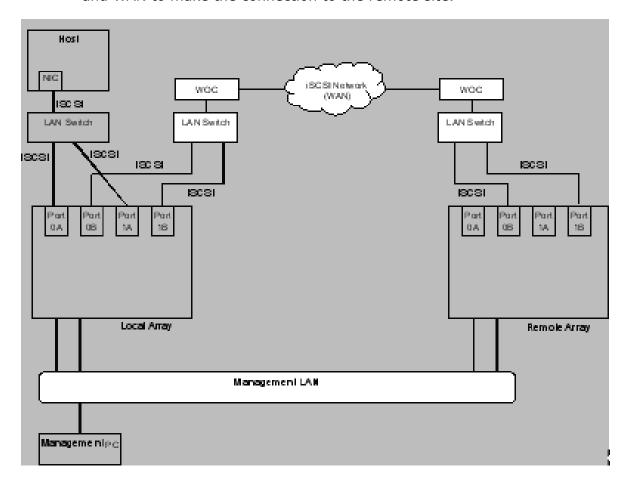


Figure 2-15: Single Switch, WOC, and WAN Connection

Recommendations

Two remote paths are recommended between local and remote arrays.
 In the event of path failure, data copying is automatically shifted to the alternate path.

Multiple LAN switch, WOC, WAN connection

Figure 2-16, illustrates two remote connections using multiple LAN switches, WOCs, and WANs to make the connection to the remote site.

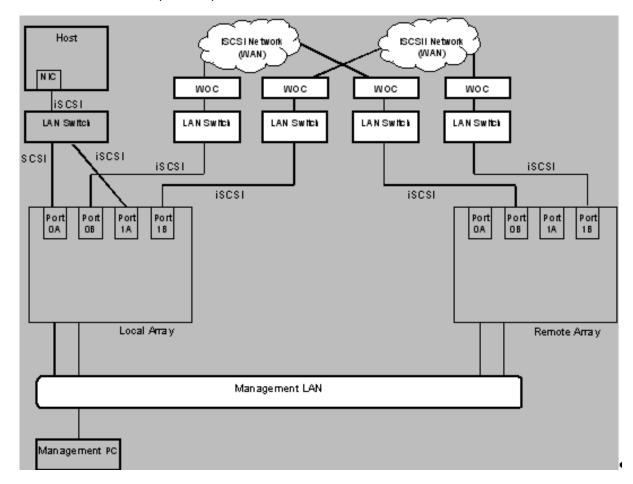


Figure 2-16: Connection Using Multiple Switch, WOC, WAN

- If a Gigabit Ethernet port (1000BASE-T) is provided on the WOC, the LAN switch to the WOC is not required. Connect array ports 0B and 1B to the WOC directly. If your WOC does not have 1Gbps ports, the LAN switch is required.
- Two paths are recommended; Using separate LAN switch, WOC and WAN for each remote path ensures that data copy automatically continues on the second path in the event of a path failure.

Multiple array, LAN switch, WOC connection with single WAN

Figure 2-17, shows two local arrays connected to two remote arrays, each via a LAN switch and WOC.

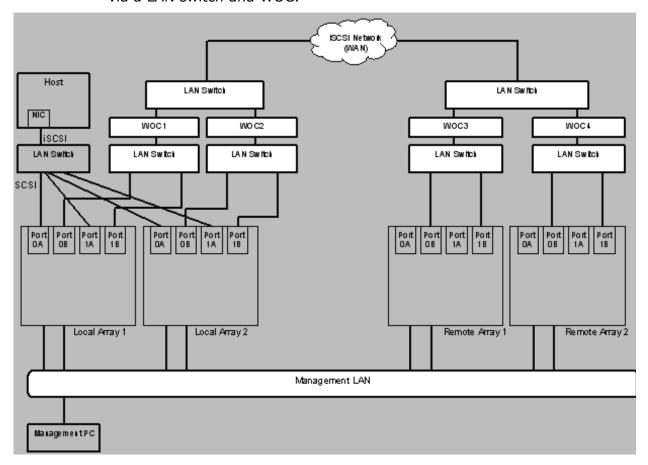


Figure 2-17: Multiple Array Connection Using Single WAN

- If a Gigabit Ethernet port (1000BASE-T) is provided on the WOC, the LAN switch to the WOC is not required. Connect array ports 0B and 1B to the WOC directly. If your WOC does not have 1Gbps ports, the LAN switch is required.
- You can reduce the number of switches by using a switch with VLAN capability. If a VLAN switch is used, port 0B of local array 1 and WOC1 should be in one LAN (VLAN1); port 0B of local array 2 and WOC3 should be in another LAN (VLAN2). Connect the VLAN2 port directly to Port 0B of the local array 2 and WOC3.

Multiple array, LAN switch, WOC connection with two WANs

Figure 2-18, shows two local arrays connected to two remote arrays, each via two LAN switches, WANs, and WOCs.

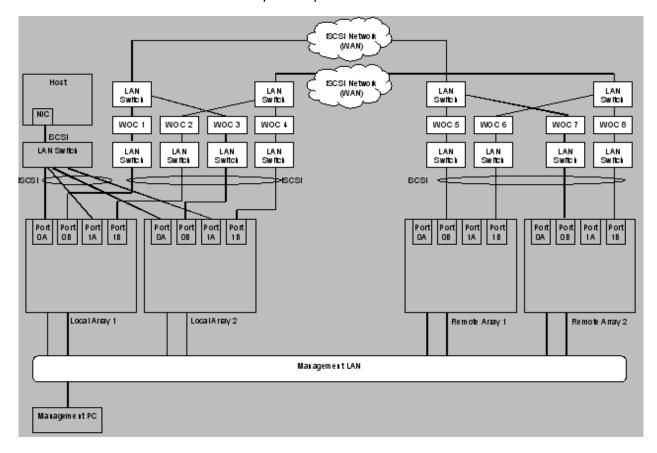


Figure 2-18: Multiple Array Connection Using Two WANs

- Two remote paths are recommended between local and remote arrays.
 In the event of path failure, data copying is automatically shifted to the alternate path.
- If a Gigabit Ethernet port (1000BASE-T) is provided on the WOC, the LAN switch to the WOC is not required. Connect array ports 0B and 1B to the WOC directly. If your WOC does not have 1Gbps ports, the LAN switch is required.
- You can reduce the number of switches by using a switch with VLAN capability. If a VLAN switch is used, port 0B of local array 1 and WOC1 should be in one LAN (VLAN1); port 0B of local array 2 and WOC3 should be in another LAN (VLAN2). Connect the VLAN2 port directly to Port 0B of the local array 2 and WOC3.

Using the remote path — best practices

The following best practices are provided to reduce and eliminate path failure.

- If both arrays are powered off, power-on the remote array first.
- When powering down both arrays, turn off the local array first.
- Before powering off the remote array, change pair status to Split. In Paired or Synchronizing status, a power-off results in Failure status on the remote array.
- If the remote array is not available during normal operations, a blockage error results with a notice regarding SNMP Agent Support Function and TRAP. In this case, follow instructions in the notice.
 - Path blockage automatically recovers after restarting. If the path blockage is not recovered when the array is READY, contact Hitachi Customer Support.
- Power off the arrays before performing the following operations:
 - Changing the micro-code program (firmware)
 - Setting or changing the fibre transfer rate

Remote processing

- Consideration for remote processing:
 - When a write I/O instruction received at a local site is executed at a remote site synchronously, performance attained at the remote site directly affects performance that is attained at the local site. The performance attained at the local site or the system is lowered when the remote site is overloaded, due to a large number of updates, etc.
 - Therefore, carefully monitor the load on the remote site as well as the local site.
- When using DP-VOL created by Dynamic Provisioning for the P-VOL and S-VOL of TrueCopy and executing I/O to the P-VOL when the pair status is Synchronizing or Paired, check that there is enough free capacity in the DP pool to which the S-VOL belongs (entire capacity of DP pool x progress of formatting - consumed capacity), and then execute it.
 - Although the format status of the belonging DP pool is formatting, if the DP-VOL formatting is completed, the DP-VOL created by Dynamic Provisioning can create a pair. If the pair status is Synchronizing or Paired and dual writing is executed in the S-VOL, a new area may be required for the S-VOL. However, if the required area cannot be secured in the DP pool, it must wait until the DP pool formatting is progressed, and the I/O performance may be extremely deteriorated due to the waiting time.

Planning for TrueCopy

Planning a TrueCopy system requires an understanding of the components that are used in the remote backup environment and awareness of their requirements, restrictions, and recommendations.

- The planning workflow
- Planning arrays
- Planning Logical Units
- Operating system recommendations and restrictions
- Calculating supported capacity
- ☐ On cascading TrueCopy with local replication software
- Concurrent use of Dynamic Provisioning

The planning workflow

Implementing a TrueCopy system requires setting up local and remote arrays, TrueCopy volumes, the remote path that connects the volumes, and interface(s).

A planning workflow can be organized in the following manner:

- Planning arrays for TrueCopy.
- Planning volume setup, which consists of:
 - Understanding TrueCopy primary and secondary volume specifications, recommendations, and restrictions,
 - Understanding how to use unified LUs to create P-VOLs and S-VOLs (optional)
 - Cascading TrueCopy volumes with ShadowImage or SnapShot volumes (optional).
 - Specifications, recommendations, and restrictions for DMLUs
 - Specifications, recommendations, and restrictions for command devices (only required if CCI is used)
- Planning remote path connections, which includes:
 - Reviewing supported path configurations (covered in Chapter 2)
 - Measuring write workload to determine the bandwidth that is required (covered in Chapter 2)

Planning arrays

Logical units on AMS 2100/2300/2500 arrays can be paired with logical units on WMS 100, AMS 200, AMS 500, and AMS 1000 arrays. Any combination of array may be used on the local and remote sides. When using the earlier model arrays, please observe the following:

- The maximum number of pairs between different model arrays is limited to the maximum number of pairs supported by the smallest array.
- The firmware version for WMS 100, AMS 200, AMS 500, or AMS 1000 must be 0780/A or later when connecting with AMS 2100, 2300, or 2500 if the hardware Rev. is 0100.
- The firmware version for a WMS 100, AMS 200, AMS 500, or AMS 1000 must be 0786/A or later when connecting with an AMS 2100, AMS 2300, or AMS 2500 if the H/W Rev. is 0200.
- The bandwidth of the remote path to a WMS 100, AMS 200, AMS 500, or AMS 1000 must be 20 Mbps or more.
- The pair operations for WMS 100, AMS 200, AMS 500, or AMS 1000 cannot be performed using the Navigator 2 GUI.
- WMS 100, AMS 200, AMS 500, or AMS 1000 cannot use functions that are newly supported by AMS 2100, 2300, or 2500.

Planning Logical Units

Please review the recommendations in the following subsections before setting up TrueCopy volumes.

Also, review:

- System requirements on page 4-2
- TrueCopy specifications on page A-2

Volume pair recommendations

- Because data written to a primary volume is also written to a secondary volume at a remote site synchronously, performance is impacted according to the distance to the remote site. Assigning LUs to primary or secondary volumes should be limited to those not required to return a quick response to a host.
- The number of volumes within the same RAID group should be limited. Pair creation or resynchronization for one of the volumes may impact I/O performance for the others because of contention between drives. When creating two or more pairs within the same RAID group, standardize the controllers for the LUs in the RAID group. Also, perform pair creation and resynchronization when I/O to other volumes in the RAID group is low.
- Assign primary volumes to a RAID group consisting of SAS drives or SSD drives. Host I/O performance is lowered during TrueCopy operations when the P-VOL is in a RAID group consisting of SATA drives.
- Assign an LU consisting of four or more data disks, otherwise host and/ or copying performance may be lowered.
- When using SATA drives, make the data disks between 4D and 6D, otherwise performance is lowered during copying.
- LUs used for pair volumes should have stripe size of 64 kB and segment size of 16 kB.
- When TrueCopy and SnapShot are cascaded, the SnapShot data pool activity influences host performance and TrueCopy copying. Assign volumes to data pools consisting of SAS drives or SSD drives and of four or more disks for best results.
- Limit the I/O load on both local and remote arrays to maximize performance. Performance on the remote array affects performance on both the local system and the synchronization of volumes.
- Synchronize Cache Execution Mode must be turned off on the remote array to prevent possible data path failure.

LUN expansion

A unified LU can be used as a TrueCopy P-VOL or S-VOL. When using TrueCopy with LUN Expansion, please observer the following:

 P-VOL and S-VOL capacities must be equal, though the number of LUs composing them (unified LUs) may differ, as shown in Figure 3-1.

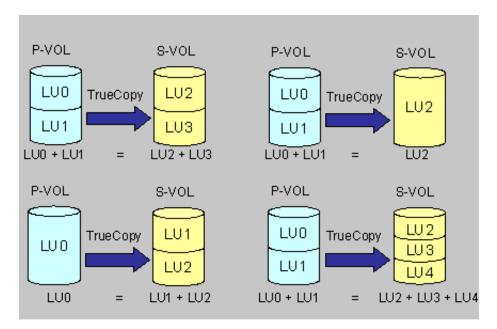


Figure 3-1: Capacity Same, Number of LUs Different

• A unified LU composed of 128 or more LUs cannot be assigned to the P-VOL or S-VOL, as shown in Figure 3-2.

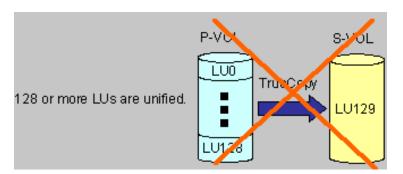


Figure 3-2: Number of LUs Restricted

 P-VOLs and S-VOLs made of unified LU(s) can be assigned to different RAID levels and have a different number of data disks, as shown in Figure 3-3.

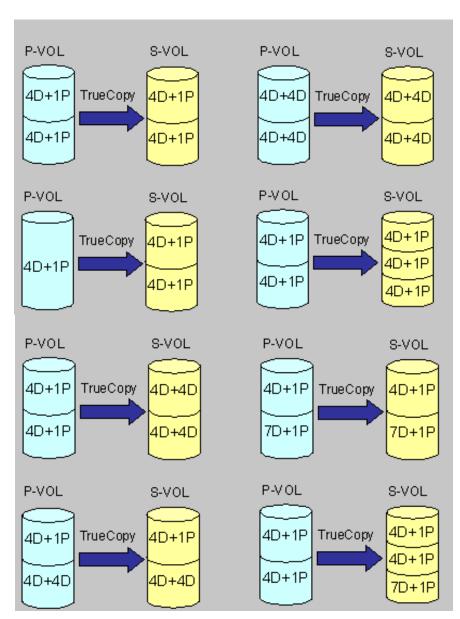


Figure 3-3: Combination of RAID Levels

- A TrueCopy P-VOL or S-VOL cannot be used to compose a unified LU.
- The LUs created in the RAID group with different drive types cannot be unified, as shown in Figure 3-4. Unify LUs consisting of drives of the same type.

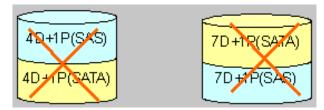


Figure 3-4: Unifying LUs of Different Drive Type not Supported

Operating system recommendations and restrictions

The following sections provide operating system recommendations and restrictions.

Host time-out

I/O time-out from the host to the array should be more than 60 seconds. You can figure I/O time-out by increasing the remote path time limit times 6. For example, if the remote path time-out value is 27 seconds, set host I/O time-out to 162 seconds (27×6) or more.

P-VOL, S-VOL recognition by same host on VxVM, AIX®, LVM

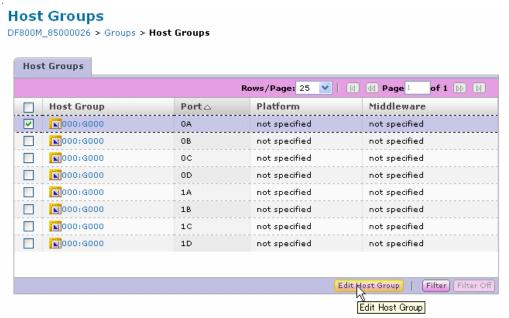
VxVM, AIX[®], and LVM do not operate properly when both the P-VOL and S-VOL are set up to be recognized by the same host. The P-VOL should be recognized one host on these platforms, and the S-VOL recognized by another.

Setting the Host Group options

When MC/Service Guard is used on HP server, connect the host group (fibre channel) or iSCSI Target to HP server as follows:

For fibre channel interfaces

- 1. In the Navigator 2 GUI, access the array and click **Host Groups** in the Groups tree view. The Host Groups screen displays.
- 2. Click the check box for the **Host Group** that you want to connect to the HP server.
- 3. Click **Edit Host Group**. The Edit Host Group screen appears.



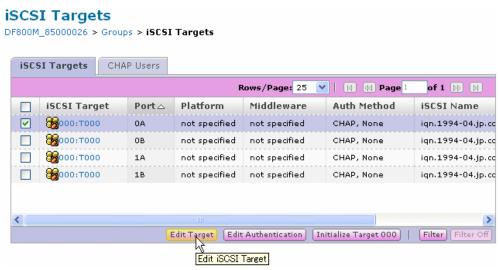
4. Select the **Options** tab.

- 5. From the Platform drop-down list, select **HP-UX**. Doing this causes "Enable HP-UX Mode" and "Enable PSUE Read Reject Mode" to be selected in the Additional Setting box.
- 6. Click OK. A message appears, click Close.



For iSCSI interfaces

- 1. In the Navigator 2 GUI, access the array and click **iSCSI Targets** in the Groups tree view. The iSCSI Targets screen displays.
- 2. Click the check box for the iSCSI Targets that you want to connect to the HP server.
- 3. Click **Edit Target**. The Edit iSCSI Target screen appears.



- 4. Select the **Options** tab.
- 5. From the Platform drop-down list, select **HP-UX**. Doing this causes "Enable HP-UX Mode" and "Enable PSUE Read Reject Mode" to be selected in the Additional Setting box.
- 6. Click OK. A message appears, click Close.

Windows 2000 Servers

- A P-VOL and S-VOL cannot be made into a dynamic disk on Windows 2000 Server.
- Native OS mount/dismount commands can be used for all platforms, except Windows 2000 Server and Windows Server 2008. The native commands on this environment do not guarantee that all data buffers are completely flushed to the volume when dismounting. In these instances, you must use Hitachi's Command Control Interface (CCI) to perform volume mount/unmount operations. See the Hitachi

AMS 2000 Family Command Control Interface Reference Guide for more information.

Windows Server 2003/2008

- A P-VOL and S-VOL cannot be made into a dynamic disk on Windows Server 2008.
- Windows Server 2008 does not operate properly when both the P-VOL and S-VOL are set up to be recognized by the same host. The P-VOL should be recognized one host, and the S-VOL recognized by another.
- Native OS mount/dismount commands can be used for all platforms, except Windows 2000 Server and Windows Server 2008. The native commands on this environment do not guarantee that all data buffers are completely flushed to the volume when dismounting. In these instances, you must use Hitachi's Command Control Interface (CCI) to perform volume mount/unmount operations. See the Hitachi AMS 2000 Family Command Control Interface Reference Guide for more information.
- (CCI only) When mounting a volume, use Volume{GUID} as an argument of the CCI mount command (if used for the operation). The Volume{GUID} can be used in CCI versions 01-23-03/08 and later.
- (CCI only) When describing a command device in the configuration definition file, specify it as Volume{GUID}.
- (CCI only) If a path detachment is caused by controller detachment or fibre channel failure, and the detachment continues for longer than one minute, the command device may not be recognized when recovery occurs. In this case, execute the "re-scanning of the disks" in Windows. If Windows cannot access the command device, though CCI recognizes the command device, restart CCI.

Identifying P-VOL and S-VOL LUs on Windows

In Navigator 2, the P-VOL and S-VOL are identified by their LU number. In Windows Server 2003 and Windows Server 2008, LUs are identified by HLUN. To map LUN to HLUN on Windows, proceed as follows. These instructions provide procedures for iSCSI and fibre channel fibre channel interfaces.

- 1. Identify the HLUN of your Windows disk.
 - a. From the Windows Server 2003 Control Panel, select Computer Management>Disk Administrator.
 - b. Right-click the disk whose HLUN you want to know, then select **Properties**. The number displayed to the right of "LUN" in the dialog window is the HLUN.
- 2. Identify HLUN-to-LUN Mapping for the iSCSI interface as follows. (If using fibre channel, skip to Step 3.)
 - a. In the Navigator 2 GUI, select the desired array.
 - b. In the array tree that displays, click the **Group** icon, then click the **iSCSI Targets** icon in the Groups tree.
 - c. On the iSCSI Target screen, select an iSCSI target.

- d. On the target screen, select the **Logical Units** tab. Find the identified HLUN. The LUN displays in the next column.
- e. If the HLUN is not present on a target screen, on the iSCSI Target screen, select another iSCSI target and repeat Step d.
- 3. Identify HLUN-to-LUN Mapping for the fibre channel interface, as follows:
 - a. In Navigator 2, select the desired array.
 - b. In the array tree that displays, click the **Groups** icon, then click the **Host Groups** icon in the Groups tree.
 - c. On the Host Groups screen, select a Host group.
 - d. On the host group screen, select the **Logical Units** tab. Find the identified H-LUN. The LUN displays in the next column.
 - e. If the HLUN is not present on a host group target screen, on the Host Groups screen, select another Host group and repeat Step d.

Dynamic Disk with Windows 2003

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

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

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

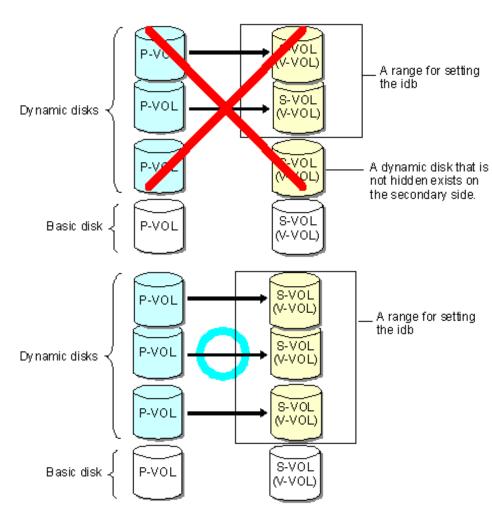


Figure 3-5: Supported, un-Supported Dynamic Disk Configurations

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

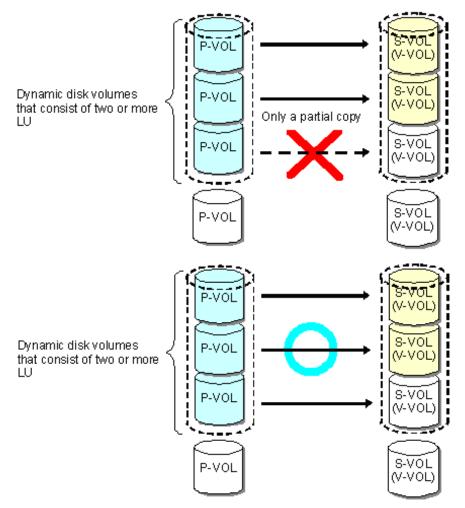


Figure 3-6: Supported Copying for Dynamic Disk Volumes

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

VMWare and TrueCopy configuration

When creating a backup of the virtual disk in the vmfs format using TrueCopy, shutdown the virtual machine which accesses the virtual disk, and slit the pair.

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

Concurrent use of Dynamic Provisioning

When the array firmware version is 0893/A or more, the DP-VOLs created by Dynamic Provisioning can be set for a P-VOL or an S-VOL in TrueCopy.

This section describes the points to remember when using TrueCopy and Dynamic Provisioning together. Refer to the *Hitachi Adaptable Modular Storage Dynamic Provisioning User's Guide* for detailed information about Dynamic Provisioning. Hereinafter, the LU created in the RAID group is called a normal LU, and the LU created in the DP pool that is created by Dynamic Provisioning is called a DP-VOL.

LU type that can be set for a P-VOL or an S-VOL of TrueCopy
 The DP-VOL created by Dynamic Provisioning can be used for a P-VOL or an S-VOL of TrueCopy. Table 3-1 shows a combination of a DP-VOL and a normal LU that can be used for a P-VOL or an S-VOL of TrueCopy.

Table 3-1: Combination of a DP-VOL and a Normal LU

TrueCopy P-VOL	TrueCopy S-VOL	Contents
DP-VOL	DP-VOL	Available. The P-VOL and S-VOL capacity can be reduced compared to the normal LU.
DP-VOL	Normal LU	Available.
Normal LU	DP-VOL	Available. When the pair status is Split, the S-VOL capacity can be reduced compared to the normal LU by deleting 0 data.

Pair status at the time of DP pool capacity depletion
 When the DP pool is depleted after operating the Truecopy pair that uses the DP-VOL created by Dynamic Provisioning, the pair status of the pair concerned may be a Failure. Table 3-2 shows the pair statuses before and after the DP pool capacity depletion. When the pair status becomes a Failure caused by the DP pool capacity depletion, add the DP pool capacity whose capacity is depleted, and execute the pair operation again.

Table 3-2: Pair Statuses before the DP Pool Capacity Depletion and Pair Statuses after the DP Pool Capacity Depletion

Pair Statuses before the DP Pool Capacity Depletion	Pair Statuses after the DP Pool Capacity Depletion belonging to P-VOL	Pair Statuses after the DP Pool Capacity Depletion belonging to data pool
Simplex	Simplex	Simplex
Synchronizing	Synchronizing Failure (See Note)	Failure
Paired	Paired Failure (See Note)	Failure
Split	Split	Split
Failure	Failure	Failure



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

DP pool status and availability of pair operation

When using the DP-VOL created by Dynamic Provisioning for a P-VOLor an S-VOL or a data pool of the TrueCopy pair, the pair operation may not be executed depending on the status of the DP pool to which the DP-VOL belongs. Table 3-3 shows the DP pool status and availability of the TrueCopy pair operation. When the pair operation fails due to the DP pool status, correct the DP pool status and execute the pair operation again.

Table 3-3: DP Pool for P-VOL Statuses and Availability of Pair Operation

	DP Pool Statuses, DP Pool Capacity Statuses, and Statuses				DP Pool Optimization	
Pair Operation	Normal	Capacity in Growth	Capacity Depletion	Regressed	Blocked	DP in Optimization
Create pair	0 (1)	0	0 (1) (2)	0	х	0
Split pair	0	0	0	0	0	0
Resync pair	0 (1)	0	0 (1) (2)	0	х	0
Swap pair	0 (2)	0	0 (2)	0	0	0
Delete pair	0	0	0	0	0	0

- (1) Refer to the status of the DP pool to which the DP-VOL of the S-VOL belongs. If the status exceeds the DP pool capacity belonging to the S-VOL by the pair operation, the pair operation cannot be performed.
- (2) Refer to the status of the DP pool to which the DP-VOL of the P-VOL belongs. If the status exceeds the DP pool capacity belonging to the P-VOL by the pair operation, the pair operation cannot be performed.
- (3) When the DP pool was created or the capacity was added, the formatting operates for the DP pool. If pair creation, pair resynchronization, or swapping is performed during the formatting, depletion of the usable capacity may occur. Since the formatting progress is displayed when checking the DP pool status, check if the sufficient usable capacity is secured according to the formatting progress, and then start the operation.
- Operation of the DP-VOL during TrueCopy use
 When using the DP-VOL created by Dynamic Provisioning for a P-VOL or an S-VOL of TrueCopy, any of the operations among the capacity growing, capacity shrinking, and LU deletion of the DP-VOL in use cannot be executed. To execute the operation, delete the TrueCopy pair of which the DP-VOL to be operated is in use, and then execute it again.
- Operation of the DP pool during TrueCopy use

When using the DP-VOL created by Dynamic Provisioning for a P-VOL or an S-VOL of TrueCopy, the DP pool to which the DP-VOL in use belongs cannot be deleted. To execute the operation, delete the TrueCopy pair of which the DP-VOL is in use belonging to the DP pool to be operated, and then execute it again. The attribute edit and capacity addition of the DP pool can be executed usually regardless of the TrueCopy pair.

Availability of TrueCopy pair creation between different firmware versions

When the firmware versions of the array differ on the local side and the remote side, if the firmware version of the array including the DP-VOL is 0893/A or more, a TrueCopy pair can be created (Figure 3-7). Table 3-4 shows pair creation availability when the firmware array version is 0893/A or more, or less than 0893/A (including WMS100/AMS200/500/1000).

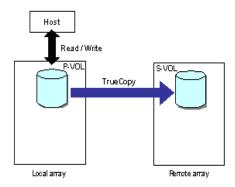


Figure 3-7: TrueCopy pair creation between different firmware versions

Table 3-4: Availability of TrueCopy Pair Creation between Different Firmware Versions

P-VOL S-VOL		Local: 0893/A or more	Local: 0893/A or more	Local: Less than 0893/A
		Remote: 0893/A or more	Remote: Less than 0893/A	Remote: 0893/A or more
Normal LU	Normal LU	0	0	0
Normal LU	DP-VOL	0	Х	0
DP-VOL	Normal LU	0	0	х
DP-VOL	DP-VOL	0	Х	Х

Calculating supported capacity

The array manages capacity in segments of 15.75 GB. Because of this, it is necessary to calculate your TrueCopy system's managed capacity and compare it to the maximum supported capacity.

Calculate managed capacity using the following formula. single maximum supported capacity is shown in Table 3-5.

TrueCopy: Maximum supported capacity value of P-VOL and S-VOL (TB)

- = Single Maximum Supported Capacity
- (Total ShadowImage S-VOL capacity)
- 17 x (Total SnapShot P-VOL capacity)

Table 3-5: Single Maximum Supported Capacity per Controller

Array Type	Mounted Memory Capacity (GB/CTL)	Single Maximum Supported Capacity (TB)
AMS2100	1	261
	2	748
	4	897
AMS 2300	1	Not supported
	2	673
	4	1,870
	8	3,740
AMS 2500	2	486
	4	1,870
	6	2,618
	8	3,366
	10	4,488
	12	5,610
	16	6,733

To calculate TrueCopy capacity

- 1. Add the sizes of all TrueCopy P-VOLs and S-VOLs.
- 2. If SnapShot is used, add 17 x the total SnapShot P-VOL size to TrueCopy's size total.
- 3. If ShadowImage is used, add total S-VOL size to TrueCopy's size total.
- 4. Add all totals. The value must be less than the single maximum supported capacity for your array/memory capacity, as shown in Table 3-5.

On cascading TrueCopy with local replication software

TrueCopy can be cascaded with ShadowImage or SnapShot. It cannot be used with TrueCopy Extended Distance.

A cascaded system is a configuration in which a TrueCopy P-VOL or S-VOL is shared with a ShadowImage or SnapShot P-VOL or S-VOL. Cascading with ShadowImage reduces performance impact to the TrueCopy host and provides data for use with secondary applications.

Figure 3-8 and Figure 3-9 show basic cascade examples.

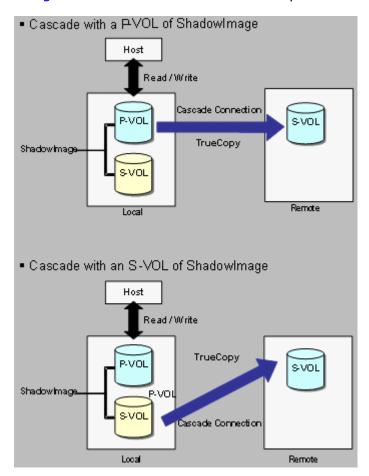


Figure 3-8: Examples of ShadowImage Cascade

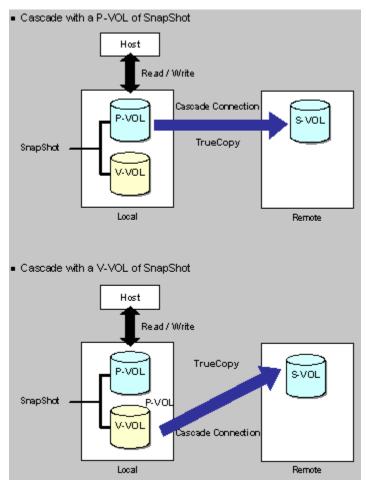


Figure 3-9: Examples of SnapShot Cascade

Many but not all configurations, operations, and statuses between TrueCopy and ShadowImage or SnapShot are supported. See Appendix D, Cascading with ShadowImage, and Appendix E, Cascading with SnapShot, for detailed information.

Cascade connection restrictions

Cascade restrictions with TrueCopy
 Cascade connection with TrueCopy is not available.

 Figure 3-10 on page 3-18 illustrates the Cascade connection restrictions.

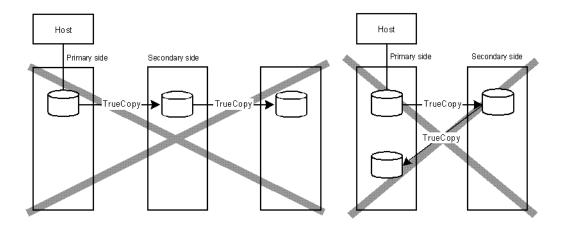


Figure 3-10: Cascade Connection Restrictions

- Cascade restrictions with data pool of SnapShot
 TrueCopy pair cannot be created using a data pool.
- Cascade restrictions with ShadowImage pair including the DP-VOLs creating by Dynamic Provisioning

When the firmware array version is less than 0893/A, the DP-VOLs created by Dynamic Provisioning cannot be set for a P-VOL or an S-VOL in TrueCopy. Therefore, a TrueCopy pair cannot be created for a ShadowImage pair including a DP-VOL.

When the firmware array version is 0893/A or more (DP-VOL side), the above-mentioned restrictions do not apply. When the firmware array version including the DP-VOL (Figure 3-11) is 0893/A or more, can be cascaded.

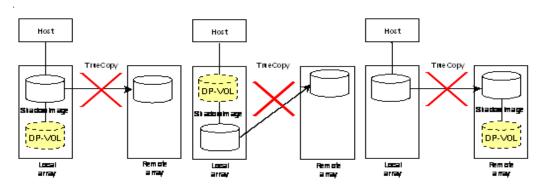


Figure 3-11: Cascade Connection Restrictions including DP-VOL Creating by Dynamic Provisioning



Requirements

This chapter describes minimum operational requirements.

■ System requirements

System requirements

The minimum requirements for TrueCopy are listed below.

• Firmware: Version 0832/B or later is required for an AMS2100 or AMS2300 array if the hardware Rev. is 0100.

Version 0840/A or later is required for an AMS2500 array if the hardware Rev. is 0100.

Version 0890/A or later is required for AMS2100, AMS2300, and AMS2500 if the hardware Rev. is 0200.

When connecting an array where the hardware Rev. is 0200, firmware version 0890/A or later is required on both local side and remote side arrays.

 Storage Navigator Modular 2: Version 3.21 or later is required for the management PC for an AMS2100 or AMS2300 array if the hardware Rev. is 0100.

Version 4.00 or later is required for management PC for AMS2500 array if the hardware Rev. is 0100.

Version 9.00 or later is required for management PC for AMS2100/AMS2300/AMS2500 if the hardware Rev. is 0200.

- Two license keys required: one local, one remote
- Number of controllers: 2 (dual configuration)
- Volume size: S-VOL block count must be equal to P-VOL block count.
- CCI version: 01-21-03/06 or later is required for Windows host only
- Command devices: Max. 128. The command device is required only when CCI is used. The command device volume size must be greater than or equal to 33 MB.
- Bandwidth for the remote path to WMS100, AMS200, AMS500, or AMS1000 must be 20 Mbps or more.

See Appendix A, Specifications for additional information.

Displaying the hardware revision

The hardware revision (Rev.) can be displayed when an individual array is selected from the Arrays list using Navigator 2, version 9.00 or later, as shown in Figure 4-1.

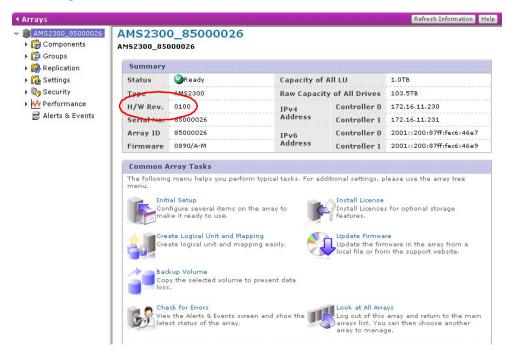


Figure 4-1: Locating the hardware revision

Requirements 4–3

Installation and setup

This chapter provides procedures for installing and setting up TrueCopy using the Navigator 2 GUI. CLI and CCI instructions are included in this manual in the appendixes.

- Installation procedures
- Setup procedures

Installation procedures

TrueCopy is an extra-cost option; it must be installed and enabled on the local and remote arrays.

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

The following sections provide instructions for installing, enabling/disabling, and uninstalling TrueCopy.

Installation

Prerequisites

- A key code or key file is required to install or uninstall TrueCopy. If you
 do not have the key file or code, you can obtain it from the download
 page on the HDS Support Portal, http://support.hds.com.
- TrueCopy cannot be installed if more than 239 hosts are connected to a port on the array.

To install TrueCopy

- 1. In the Navigator 2 GUI, click the check box for the array where you want to install TrueCopy, then click the **Show & Configure Array** button.
- 2. Under Common Array Tasks, click **Install License**. The Install License screen displays.
- 3. Select the **Key File** or **Key Code** radio button, then enter the file name or key code. You may browse for the Key File.
- 4. Click OK.
- 5. Click **Confirm** on the subsequent screen to proceed, then click **Close** on the installation complete message.

Enabling, disabling

TrueCopy is automatically enabled when it is installed. You can disable and re-enable it.

Prerequisites

Before disabling TrueCopy:

- TrueCopy pairs must be deleted and the status of the logical units must be Simplex.
- The remote path must be deleted.
- TrueCopy cannot be enabled if more than 239 hosts are connected to a port on the array.

To enable/disable TrueCopy

- 1. In the Navigator 2 GUI, click the check box for the array, then click the **Show & Configure Array** button.
- 2. In the tree view, click **Settings**, then click **Licenses**.

- 3. Select **TrueCopy** in the Licenses list.
- 4. Click **Change Status**. The Change License screen displays.
- To disable, clear the **Enable: Yes** check box.To enable, check the **Enable: Yes** check box.
- 6. Click OK.
- 7. A message appears confirming that TrueCopy is disabled. Click **Close**.

Uninstalling TrueCopy

Prerequisite

- TrueCopy pairs must be deleted. Volume status must be Simplex.
- A key code is required. If you do not have the key code, you can obtain it from the download page on the HDS Support Portal, http:// support.hds.com.

To uninstall TrueCopy

- 1. In the Navigator 2 GUI, click the check box for the array, then click the **Show & Configure Array** button.
- 2. In the navigation tree, click **Settings**, then click **Licenses**.
- 3. On the Licenses screen, select **TrueCopy** in the Licenses list and click the **De-install License** button.
- 4. On the De-Install License screen, enter the file or code in the **Key File** or **Key Code** box, and then click **OK**.
- 5. On the confirmation screen, click **Close**.

Setup procedures

The following sections provide instructions for setting up the DMLU and Remote Path.

(For CCI users, TrueCopy/CCI setup includes configuring the command device, the configuration definition file, the environment variable, and LU mapping. See Appendix C, Operations using CCI for instructions.)

Defining the DMLU

The DMLU (differential management-logical unit) must be set up prior to using TrueCopy. The DMLU is used by the system for storing TrueCopy information when the array is powered down.

Prerequisites

- The logical unit used for the DMLU must be set up and formatted.
- The logical unit used for the DMLU must be at least 10 GB.
- DMLUs must be set up on both the local and remote the arrays.
- One DMLU is required on each array; two are recommended, the second used as backup. No more than two DMLUs can be installed.
- Please also review specifications for DMLUs in Appendix A, Specifications.

To define the DMLU

- 1. In the Navigator 2 GUI, select the array where you want to set up the DMLU.
- 1. In the navigation tree, click **Settings**, then click **DMLU**. The DMLU screen displays.
- 2. Click **Add DMLU**. The Add DMLU screen displays.
- 3. Select the **LUN(s)** that you want to assign as DMLUs, and then click **OK**. A confirmation message displays.
- 4. Select the **Yes, I have read ...** check box, then click **Confirm**. When a success message displays, click **Close**.

Adding or changing the remote port CHAP secret

(For iSCSI only)

Challenge-Handshake Authentication Protocol (CHAP) provides a level of security at the time that a link is established between the local and remote arrays. Authentication is based on a shared secret that validates the identity of the remote path. The CHAP secret is shared between the local and remote arrays.

Prerequisites

• The array IDs for local and remote arrays are required.

To add a CHAP secret

This procedure is used to add CHAP authentication manually on the remote array.

- On the remote array, navigate down the GUI tree view to Replication/ Setup/Remote Path. The Remote Path screen displays. (Though you may have a remote path set, it does not show up on the remote array. Remote paths are set from the local array.)
- 2. Click the **Remote Port CHAP** tab. The Remote Port CHAP screen displays.
- 3. Click the **Add Remote Port CHAP** button. The Add Remote Port CHAP screen displays.
- 4. Enter the Local Array ID.
- 5. Enter CHAP **Secrets** for Remote Path 0 and Remote Path 1, following onscreen instructions.
- 6. Click **OK** when finished.

To change a CHAP secret

- 1. Split the TrueCopy pairs, after confirming first that the status of all pairs is Paired.
 - To confirm pair status, see Monitoring pair status on page 8-3.
 - To split pairs, see Splitting pairs on page 6-5.
- 2. On the local array, delete the remote path. Be sure to confirm that the pair status is Split before deleting the remote path. See Deleting the remote path on page 6-8.
- 3. Add the remote port CHAP secret on the remote array. See the instructions above.
- 4. Re-create the remote path on the local array. See Defining the remote path on page 5-5.
 - For the CHAP secret field, select **manually** to enable the CHAP Secret boxes so that the CHAP secrets can be entered. Use the CHAP secret added on the remote array.
- 5. Resynchronize the pairs after confirming that the remote path is set. See Resynchronizing pairs on page 6-6.

Defining the remote path

Data is transferred between the P-VOL and S-VOL on the remote path. The remote path is set up on the local array.

Prerequisites

- Two paths are recommended; one from controller 0 and one from controller 1.
- Remote path information cannot be edited after the path is set up. To make changes, it is necessary to delete the remote path then set up a new remote path with the changed information.

- Both local and remote arrays must be connected to the network for the remote path.
- The remote array ID will be required on the GUI screen. The remote array ID is shown on the main array screen.
- Network bandwidth will be required.
- For iSCSI, the following additional information is required:
 - Remote IP address, listed in the remote array's GUI Settings/IP Settings
 - TCP port number. You can see this by navigating to the remote array's GUI's Settings/IP Settings/selected port screen.
 - CHAP secret (if specified on the remote array—see Adding or changing the remote port CHAP secret on page 5-4 for more information).

To set up the remote path

- 1. On the local array, from the navigation tree, click **Replication**, then click **Setup**. The Setup screen displays.
- 2. Click **Remote Path**; on the Remote Path screen click the **Create Remote Path** button. The Create Remote Path screen displays.
- 3. For Interface Type, select **Fibre** or **iSCSI**.
- 4. Enter the **Remote Array ID**.
- Enter Bandwidth.
- (iSCSI only) In the CHAP secret field, select **Automatically** to allow TCE to create a default CHAP secret, or select **manually** to enter previously defined CHAP secrets. The CHAP secret must be set up on the remote array.
- 7. In the two remote path boxes, Remote Path 0 and Remote Path 1, select local ports. For iSCSI, enter the **Remote Port IP Address** and **TCP Port No.** for the remote array's controller 0 and 1 ports.
- 8. Click Ok.

Using TrueCopy

This chapter provides procedures for performing basic TrueCopy operations using the Navigator 2 GUI. Appendixes with CLI and CCI instructions for the same operations are included in this manual.

- TrueCopy operations
- Checking pair status
- Creating pairs
- Splitting pairs
- Resynchronizing pairs
- Swapping pairs
- Editing pairs
- Deleting pairs
- Deleting the remote path

TrueCopy operations

Basic TrueCopy operations consists of the following. See Chapter 7, Disaster recovery for disaster recovery procedures.

- Always check pair status. Each operation requires the pair to be in a specific status.
- Create the pair, in which the S-VOL becomes a duplicate of the P-VOL.
- Split the pair, which separates the P-VOL and S-VOL and allows read/ write access to the S-VOL.
- Re-synchronize the pair, in which the S-VOL again mirrors the on-going, current data in the P-VOL.
- Swap pairs, which reverses pair roles.
- Delete a pair.
- Edit pair information.

Checking pair status

Each TrueCopy operation requires a specific pair status. Before performing any operation, check pair status.

- Find status requirements for each operation under the Prerequisites sections.
- To view a pair's current status in the GUI, please refer to Monitoring pair status on page 8-3

Creating pairs

A TrueCopy pair consists of primary and secondary volume whose data stays synchronized until the pair is split. During the create pair operation, the following takes place:

- All data in the local P-VOL is copied to the remote S-VOL.
- The P-VOL remains available to the host for read/write throughout the copy operation.
- Pair status is Synchronizing while the initial copy operation is in progress.
- Status changes to Paired when the initial copy is complete.
- New writes to the P-VOL continue to be copied to the S-VOL in the Paired status.

Prerequisite information and best practices

- Logical units for the primary and secondary volumes must be created and formatted. They must have identical block counts.
- Logical units for volumes to be paired must be in Simplex status.
- DMLUs must be set up.

- The create pair and resynchronize operations affect performance on the host. Therefore:
 - Perform the operation when I/O load is light.
 - Limit the number of pairs that you create simultaneously within the same RAID group to two.
 - If a TrueCopy pair is cascaded with ShadowImage, and the pair of one or the other is in Paired or Synchronizing status, place the other in Split status to lower the impact on performance.
 - If you have two TrueCopy pairs on the same two arrays and the pairs are bi-directional, perform copy operations at different times to lower the impact on performance.
 - Monitor write-workload on the remote array as well as on the local array. Performance on the remote array affects performance on the local array, since TrueCopy operations are slowed down by unrelated remote operations. Performance backup reverberates across the two systems.
 - Use a copy pace that matches your priority for either performance or copying speed.

The following sections discuss options in the Create Pair procedure.

Copy pace

Copy pace is the speed at which data is copied during pair creation or resynchronization. You select the copy pace on the GUI procedure when you create or resync a pair (if using CLI, you enter a copy pace parameter).

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

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

Fence level

The Fence Level setting determines whether the host is denied access to the P-VOL if a TrueCopy pair is suspended due to an error.

You must decide whether you want to bring the production application(s) to a halt if the remote site is down or inaccessible.

There are two synchronous fence-level settings:

- Never The P-VOL will never be fenced. "Never" ensures that a host never loses access to the P-VOL, even if all TrueCopy copy operations are stopped. Once the failure is corrected, a full re-copy may be needed to ensure that the S-VOL is current. This setting should be used when I/O performance out-weighs data recovery.
- Data The P-VOL will be fenced if an update copy operation fails. "Data" insures that the S-VOL remains identical to the P-VOL. This is done by preventing the host from writing to the P-VOL during a failure. This setting should be used for critical data.

Create pair procedure

To create a pair

- 1. In the Navigator 2 GUI, click the check box for the local array, then click the **Show & Configure Array** button.
- 2. In the navigation tree, click **Replication**, then click **Remote Replication**. The Remote Replication screen displays.
- 3. Click the **Create Pair** button. The Create Pair screen displays.
- 4. Enter a **Pair Name**, if desired. Omitting the pair name results in a default name created for you (which can be changed via the Edit Pair screen).
- 5. Select the **LUN** to become the P-VOL.
- 6. Enter the **LUN** for the S-VOL. Keep in mind that the P-VOL and S-VOL must have matching sizes.
- 7. Click the **Advanced** tab.
- 8. From the **Copy Pace** dropdown list, select the speed at which copies will be made. Select Slow, Medium, or Fast. See Copy pace, above, for more information.
- 9. In the **Group Assignment** area, you have the option of assigning the new pair to a consistency group. (For a description, see Consistency Group (CTG) on page 1-4.) Do one of the following:
 - If you do not want to assign the pair to a consistency group, leave the **Ungrouped** button selected.
 - To create a group and assign the new pair to it, click the New or existing Group Number button and enter a new number for the group in the box.
 - To assign the pair to an existing group, enter its number in the Group Number box, or enter the group name in the Existing Group Name box.



NOTE: When a group is created, future pairs can be added to it. You can also name the group on the Edit Pair screen. See Editing pairs on page 6-7 for details.

10. In the **Do initial copy from the primary volume ...** field, leave **Yes** checked to copy the primary to the secondary volume.

Clear the check box to create a pair without copying the P-VOL at this time. Do this when the S-VOL is already a copy of the P-VOL.

- 11. Select a **Fence Level** of Never or Data. See Fence level on page 6-3 for more information.
- 12. Click OK.
- 13. Check the **Yes, I have read ...** message then click Confirm.
- 14. Click Close on the success message.

Splitting pairs

All data written to the P-VOL is copied to the S-VOL when the pair is in Paired status. This continues until the pair is split. Then, updates continue to be written to the P-VOL but not the S-VOL. Data in the S-VOL is frozen at the time of the split, and the pair is no longer synchronous.

When a pair is split:

- Data copying to the S-VOL is completed so that the data is identical with P-VOL data. The time it takes to perform the split depends on the amount of differential data copied to the S-VOL.
- If the pair is included in a consistency group, all pairs in the group are split.

After the Split Pair operation:

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

Prerequisites

The pair must be in Paired status.

To split the pair

- 1. In Navigator 2 GUI, select the desired array, then click the **Show & Configure Array** button.
- 2. From the Replication tree, select the **Remote Replication** icon. The **Pairs** screen displays.
- 3. Click the check box for the pair you want to split, then click **Split Pair**. The Split Pair screen appears.
- 4. Select the **Option** you want for the S-VOL, Read/Write, which makes the S-VOL available to be written to by a secondary application, or Read Only, which prevents it from being written to by a secondary application.
- 5. Click **OK** and **Close**.

Resynchronizing pairs

Re-synchronizing a pair that has been split updates the S-VOL so that it is again identical with the P-VOL. Differential data accumulated on the local array since the last pairing is updated to the S-VOL.

- Pair status during a re-synchronizing is Synchronizing.
- Status changes to Paired when the resync is complete.
- If P-VOL status is Failure and S-VOL status is Takeover or Simplex, the pair cannot be recovered by resynchronizing. The pair must be deleted and created again.

Prerequisites

- The pair must be in Split status.
- The prerequisites for creating a pair apply to resynchronizing. See Creating pairs on page 6-2.

To resync the pair

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

Swapping pairs

In a pair swap, primary and secondary-volume roles are reversed. Data flows from remote to local or new array.

A pair swap is performed when data in the S-VOL must be used to restore the local array, or possibly to a new array/volume following disaster.

Prerequisites and Notes

- The pair must be in Paired, Split, or Takeover status.
- The pair swap is executed on the remote array.
- A remote path must be created on the remote array. This is done using Navigator 2 GUI by connecting to the remote array and creating a remote path to the local array.
- If a pair is swapped with a remote path failure, the primary/secondary roles are reversed but data cannot be restored.
- When a pair that is in a consistency group is swapped, all pairs in the group are swapped.

To swap TrueCopy pairs

 In Navigator 2 GUI, connect to the remote array, then click the Show & Configure Array button.

- 2. From the Replication tree, select the **Remote Replication** icon. The **Pairs** screen displays.
- 3. Select the pair you want to swap.
- 4. Click the **Swap Pair** button.
- 5. On the message screen, check the **Yes, I have read ...** box, then click **Confirm**.
- 6. Click **Close** on the confirmation screen.

Editing pairs

You can edit the name, group name, and copy pace for a pair. A group created with no name can be named from the Edit Pair screen.

To edit pairs

- In Navigator 2 GUI, select the desired array, then click the Show & Configure Array button.
- 2. From the Replication tree, select the **Remote Replication** icon. The **Pairs** screen displays.
- 3. Select the pair that you want to edit.
- 4. Click the Edit Pair button.
- 5. Make any changes, then click **OK**.
- 6. On the confirmation message, click **Close**.



NOTE: Edits made on the local array are not reflected on the remote array. To have the same information reflected on both arrays, it is necessary to edit the pair on the remote array also.

Deleting pairs

When a pair is deleted, transfer of differential data from P-VOL to S-VOL is completed, then the volumes become Simplex. The pair is no longer displayed in the Pairs list on Navigator 2 GUI.

- A pair can be deleted regardless of its status. However, data consistency is not guaranteed unless status prior to deletion is Paired.
- If the operation fails, the P-VOL nevertheless becomes Simplex. Transfer of differential data from P-VOL to S-VOL is terminated.
- Normally, a Delete Pair operation is performed on the local array where the P-VOL resides. However, it is possible to perform the operation from the remote array, though with the following results:
 - Only the S-VOL becomes Simplex.
 - Data consistency in the S-VOL is not guaranteed.
 - The P-VOL does not recognize that the S-VOL is in Simplex status. When the P-VOL tries to send differential data to the S-VOL, it recognizes that the S-VOL is absent and the pair becomes Failure.
 - When the pair status changes to Failure, the status of the other pairs in the group also becomes Failure.

 From the remote array, this Failure status is not seen and pair status remains Paired.

To delete a pair

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

Deleting the remote path

Prerequisites

 The pair status of the volumes using the remote path to be deleted must be Simplex or Split.

To delete the remote path

- 1. In Navigator 2 GUI, select the desired array, then click the **Show & Configure Array** button.
- 2. From the Replication tree, select the **Setup** icon, then select **Remote Path**. The Remote Path screen displays.
- 3. Select the remote path you want to delete in the Remote Path list and click **Delete Path**.
- 4. A message appears. Click Close.

Disaster recovery

TrueCopy is designed for you to create and maintain a viable copy of the production volume at a remote location. If the data path or local host or array go down, data in the remote copy is available for restoring the primary volume, or for host operations to be switched to the remote site to continue operations.

This chapter provides procedures for disaster recovery under varying circumstances.

- Disaster recovery
- Resynchronizing the pair
- Data path failure and recovery
- Host server failure and recovery
- Production site failure and recovery
- Special problems and recommendations

Disaster recovery

This chapter describes restoration procedures without and with a ShadowImage or SnapShot backup. Four scenarios are presented with recovery procedures.



NOTE: In the following procedures, ShadowImage or SnapShot pairs are cascaded with TrueCopy and are referred to as the "backup pair".

Also, CCI is located on a host management server, and the production applications are located on a host production server.

Resynchronizing the pair

In this scenario, an error has occurred in the replication environment, which causes the TrueCopy pair to suspend with an error. To recover TrueCopy pair operations, proceed as follows:

- 1. If ShadowImage or SnapShot backups exist on the remote array, resynchronize the backup pair.
- 2. Split the backup pair.
- 3. Confirm that the remote path is operational.
- 4. Resynchronize the TrueCopy pair.

Data path failure and recovery

In this scenario, a power outage has disabled the local site and the remote path. A decision is made to move host applications to the remote site, where the TrueCopy S-VOL will become the primary volume. When the data path is unavailable, CCI cannot communicate with the opposite server, and the horctakeover command cannot perform the pair resync/swap operation.

- 1. Shutdown applications on the local production site.
- 2. Execute the CCI horctakeover command from the remote site. This transfers access to the remote array. The S-VOL status becomes SSWS, which is read/write enabled.
- 3. At this time, horctakeover attempts a "swaps resync". This fails because the data path is unavailable.
- 4. Execute the CCI pairdisplay -fc command to confirm the SSWS status.
- 5. At the remote site, mount the volumes and bring up the applications.
- 6. When the local site and/or data path is again operational, prepare to reestablish TrueCopy operations, as follows:
 - a. If ShadowImage or SnapShot are used, re-synchronize the backup pair on the remote array. Use the Navigator 2 GUI or CLI to do this if the production server is not available to run the CCI command.
 - b. Confirm that the backup pair status is Paired, then split the pair.
 - c. Create a data path, from the remote array to the local array. See Defining the remote path on page 5-5 for instructions. Confirm that the data path is operational.

- 7. Perform the CCI pairresync -swaps command at the remote site. This completes the reversal of the PVOL-SVOL relationship. TrueCopy operations now flow from the remote array to the local array.
- 8. Execute the CCI pairdisplay command to confirm that the P-VOL pair is on the remote array.

The remote site is now the production site.

Next, the production environment is returned to its normal state at the local site. Applications should be moved back to the local production site in a staged manner. The following procedure may be performed hours, days, or weeks after the completion of the previous steps.

- 9. Shut down the application(s) on the remote server then unmount the volumes.
- 10. Boot the local production servers. DO NOT MOUNT the volumes or start the applications.
- 11. Execute the CCI horctakeover command on the local management server. Because the data path is operational, this command includes the pair resync/swap operation, which reverses TrueCopy pair roles. The S-VOL becomes read/write disabled.
- 12. Execute the CCI pairdisplay command to confirm that the P-VOL pair is now on the local array.
- 13. Mount the volumes and start the applications on the local production server.

Host server failure and recovery

In this scenario, the host server fails and no local standby server is available. A decision is made to move the application to the remote site. The data path and local array are still operational.

- 1. Confirm that the remote server is available and that all CCI horem instances are started on it.
- 2. Ensure that a live replication path (remote path) from the remote to the local array is in place.
- 3. Execute the CCI command horctakeover from the remote server. Because the data path link is operational, this command includes pair resync/swap, which reverses TrueCopy pair roles.
- 4. Execute the CCI command pairdisplay to confirm that the P-VOL pair is on the remote site.
- 5. Mount the volumes and start the applications.

The remote site is now the production site.

At a later date, the host server is restored, and the applications are moved back to the production site in a staged manner. The following procedure may be performed hours, days, or weeks after the completion of the previous steps.

- 6. Shutdown applications and unmount volumes on remote server
- 7. Boot the local production server. DO NOT mount the volumes or start the applications.

- 8. Execute the CCI horctakeover command from the local management server. Because the data path is operational, this command includes the pair resync/swap operation, which reverses TrueCopy pair roles. The S-VOL becomes read/write disabled.
- 9. Execute the CCI pairdisplay command to confirm that the P-VOL pair is on the local site.
- 10. Mount the volumes and start the applications.

Production site failure and recovery

In this scenario, the production site is unavailable and a decision is made to move all operations to the remote site.

- 1. Confirm that the necessary remote servers are available and that all CCI horem instances are started.
- 2. Execute the CCI horctakeover command from the remote server.
- 3. Execute the CCI pairdisplay -fc command to confirm that the P-VOL pair is on the remote site, or, if the data path is unavailable, that the S-VOL is in SSWS status, which enables read/write.
- 4. Mount the volumes and start the applications on the remote site.
 - The remote site is now the production site.
 - When the local site is again operational, TrueCopy operations can be restarted.
- 5. Make sure the data path is deleted on the local array.
- 6. If ShadowImage is cascaded on the original remote site, re-synchronize the pair.
- 7. Split the ShadowImage pair.
- 8. Set up the remote path, on the remote array.
- 9. Execute the CCI pairresync -swaps command from the remote management server.
- 10. If Step 9 fails, you may need to recreate the paths and delete and create new TrueCopy pairs, which will require full synchronization on all LUNs. You may need to contact the support center.
- 11. When TrueCopy operations are normal and data is flowing from the remote P-VOL to the local S-VOL, production operations can be switched back to the local side again. To do this, start at Step 6 in the Host server failure and recovery section page 7-3.

Special problems and recommendations

- Before path management software (such as Hitachi Dynamic Link Manager) handles a recovery, the remote path must be manually recovered and the pair resynchronized. Otherwise the management software freezes when attempting to write to the P-VOL.
- The file systems for UNIX and Windows 2000 Server /Windows Server 2003/2008 do not have writing logs or journal files. Though "data" may be specified as the fence level, the file sometimes does not correspond to the directory. In this case, "fsck" and "chkdsk" must be executed before the S-VOL can be used, though data consistency cannot be completely guaranteed. The work-around is to cascade TrueCopy with ShadowImage, which saves the complete P-VOL or S-VOL to the ShadowImage S-VOL.

For more information on performing system recovery using CCI, see the *Hitachi AMS 2000 Family Command Control Interface (CCI) Reference Guide*.

For information on cascading with ShadowImage or SnapShot, see Appendix D or Appendix E.

Monitoring, maintenance, troubleshooting

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

- Monitoring and maintenance
- Troubleshooting

Monitoring and maintenance

Monitoring the TrueCopy system is an ongoing operation that should be performed to maintain your pairs as intended.

- When you want to perform a pair command, first check the pair's status. Each operation requires a specific status or statuses.
- Pair-status changes when an operation is performed. Check status to see that TrueCopy pairs are operating correctly and that data is updated from P-VOLs to S-VOLs in the Paired status, or that differential data management is performed in the Split status.

When a hardware failure occurs, a pair failure may occur as a result. When a pair failure occurs, the processes in Table 8-1 are executed:

Table 8-1: Processes when Pair Failure Occurs

Management Software	Results
Navigator 2	A message is displayed in the event log.
	The pair status is changed to Failure.
CCI	An error message is output to the system log file, as shown in Table 8-2.
	The pair status is changed to PSUE. (On UNIX the syslog file displays; on Windows 2000, the eventlog file display.)
SNMP Agent Support Function	A trap is reported.

Table 8-2: CCI System Log Message when PSUE

Message ID	Condition	Cause
HORCM_102		The pair status was suspended due to code 0006.

See the maintenance log section in the *Hitachi AMS Command Control Interface (CCI) Reference Guide* for more information.

A sample script is provided for CLI users in Pair-monitoring script on page B-14.

Monitoring pair status

Monitoring using the GUI is done at the user's discretion. Monitoring should be repeated frequently. Email notifications of problems can be set up using the GUI.

To monitor pair status using the GUI

- 1. In Navigator 2 GUI, select the desired array, then click the **Show & Configure Array** button.
- 2. From the Replication tree, select the **Remote Replication** icon. The **Pairs** screen displays, as shown in Figure 8-1.

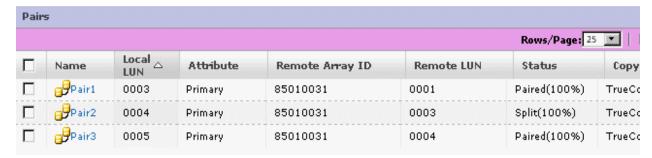


Figure 8-1: Pairs Screen

- 3. Locate the pair whose status you want to review in the **Pair** list. Status descriptions are provided in Table 8-3. You can click the **Refresh Information** button (not in view) to make sure data is current.
 - The percentage that displays with each status shows how close the S-VOL is to being completely paired with the P-VOL.
 - The Attribute column shows the pair volume for which status is shown.

Table 8-3: TrueCopy Pair Status

Pair Status	Description	P-VOL Access	S-VOL Access
Simplex	The volume is not assigned to the pair. The array accepts read and write operations for Simplex volumes.	Read: Yes Write: Yes	Read: Yes Write: Yes
Synch- ronizing	The copy operation is in progress.	Read: Yes Write: Yes	Read: Yes, mount operation disabled. Write: No
Paired	The copy operation is complete.	Read: Yes Write: Yes	Read: Yes, mount operation disabled. Write: No
Split	The copy operation is suspended. The array starts accepting write operations for P-VOL and S-VOL. When the pair is resynchronized, the array executes the differential data copying from P-VOL to S-VOL.	Read: Yes Write: Yes	Read: Yes Write: Yes

Table 8-3: TrueCopy Pair Status

Pair Status	Description	P-VOL Access	S-VOL Access
Takeover	Takeover is a transitional status after Swap Pair is executed. Immediately after the pair is changed to Takeover status, the pair relationship is swapped and copy from the new P-VOL to the new S-VOL is started.	_	Read: Available Write: Available
Failure	A failure occurs and copy operations are suspended forcibly. If Data is specified as the fence level, the array rejects all write I/O. Read I/O is also rejected if PSUE Read Reject is specified. If Never is specified, read/write I/O continues as long as the LU is unblocked S-VOL read operations are accepted but not write operations. To recover, resynchronize the pair (might require copying entire P-VOL).' See Fence level on page 6-3 for more information.	Read: Yes/ No Write: Yes/No	Read: Yes Write: No

For CCI status see Confirming pair status on page C-9.

Status Narrative: If a volume is not assigned to a TrueCopy pair, its status is Simplex. When a TrueCopy pair is being created, the status of the P-VOL and S-VOL is Synchronizing. When the copy operation is complete, status becomes Paired. If the system cannot maintain Paired status for any reason, the pair status changes to Failure. When the Split Pair operation is complete, the pair status changes to Split and the S-VOL can be written to. When you start a Resync Pair operation, the pair status changes to Synchronizing. When the operation is completed, the pair status changes to Paired. When you delete a pair, pair status changes to Simplex.



NOTE: Pair status for the P-VOL can differ from status for the S-VOL. If the remote path breaks down when pair status is Paired, the local array becomes Failure because it cannot send data to the remote array. The remote array remains Paired though there is no write I/O from the P-VOL.

Monitoring the remote path

Monitor the remote path to ensure that data copying is unimpeded. If a path is blocked, the status is Detached, and data cannot be copied.

You can adjust remote path bandwidth to improve data transfer rate.

To monitor the remote path

1. In the Replication tree, click **Setup**, then **Remote Path**. The Remote Path screen displays.

- 2. Review statuses and bandwidth. Path statuses can be Normal, Blocked, or Diagnosing. When Blocked or Diagnosing is displayed, data cannot be copied.
- 3. Take corrective steps as needed.

Troubleshooting

Pair failure

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

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

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

To restore pairs after a hardware failure

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

To restore pairs after forcible release operation

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



NOTE: A TrueCopy operation may stop when the format command is run on Navigator 2 or Windows 2000 Server and Windows Server 2003/2008. This operation causes the Format, Synchronize Cache, or Verify operations to be performed, which in turn cause TrueCopy path blockage. The copy operation restarts when the Verify, Format, or Synchronize Cache operation is completed.



Specifications

This appendix provides specifications for TrueCopy.

■ TrueCopy specifications

Specifications **A-1**

TrueCopy specifications

Table A-1 lists external specifications for TrueCopy.

Table A-1: Specifications

Parameter	TrueCopy Requirement
License	Key code required for installation.
User Interface	 Navigator 2 GUI and/or CLI: Setup and pair operations. CCI: Pair operations. Certain related operations only available using CCI.
Command device	 Required for CCI, one per array. Up to 128 allowed per array. Must be 65,538 blocks or more (1 block = 512 bytes) (33 M bytes or more).
Host Interface	AMS 2100, 2300, or 2500: Fibre Channel or iSCSI
Controller configuration	Configuration of dual controller is required.
Remote path	 The interface type of two remote paths between arrays must be the same, fibre channel or iSCSI. One remote path per controller is required, with a total of two between the arrays in a the dual controller configuration.
Port modes	Initiator and target intermix mode. One port may be used for host I/O and TrueCopy at the same time.
Bandwidth supported	1.5 M bps or more (100 M bps or more is recommended.) Low transfer rate results in greater time for TrueCopy operations and reduced host I/O performance.
DMLU	 Minimum size: 10 GB One DMLU is required, two (maximum) are recommended. When using two, each must be: In different RAID groups than the other Under different controller than the other
Unit of pair management	LUs are the target of TrueCopy pairs, and are managed per logical unit.
Maximum number of LUs in which pair can be created	 AMS2100: 2,046 AMS2300, AMS2500: 4,094 When different arrays are used, maximum is the smaller number.
Pair structure	One copy (S-VOL) per P-VOL.
Supported RAID level	 RAID 1 (1D+1D), RAID 5 (2D+1P to 15D+1P) RAID 1+0 (2D+2D to 8D+8D) RAID 6 (2D+2P to 28D+2P)
Combination of RAID levels	All combinations supported. The number of data disks does not have to be the same.
P-VOL:S-VOL size	1:1, P-VOL = S-VOL.

A–2 Specifications

Table A-1: Specifications (Continued)

Parameter	TrueCopy Requirement
Drive Types for P-VOL/S-VOL	If the drive types are supported by the array, they can be set for a P-VOL and an S-VOL. However, SAS, or SSD drives are recommended, especially for the P-VOL. When a pair is created using two LUs configured by the SATA drives, requirements for using the SATA drives may differ.
Supported capacity value of P-VOL and S-VOL	The capacity for TrueCopy is limited. See Calculating supported capacity on page 3-15.
Consistency Groups (CTG) supported	 Up to 256 CTGs per array. Maximum number of pairs one CTG can manage is 2,046 for AMS2100, 4,094 for AMS2300/2500.
Management of LUs while using TrueCopy	 A TrueCopy pair must be deleted before the following operations: Deletion of the pair's RAID group Deletion of an LU Deletion of the pair's data pool Formatting, growing, or shrinking an LU
Restrictions during RAID group expansion	A RAID group with a TrueCopy P-VOL or S-VOL can be expanded only when the pair status is Simplex or Split.
Pair creation of a unified LU	A TrueCopy pair can be created by specifying a unified LU. However, unification of LUs or release of the unified LU cannot be done for the paired LUs.
Failures	When a failure of the copy operation from P-VOL to S-VOL occurs, TrueCopy will suspend the pair (Failure). If a LU failure occurs, TrueCopy suspend the pair. If a drive failure occurs, the TrueCopy pair status is not affected because of the RAID architecture.
Concurrent use of Data Retention Utility	 Yes, but note following: When S-VOL Disable is set for an LU, the LU cannot be used in a pair. When S-VOL Disable is set for an LU that is already an S-VOL, no suppression of the pair takes place, unless the pair status is split.
Concurrent use of Cache Residency Manager	Yes, but an LU specified by Cache Residency Manager cannot be used as a P-VOL or S-VOL.
Concurrent use of SNMP Agent	Yes. A trap is transmitted when: • A failure occurs in the remote path • Pair status changes to Failure.
Concurrent use of Volume Migration	Yes, but, a Volume Migration P-VOL, S-VOL, and Reserved LU of cannot be specified as a TrueCopy P-VOL or an S-VOL
Concurrent use of TCE	No.
Concurrent use of Cache Partition Manager	Not available. The DP-VOLs created by Dynamic Provisioning cannot be used in TrueCopy. When a DP-VOL is included in a ShadowImage pair, the ShadowImage pair and TrueCopy cannot be cascaded.
Concurrent use of ShadowImage	Yes. TrueCopy can be used together with ShadowImage and cascaded with ShadowImage.

Specifications A-3

Table A-1: Specifications (Continued)

Parameter	TrueCopy Requirement
Concurrent use of SnapShot	Yes. TrueCopy can be used together with SnapShot and cascaded with SnapShot.
Concurrent use of Power Saving	Yes. However, when a P-VOL or an S-VOL is included in a RAID group for which the Power Saving is specified, only a TrueCopy pair split and the pair delete can be performed.
Concurrent use of Dynamic Provisioning	 Available. However, if the firmware array version is less than 0893/A, The DP-VOLs created by Dynamic Provisioning cannot be used in TrueCopy. A normal LU can be set for a P-VOL or an S-VOL of TrueCopy. A DP-VOL is included in a ShadowImage pair; the ShadowImage pair and TrueCopy cannot be cascaded. For more details, see Cascade connection restrictions on page 3-17.
Reduction of memory	Memory cannot be reduced when the ShadowImage, SnapShot, TrueCopy, or Volume Migration function are enabled. Reduce memory after disabling the functions.



Operations using CLI

This appendix describes CLI procedures for setting up and performing TrueCopy operations.

- Installation and setup
- Pair operations
- Sample scripts



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

Installation and setup

This section provides installation/uninstalling, enabling/disabling, and setup procedures.

TrueCopy is an extra-cost option and must be installed using a key code or file. Obtain it from the download page on the HDS Support Portal, http://support.hds.com. For prerequisites, see the GUI-based instructions in Installation procedures on page 5-2.

Installing

TrueCopy cannot be installed if more than 239 hosts are connected to a port on the array.

To install TrueCopy

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

```
% auopt -unit array-name -lock off -keycode manual-attached-keycode 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.
```

3. Execute the auopt command to confirm whether TrueCopy has been installed. For example:

TrueCopy is installed and enabled. Installation is complete.

Enabling or disabling

TrueCopy can be disabled or enabled. When TrueCopy is first installed it is automatically enabled.

Prerequisites for disabling

- TrueCopy pairs must be released (the status of all LUs must be Simplex).
- The remote path must be released.
- TrueCopy cannot be enabled if more than 239 hosts are connected to a port on the array.

To enable/disable TrueCopy

1. From the command prompt, register the array in which the status of the feature is to be changed, and then connect to the array.

2. Execute the auopt command to change TrueCopy 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.

```
% auopt -unit array-name -option TRUECOPY -st disable
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 that the status has been changed. For example:

Uninstalling

To uninstall TrueCopy, the key code provided for optional features is required.

Prerequisites for uninstalling

- All TrueCopy pairs must be released (the status of all LUs must be Simplex).
- The remote path must be released.

To uninstall TrueCopy

- 1. From the command prompt, register the array in which the TrueCopy is to be uninstalled, and then connect to the array.
- 2. Execute the auopt command to uninstall TrueCopy. For example:

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

3. Execute the auopt command to confirm that TrueCopy is uninstalled. For example:

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

Setting the Differential Management-Logical Unit

The DMLU must be set up before TrueCopy copies can be made. Please see the prerequisites under Defining the DMLU on page 5-4 before proceeding.

To set up the DMLU

- 1. From the command prompt, register the array to which you want to set the DMLU. Connect to the array.
- 2. Execute the audmlu command. This command first displays LUs that can be assigned as DMLUs and then creates a DMLU. For example:

Release a DMLU

Observe the following when releasing a DMLU for TrueCopy:

- When only one DMLU is set, it cannot be released.
- When two DMLUs are set, only one can be released.

To release a TrueCopy DMLU

Use the following example:

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

Setting the remote port CHAP secret

For iSCSI, the remote path can employ a CHAP secret. Set the CHAP secret mode on the remote array. For more information on the CHAP secret, see Adding or changing the remote port CHAP secret on page 5-4.

The setting procedure of the remote port CHAP secret is shown below.

- 1. From the command prompt, register the array in which you want to set the remote path, and then connect to the array.
- 2. Execute the aurmtpath command with the –set option and perform the CHAP secret of the remote port. The input example and the result are shown below. For example:

```
% aurmtpath -unit array-name -set -target -local 85000027 -
secret
Are you sure you want to set the remote path information?
  (y/n[n]): y
Please input Path 0 Secret.
Path 0 Secret:
Re-enter Path 0 Secret:
Please input Path 1 Secret.
Path 1 Secret:
Re-enter Path 1 Secret:
The remote path information has been set successfully.
%
```

The setting of the remote port CHAP secret is completed.

Setting the remote path

Data is transferred from the local to the remote array over the remote path.

Please review Prerequisites in GUI instructions in Defining the remote path on page 5-5 before proceeding.

To set up the remote path

- 1. From the command prompt, register the array in which you want to set the remote path, and then connect to the array.
- 2. The following shows an example of referencing the remote path status where remote path information is not yet specified.

Fibre channel example:

```
% aurmtpath -unit array-name -refer
Initiator Information
Local Information
          : 85000026
 Array ID
 Distributed Mode : N/A
Path Information
 Interface Type
                    : ---
 Remote Array ID
 Remote Path Name
 Bandwidth [0.1 Mbps] : ---
 iSCSI CHAP Secret : ---
                                  Remote Port TCP Port No. of
                             Remote IP Address
 Path Status
                    Local
                                                     Remote Port
   0 Undefined
    1 Undefined
```

iSCSI example:

```
% aurmtpath -unit array-name -refer
Initiator Information
 Local Information
   Array ID : 85000026
   Distributed Mode : N/A
Path Information
                : ---
 Interface Type
 Remote Array ID
 Remote Path Name
                  : ---
 Bandwidth [0.1 Mbps] : ---
 iSCSI CHAP Secret : ---
                                Remote Port TCP Port No. of
  Path Status Local
                           Remote IP Address Remote Port
    0 Undefined
                   ---
                            ___
    1 Undefined
                   ---
                            ---
                                  ---
Target Information
 Local Array ID
```

3. Execute the aurmtpath command to set the remote path.

Fibre channel example:

iSCSI example:

4. Execute the aurmtpath command to confirm whether the remote path has been set. For example:

Fibre channel example:

```
% aurmtpath -unit array-name -refer
Initiator Information
Local Information
 Array ID : 85000026
 Distributed Mode : N/A
Path Information
                 : FC
 Interface Type
 Remote Array ID
                   : 85000027
 Remote Path Name : Array 85000027
 Bandwidth [0.1 Mbps] : 15
 iSCSI CHAP Secret : N/A
                              Remote Port TCP Port No. of
 Path Status
                 Local Remote IP Address Remote Port
   0 Normal
                   0.A
                           OA N/A
                                                       N/A
   1 Normal
                    1A
                             1B N/A
                                                       N/A
응
```

iSCSI example:

```
% aurmtpath -unit array-name -refer
Initiator Information
 Local Information
   Array ID : 85000026
   Distributed Mode : N/A
 Path Information
                  : iSCSI
   Interface Type
                     : 85000027
   Remote Array ID
   Remote Path Name : Array 85000027
   Bandwidth [0.1 Mbps] : 100
   iSCSI CHAP Secret : Disable
                                 Remote Port TCP Port No. of
                            Remote IP Address Remote Port
  Path Status
                   Local
    0 Normal
                     0B
                            N/A 192.168.0.201
                                                         3260
    1 Normal
                     1B
                            N/A 192.168.0.209
                                                         3260
Target Information
 Local Array ID
               : 85000026
```

Deleting the remote path

When shutdown of the arrays is necessary, the remote path must be deleted first. The status of TrueCopy logical units must be Simplex or Split.

To delete the remote path

- 1. From the command prompt, register the array in which you want to delete the remote path, and then connect to the array.
- 2. Execute the aurmtpath command to delete the remote path. For example:

```
% aurmtpath -unit array-name -rm -remote 85000027
Are you sure you want to delete the remote path information?
  (y/n[n]): y
The remote path information has been deleted successfully.
%
```

3. Execute the aurmtpath command to confirm that the path is deleted. For example:

```
% aurmtpath -unit array-name -refer
Initiator Information
Local Information
 Array ID : 85000026
 Distributed Mode : N/A
 Path Information
   Interface Type : ---
   Remote Array ID
                      : ---
   Remote Path Name : ---
   Bandwidth [0.1 Mbps] : ---
   iSCSI CHAP Secret : ---
                            Remote Port TCP Port No. of
 Path Status Local Remote IP Address Remote
Port
  0 Undefined
1 Undefined
                 ___
                          ---
                                                      ___
```

Pair operations

The following sections describe the CLI procedures and commands for performing TrueCopy operations.

Displaying status for all pairs

To display all pair status

- 1. From the command prompt, register the array to which you want to display the status of paired logical volumes. Connect to the array.
- 2. Execute the aureplicationremote -refer command. For example:

```
% aureplicationremote -unit local array-name -refer
Pair name Local LUN Attribute Remote LUN Status

Copy Type Group Name

TC_LU0000_LU0000 0 P-VOL 0 Paired(100%)

TrueCopy 0:

TC_LU0001_LU0001 1 P-VOL 1 Paired(100%)

TrueCopy 0:
```

Displaying detail for a specific pair

To display pair details

- 1. From the command prompt, register the array to which you want to display the status and other details for a pair. Connect to the array.
- 2. Execute the aureplicationremote -refer -detail command to display the detailed pair status. For example:

```
% aureplicationremote -unit local array-name -refer -detail -pvol 0 -svol 0
                        -locallun pvol -remote 85000027
Pair Name
                     : TC LU0000 LU0000
Local Information
                     : 0
 Attribute : P-VOL
Remote Information
Array ID : 85000027
Path Name : Array_85000027
Capacity : 50.0 GB
Status
                     : Paired(100%)
              : TrueCopy
Copy Type
Group Name : 0:
Data Pool : N/A
Data Pool Usage Rate : N/A
Consistency Time : N/A
Difference Size : N/A
Copy Pace : Prior
Fence Level : Never
Previous Cycle Time : N/A
```

Creating a pair

See prerequisite information under Creating pairs on page 6-2 before continuing.

To create a pair

- 1. From the command prompt, register the local array in which you want to create pairs, and then connect to the array.
- 2. Execute the aureplicationremote -refer -availablelist command to display logical units available for copy as the P-VOL. For example:

```
% aureplicationremote -unit local array-name -refer -availablelist -tc -pvol
Available Logical Units
LUN Capacity RAID Group DP Pool RAID Level Type Status
0 10.0 GB 0 N/A 6(9D+2P) SAS Normal
%
```

3. Execute the aureplicationremote -refer -availablelist command to display logical units on the remote array that are available as the S-VOL. For example:

```
% aureplicationremote -unit remote array-name -refer -availablelist -tc -pvol
Available Logical Units
LUN Capacity RAID Group DP Pool RAID Level Type Status
0 10.0 GB 0 N/A) 6(9D+2P) SAS Normal
%
```

4. Specify the logical units to be paired and create a pair using the aureplicationremote -create command. For example:

Creating pairs belonging to a group

To create multiple pairs that belongs to a group

1. Create a pair that belongs to a group by specifying a group number for the <code>-gno</code> command option. The new group and new pair are created at the same time. For example:

2. Create new pairs as needed and assign them to the group.

Splitting a pair

A pair split operation on a pair belonging to a group results in all pairs in the group being split.

To split a pair

- 1. From the command prompt, register the local array in which you want to split pairs, and then connect to the array.
- 2. Execute the aureplicationremote -split command to split the specified pair. For example:

Resynchronizing a pair

To resynchronize a pair

- 1. From the command prompt, register the local array in which you want to re-synchronize pairs, and then connect to the array.
- 2. Execute the aureplicationremote -resync command to resynchronize the specified pair. For example:

Swapping a pair

Please review the Prerequisites in Swapping pairs on page 6-6.

To swap the pairs, the remote path must be set to the local array from the remote array.

To swap a pair

- 1. From the command prompt, register the remote array in which you want to swap pairs, and then connect to the array.
- 2. Execute the aureplicationremote -swaps command to swap the specified pair. For example:

```
% aureplicationremote -unit remote array-name -swaps -tc -svol
2002
Are you sure you want to swap pair?
  (y/n [n]): y
The pair has been swapped successfully.
%
```

Deleting a pair

To delete a pair

- 1. From the command prompt, register the local array in which you want to delete pairs, and then connect to the array.
- 2. Execute the aureplicationremote -simplex command to delete the specified pair. For example:

3. Verify the pair status. For example:

```
% aureplicationremote -unit local array-name -refer
DMEC002015: No information displayed.
%
```

Changing pair information

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

To change pair information

- 1. From the command prompt, register the local array on which you want to change the TrueCopy pair information, and then connect to the array.
- 2. Execute the aureplicationlocal -chg command to change the TrueCopy pair information. In the following example, the copy pace is changed from normal to slow.

Sample scripts

This section provides sample CLI scripts for executing a backup and for monitoring pair status.

Backup script

The following example provides sample script commands for backing up a volume on a Windows host.

```
echo off
REM Specify the registered name of the arrays
set UNITNAME=Array1
REM Specify the group name (Specify "Ungroup" if the pair doesn't
belong to any group)
set G NAME=Ungrouped
REM Specify the pair name
set P NAME=TC LU0001 LU0002
REM Specify the directory path that is mount point of P-VOL and
set MAINDIR=C:\main
set BACKUPDIR=C:\backup
REM Specify GUID of P-VOL and S-VOL
PVOL GUID=xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxx
REM Unmounting the S-VOL
pairdisplay -x umount %BACKUPDIR%
REM Re-synchoronizeing pair (Updating the backup data)
aureplicationremote -unit %UNITNAME% -tc -resync -pairname
%P NAME% -qname %G NAME%
aureplicationmon -unit %UNITNAME% -evwait -tc -pairname %P NAME%
-gname %G NAME% -st paired -pvol
REM Unmounting the P-VOL
pairdisplay -x umount %MAINDIR%
REM Splitting pair (Determine the backup data)
aureplicationremote -unit %UNITNAME% -tc -split -pairname
%P NAME% -gname %G NAME%
aureplicationmon -unit %UNITNAME% -evwait -tc -pairname %P NAME%
-gname %G NAME% -st split -pvol
REM Mounting the P-VOL
pairdisplay -x mount %MAINDIR% Volume{%PVOL GUID%}
REM Mounting the S-VOL
pairdisplay -x mount %BACKUPDIR% Volume{%SVOL GUID%}
< The procedure of data copy from C:\backup to backup appliance>
```

When Windows 2000 Server and Windows Server 2008 are used, the CCI mount command is required when mounting or un-mounting a volume. The GUID, which is displayed by mountvol command, is needed as an argument when using the mount command. For more information, see the *Hitachi AMS Command Control Interface (CCI) Reference Guide*.

Pair-monitoring script

The following is a sample script for monitoring two TrueCopy pairs (TC_LU0001_LU0002 and TC_LU0003_LU0004). The script includes commands for informing the user when pair failure occurs. The script is reactivated after several minutes. The array must be registered.

```
REM Specify the registered name of the arrays
set UNITNAME=Array1
REM Specify the name of target group (Specify "Ungroup" if the
pair doesn't belong to any group)
set G NAME=Ungrouped
REM Specify the name of target pair
set P1 NAME=TC LU0001 LU0002
set P2 NAME=TC LU0003 LU0004
REM Specify the value to inform "Failure"
set FAILURE=14
REM Checking the first pair
:pair1
aureplicationmon -unit %UNITNAME% -evwait -tc -pairname %P1 NAME%
-gname %G NAME% -nowait
if errorlevel %FAILURE% goto pair1 failure
goto pair2
:pair1 failure
<The procedure for informing a user>*
REM Checking the second pair
aureplicationmon -unit %UNITNAME% -evwait -tc -pairname %P2 NAME%
-gname %G NAME% -nowait
if errorlevel %FAILURE% goto pair2_failure
goto end
:pair2_failure
<The procedure for informing a user>*
:end
```



Operations using CCI

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

- Preparing for CCI operations
- Pair operations

Preparing for CCI operations

The following must be set up for TrueCopy CCI operations

- Command device
- LU mapping
- Environment variable
- Configuration definition file

Enter CCI commands from the command prompt on the host where CCI is installed.

Setting the command device

CCI interfaces with AMS through the *command device*, a dedicated logical volume located on the local and remote arrays. CCI software on the host issues read/write commands to the command device. The command device is accessed as a raw device (no file system, no mount operation). It is dedicated to CCI communications and cannot be used by any other applications.

- The command device is defined in the HORCM_CMD section of the configuration definition file for the CCI instance on the attached host.
- 128 command devices can be designated for the AMS array.
- Logical units used as command devices must be recognized by the host.
- The command device must be 33 MB or greater.



If a command device fails, all commands are terminated. CCI supports an alternate command device function, in which two command devices are specified within the same array, to provide a backup. For details on the alternate command device function, refer to the *Hitachi Adaptable Modular Storage Command Control Interface (CCI) User's Guide.*

To designate a command device

- 1. From the command prompt, register the array to which you want to set the command device, and then connect to the array.
- 2. Execute the aucmddev command to set a command device. When this command is run, logical units that can be assigned as a command device display, then the command device is set. To use the CCI protection function, enter enable following the -dev option. The following is an example of specifying LU 2 for command device 1.

3. Execute the aucmddev command to verify that the command device is set. For example:

4. To release a command device, follow the example below, in which command device 1 is released.

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

5. To change a command device, first release it, then change the LU number. The following example of specifies LU 3 for command device 1.

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

Setting LU mapping

For iSCSI, use the autargetmap command instead of the auhgmap command.

To set up LU Mapping

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

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

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

Defining the configuration definition file

The Configuration Definition file describes the system configuration for CCI. It must be set by the user. The configuration definition file is a text file created and/or edited using any standard text editor. It can be defined from the PC where CCI software is installed.

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

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

To define the configuration definition file

The following example defines the configuration definition file with two instances on the same Windows host.

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

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

- 3. Open **horcm0.conf** using the text editor.
- 4. In the HORCM_MON section, set the necessary parameters.

Important: A value more than or equal to 6000 must be set for poll(10ms). Specifying the value incorrectly may cause resource contention in the internal process, resulting the process temporarily suspending and pausing the internal processing of the array.

5. In the HORCM_CMD section, specify the physical drive (command device) on the array. Figure C-1 shows an example of the horcm0.conf file.

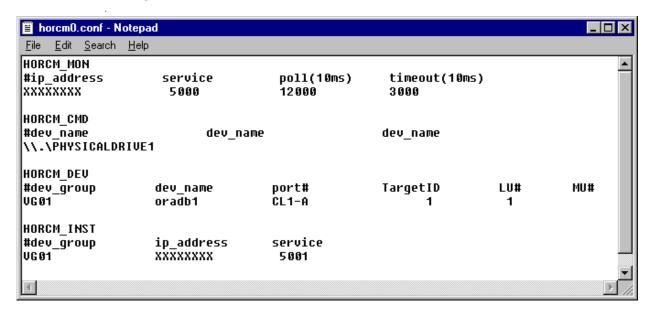


Figure C-1: Horcm0.conf Example

- 6. Save the configuration definition file and use the horemstart command to start CCI.
- 7. Execute the raidscan command; in the result, note the target ID.
- 8. Shut down CCI and open the configuration definition file again.
- 9. In the HORCM_DEV section, set the necessary parameters. For the target ID, enter the ID of the raidscan result. For MU#, do not set a parameter.
- 10.In the HORCM_INST section, set the necessary parameters, and then save (overwrite) the file.
- 11. Repeat Steps 3 to 10 for the **horcm1.conf** file. Figure C-2 shows an example of the horcm1.conf file.

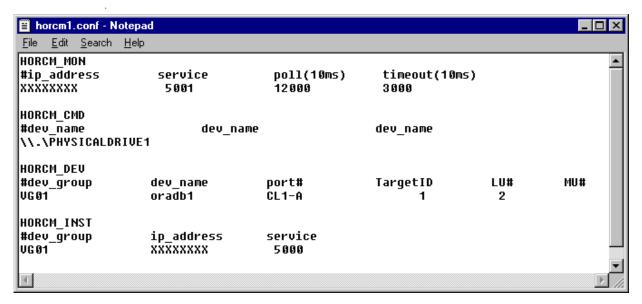


Figure C-2: Horcm1.conf Example

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

```
C:\>cd horcm\etc>
C:\horcm\etc>echo hd1-3 | .\inqraid
Harddisk 1 -> [ST] CL1-A Ser =85000174 LDEV = 0 [HITACHI ] [DF600F-CM ]
Harddisk 2 -> [ST] CL1-A Ser =85000174 LDEV = 1 [HITACHI ] [DF600F ]
HORC = SMPL HOMRCF[MU#0 = NONE MU#1 = NONE MU#2 = NONE]
RAID5[Group 1-0] SSID = 0x0000
Harddisk 3 -> [ST] CL1-A Ser =85000175 LDEV = 2 [HITACHI ] [DF600F ]
HORC = SMPL HOMRCF[MU#0 = NONE MU#1 = NONE MU#2 = NONE]
RAID5[Group 2-0] SSID = 0x0000
C:\horcm\etc>
```

Setting the environment variable

The environment variable must be set up for the execution environment. The following describes an example in which two instances (0 and 1) are configured on the same Windows server.

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

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

2. Execute the horcmstart script, and then execute the pairdisplay command to verify the configuration. For 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) ,Seq#,LDEV#.P/S,Status,Fence, Seq#,P-LDEV# M
VG01 oradb1(L) (CL1-A , 1, 1 )85000174 1.SMPL ---- ---- ----
VG01 oradb1(R) (CL1-A , 1, 2 )85000175 2.SMPL ---- ----- ----
```

Pair operations

This section provides information and instructions for performing TrueCopy operations.

Multiple CCI requests and order of execution

Commands issued from CCI are generally executed in the order they are received in the LUs. However, when more than five commands are issued per controller, the execution order changes.

For example, LUN 0, 1, 2, 3, 5, 6, 7, 8 begin multi-transmission and start the copy operation almost at the same time. LU#4 starts copying when one of the LUs (0 to 3) completes the operation. LU#9 starts copying when one of the LUs (5 to 8) completes the operation.

A single controller provides a data-transmission-area-resource for four pairs multi-transmission. If the controller receives more than that of a command request, the last request will be processed upon completion of the previous ones.

Operations and pair status

Each TrueCopy operation requires a specific pair status. Figure C-3 shows the relationship between status and operations.

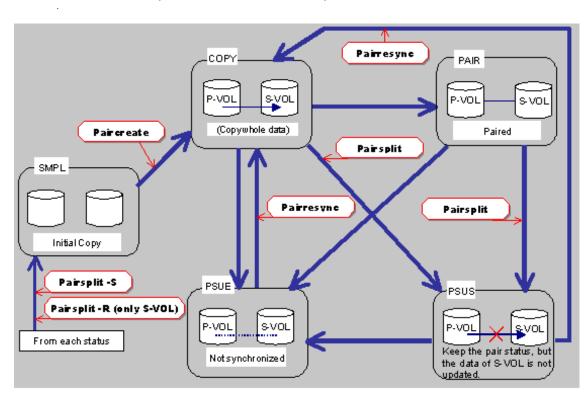


Figure C-3: TrueCopy Pair Status Transitions

Confirming pair status

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

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

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

To confirm TrueCopy pairs

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

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

```
c:\HORCM\etc>pairdisplay -g VG01
Group PairVol(L/R) (Port#,TID, LU) ,Seq#,LDEV#.P/S,Status,Fence, Seq#,P-LDEV# M
VG01 oradb1(L) (CL1-A , 1, 1 )85000174    1.P-VOL COPY Never ,85000175    2 -
VG01 oradb1(R) (CL1-A , 1, 2 )85000175    2.S-VOL COPY Never ,----    1 -
```

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

Creating pairs (paircreate)

In the examples, VG01 is the group name in the configuration definition file.

To create TrueCopy pairs



NOTE: A pair created using CCI and defined in the configuration definition file appear unnamed in the Navigator 2 GUI.

1. Execute the pairdisplay command to verify that the status of the possible volumes to be copied is SMPL. For example:

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

Group PairVol(L/R) (Port#,TID, LU) ,Seq#,LDEV#.P/S,Status,Fence, Seq#,P-LDEV# M

VG01 oradb1(L) (CL1-A , 1, 1 )85000174 1.SMPL ---- ---- ---- ----

VG01 oradb1(R) (CL1-A , 1, 2 )85000175 2.SMPL ---- ---- ----
```

- 2. Execute the paircreate command. The -c option (medium) is recommended when specifying copying pace. See Copy pace on page 6-3 for more information.
- 3. Execute the pairevtwait command to verify that the status of each volume is PAIR. The following example shows the paircreate and pairevtwait commands.

```
C:\HORCM\etc>paircreate -g VG01 -f never -vl -c 10 C:\HORCM\etc>pairevtwait -g VG01 -s pair -t 300 10 pairevtwait : Wait status done.
```

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

```
c:\HORCM\etc>pairdisplay -g VG01
Group PairVol(L/R) (Port#,TID, LU) ,Seq#,LDEV#.P/S,Status,Fence, Seq#,P-LDEV# M
VG01 oradb1(L) (CL1-A , 1, 1 )85000174   1.P-VOL COPY Never ,85000175  2 -
VG01 oradb1(R) (CL1-A , 1, 2 )85000175   2.S-VOL COPY Never ,----  1 -
```

Creating pairs that belong to a consistency group

In the examples, VG01 is the group name in the configuration definition file.

The examples below show how to create a pair using the consistency group (CTG).



NOTE: Consistency groups created using CCI and defined in the configuration definition file are not seen in the Navigator 2 GUI. Also, pairs assigned to groups using CCI appear ungrouped in the Navigator 2 GUI.

1. Execute the pairdisplay command to verify that volume status is SMPL. For example:

2. Execute the paircreate -fg command, then execute the pairevtwait command to verify that the status of each volume is PAIR. For example:

```
C:\HORCM\etc>paircreate -g VG01 -f never -vl -m fg -c 10 C:\HORCM\etc>pairevtwait -g VG01 -s pair -t 300 10 pairevtwait: Wait status done.
```

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

```
C:\HORCM\etc>pairdisplay -g VG01
Group PairVol(L/R) (Port#, TID, LU) , Seq#, LDEV#.P/S, Status, Fence, Seq#, P-LDEV# M
VG01
                (CL1-A , 1, 1 )85000174 1.P-VOL COPY Never ,85000175 2 -
      oradb1(L)
                  (CL1-A , 1, 2 )85000175
VG01
      oradb1(R)
                                           2.S-VOL COPY Never ,----
                                                                       1
                  (CL1-A , 1, 3 )85000174 3.P-VOL COPY Never ,85000175
VG01
     oradb2(L)
                                                                        4 -
VG01 oradb2(R)
                (CL1-A , 1, 4 )85000175
                                           4.S-VOL COPY Never ,----
                                                                        3 -
VG01
    oradb3(L)
                  (CL1-A , 1, 5 )85000174 5.P-VOL COPY Never ,85000175
                  (CL1-A , 1, 6)85000175 6.S-VOL COPY Never ,----
VG01
      oradb3(R)
                                                                        5 -
```

Splitting pairs (pairsplit)

In the examples, VG01 is the group name in the configuration definition file.

Two or more pairs can be split at the same time if they are in the same consistency group.

To split pairs

1. Execute the pairsplit command to split the TrueCopy pair in the PAIR status. For example:

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

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

```
c:\horcm\etc>pairdisplay -g VG01
Group PairVol(L/R) (Port#,TID, LU) ,Seq#,LDEV#.P/S,Status,Fence, Setup-LDEV# M
VG01 oradb1(L) (CL1-A , 1, 1 )85000174  1.P-VOL PSUS Never ,85000175  2 -
VG01 oradb1(R) (CL1-A , 1, 2 )85000175  2.S-VOL SSUS Never ,---- 1 -
```

Resynchronizing pairs (pairresync)

In the examples, VG01 is the group name in the configuration definition file.

To resynchronize pairs

- 1. Execute the pairresync command. Enter between 1 to 15 for copy pace, 1 being slowest (and therefore best I/O performance), and 15 being fastest (and therefore lowest I/O performance). A medium value is recommended.
- 2. Execute the pairevtwait command to verify that the status of each volume is PAIR. The following example shows the pairresync and the pairevtwait commands.

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

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

Suspending pairs (pairsplit -R)

In the examples, VG01 is the group name in the configuration definition file.

To suspend pairs

1. Execute the pairdisplay command to verify that the pair to be suspended is in PAIR status. For example:

```
c:\horcm\etc>pairdisplay -g VG01
Group PairVol(L/R) (Port#,TID, LU) ,Seq#,LDEV#.P/S,Status,Fence, Seq#,P-LDEV# M
VG01 oradb1(L) (CL1-A , 1, 1 )85000174 1.P-VOL PAIR NEVER ,85000175 2 -
VG01 oradb1(R) (CL1-A , 1, 2 )85000175 2.S-VOL PAIR NEVER ,---- 1 -
```

2. Execute the pairsplit -R command to split the pair. For example:

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

3. Execute the pairdisplay command to verify that the P-VOL pair status changed to PSUE. For example:

```
c:\horcm\etc>pairdisplay -g VG01
Group PairVol(L/R) (Port#,TID, LU) ,Seq#,LDEV#.P/S,Status,Fence, Seq#,P-LDEV# M
VG01 oradb1(L) (CL1-A , 1, 1 )85000174 1.P-VOL PSUE NEVER ,85000175 2 -
VG01 oradb1(R) (CL1-A , 1, 2 )85000175 2.SMPL ----- ,----- ---
```

Releasing pairs (pairsplit -S)

In the examples, VG01 is the group name in the configuration definition file.

To release pairs and change status to SMPL

1. Execute the pairsplit -S command to release the pair. For example:

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

2. Execute the pairdisplay command to verify that the pair status changed to SMPL. For example:

Mounting and unmounting a volume

When Windows 2000 Server and Windows Server 2008 are used, the CCI mount command is required to mount or un-mount a volume. The GUID, which is displayed by mountvol command, is needed as an argument when using the mount command. For more information, see the *Hitachi AMS Command Control Interface (CCI) Reference Guide*.



Cascading with ShadowImage

In a cascaded system the TrueCopy P-VOL or S-VOL is shared with a ShadowImage P-VOL or S-VOL. Cascading is usually done to provide a backup that can be used if the volumes are damaged or have inconsistent data. This appendix discusses the configurations, operations, and statuses that are allowed in the cascaded systems.

- Cascade overview
- Cascade configurations
- Replication operations supported
- ☐ Read/write operations, combined status supported
- Guidelines and restrictions
- ☐ Creating a backup with ShadowImage

Cascade overview

TrueCopy's main function is to maintain a copy of the production volume in order to fully restore the P-VOL in the event of a disaster. A ShadowImage backup is another copy—of either the local production volume or the remote S-VOL. A backup ensures that the TrueCopy system:

- Has access to reliable data that can be used to stabilize inconsistencies between the P-VOL and S-VOL, which can result when a sudden outage occurs.
- Can complete the subsequent recovery of the production storage system.

Cascading TrueCopy with ShadowImage ensures the following:

- When ShadowImage is cascaded on the local side, TrueCopy operations can be conducted from the local ShadowImage S-VOL. In this case, the latency associated with the TrueCopy backup is lowered, improving host I/O performance.
- When ShadowImage is cascaded on the remote side, data in the ShadowImage S-VOL can be used as a backup for the TrueCopy S-VOL, which may be required in the event of failure during a TrueCopy resynchronization. The backup data is used to restore the TrueCopy S-VOL, if necessary, from which the local P-VOL can be restored.
- A full-volume copy can be used development, reporting, and so on.

For a discussion on impact to I/O performance versus data recovery, see Performance, data recovery with ShadowImage on page 2-6.

Cascade configurations

Cascade configurations can consist of P-VOLs and S-VOLs, in both TrueCopy and ShadowImage.

The following sections show supported configurations.

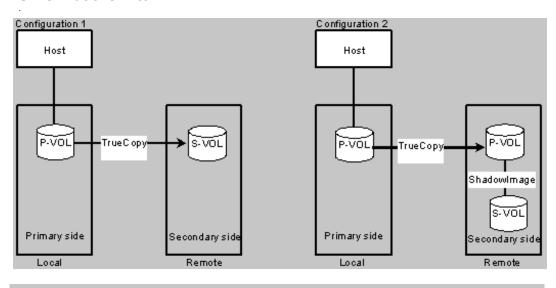


NOTE: Cascading TrueCopy with another TrueCopy system or TrueCopy Extended Distance is not supported.

Configurations with ShadowImage P-VOLs

The ShadowImage P-VOL can be shared with the TrueCopy P-VOL or S-VOL.

Figure D-1 shows cascade configurations where the ShadowImage P-VOL to S-VOL ratio is 1-to-1.



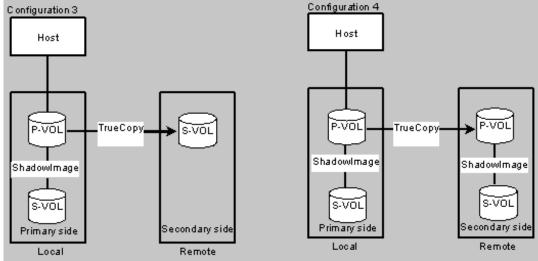


Figure D-1: Cascade Using ShadowImage P-VOL, P-VOL:S-VOL = 1:1

Figure D-2 shows cascade configurations where the ShadowImage P-VOL to S-VOL ratio is 1-to-3.

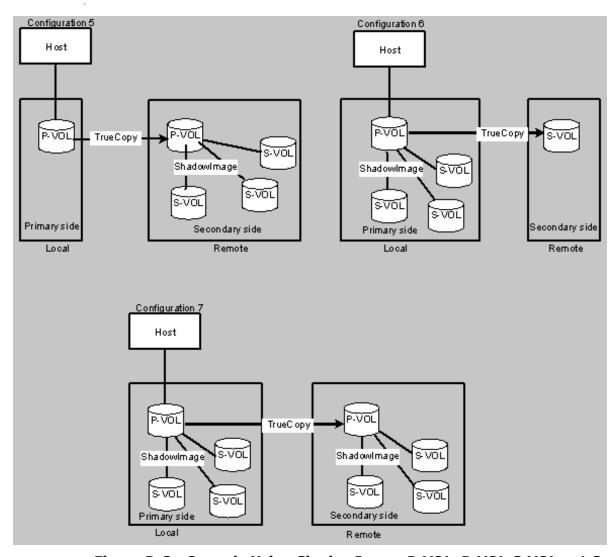


Figure D-2: Cascade Using ShadowImage P-VOL, P-VOL:S-VOL = 1:3

Configurations with ShadowImage S-VOLs

The ShadowImage S-VOL can be shared with the TrueCopy P-VOL only. The TrueCopy remote copy is created from the-local side ShadowImage S-VOL. This results in an asynchronous TrueCopy system.

Figure D-3 shows cascade configurations where the ShadowImage P-VOL to S-VOL ratio is 1-to-1.

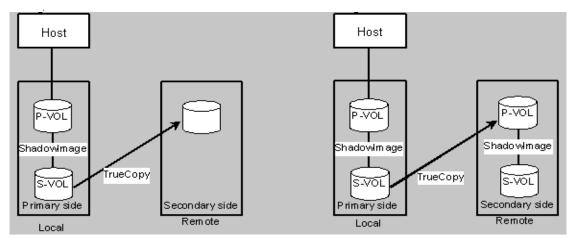


Figure D-3: Cascade with ShadowImage S-VOL — P-VOL: S-VOL = 1:1

Figure D-4 shows cascade configurations where the ShadowImage P-VOL to S-VOL ratio is 1-to-3.

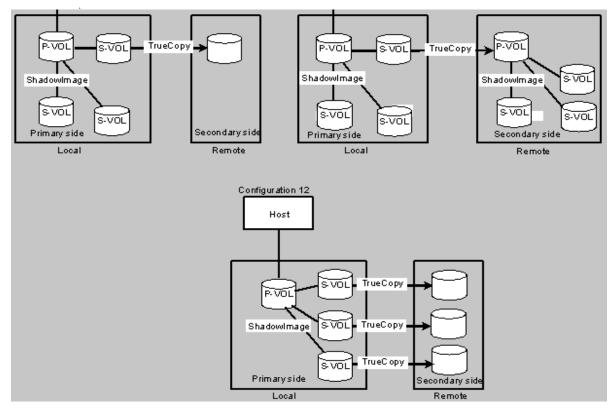


Figure D-4: Cascade with ShadowImage S-VOL— P-VOL: S-VOL = 1:3

Configurations with ShadowImage P-VOLs and S-VOLs

This sections shows configurations in which TrueCopy is cascaded with ShadowImage P-VOLs and S-VOLs.

The lower pair in Figure D-5 shows both ShadowImage P-VOL and S-VOL cascaded with the TrueCopy P-VOL; on the right-side one of the pairs is reversed due to a pair swap.

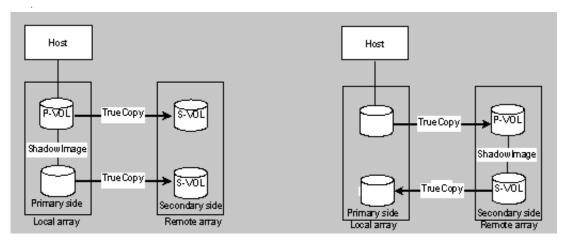


Figure D-5: Swapped Pair Configuration

Figure D-6 shows multiple cascade volumes. The right side configuration shows pairs that have been swapped.

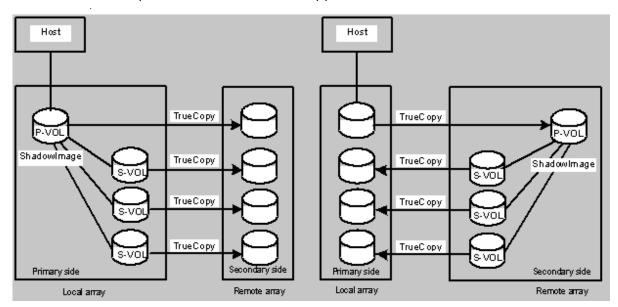


Figure D-6: Multiple Swapped Pair Configuration

Replication operations supported

Generally, a TrueCopy operation can only be performed when the pair is in the appropriate pair status. A split operation can only be performed when the TrueCopy pair is in Paired status, for example.

With cascaded volumes, ShadowImage's pair status must also be taken into account before an operation can be performed.

The tables in this section show TrueCopy and ShadowImage operations that may be performed.

TrueCopy operations supported

Table D-2 and Table D-4 show the TrueCopy operations that can be performed on the shared volume for the displayed ShadowImage status.

Table D-1: When TrueCopy P-VOL/ShadowImage P-VOL Cascaded

ShadowImage	TrueCopy Operations						
Status	Pair Creation	Pair Split	Pair Resync	Swap	Pair Deletion		
Paired	Yes	Yes	Yes	Yes	Yes		
Synchronizing	Yes	Yes	Yes	Yes	Yes		
Synchronizing (Restore)	No	Not possible	No	No	Yes		
Split	Yes	Yes	Yes	Yes	Yes		
Failure	Yes	Yes	Yes	Yes	Yes		
Failure(R)	No	Not possible	No	No	Yes		

Table D-2: When TrueCopy P-VOL/ShadowImage S-VOL Cascaded

ShadowImage	TrueCopy Operations						
Status	Pair Creation	Pair Split	Pair Resync	Swap	Pair Deletion		
Paired	Yes	Not possible	No	No	Yes		
Synchronizing	No	Not possible	No	No	Yes		
Synchronizing (Restore)	No	Not possible	No	No	Yes		
Split	Yes	Yes	Yes	Yes	Yes		
Failure	No	Not possible	No	No	Yes		
Failure(R)	No	Not possible	No	No	Yes		

Table D-3: When TrueCopy S-VOL/ShadowImage P-VOL Cascaded

ShadowImage Status	TrueCopy Operations						
	Pair Creation	Pair Split	Pair Resync	Swap	Pair Deletion		
Paired	Yes	Yes	Yes	Yes	Yes		
Synchronizing	Yes	Yes	Yes	Yes	Yes		
Synchronizing (Restore)	No	Not possible	No	No	Yes		
Split	Yes	Yes	Yes	Yes	Yes		
Failure	Yes	Yes	Yes	Yes	Yes		
Failure(R)	No	Not possible	No	No	Yes		

Table D-4: When TrueCopy S-VOL/ShadowImage S-VOL Cascaded

ShadowImage	TrueCopy Operations						
Status	Pair Creation	Pair Split	Pair Resync	Swap	Pair Deletion		
Paired	Yes	Not possible	No	No	Yes		
Synchronizing	Yes	Not possible	No	No	Yes		
Synchronizing (Restore)	Yes	Not possible	No	No	Yes		
Split	Yes	Yes	Yes	Yes	Yes		
Failure	Yes	Not possible	No	No	Yes		
Failure(R)	Yes	Not possible	No	No	Yes		

ShadowImage operations supported

Table D-6 and Table D-8 show the ShadowImage operations that can be performed on the shared volume for the displayed TrueCopy status.

Table D-5: When TrueCopy P-VOL/ShadowImage P-VOL Cascaded

		ShadowImage Operations							
TrueCopy Status	Pair Creation	Pair Split	Pair Resync	Pair Restore	Pair Deletion				
Paired	Yes	Yes	Yes	No	Yes				
Synchronizing	Yes	Yes	Yes	No	Yes				
Split	Yes Yes Yes		Yes Yes		Yes				
Failure	Yes	Yes	Yes	No	Yes				

Table D-6: When TrueCopy P-VOL/ShadowImage S-VOL Cascaded

		ShadowImage Operations							
TrueCopy Status	Pair Creation	Pair Split	Pair Resync	Pair Restore	Pair Deletion				
Paired	Yes	Not possible	No	No	Yes				
Synchronizing	Yes	Not possible	No	No	Yes				
Split	Yes	Yes	Yes	Yes	Yes				
Failure	Yes	Not possible	No	Not possible	Yes				

Table D-7: When TrueCopy S-VOL/ShadowImage P-VOL Cascaded

	ShadowImage Operations						
TrueCopy Status	Pair Creation	Pair Split	Pair Resync	Pair Restore	Pair Deletion		
Paired	Yes	Yes	Yes	No	Yes		
Synchronizing	Yes	Yes	Yes	No	Yes		
Split	Yes	Yes	Yes	Yes (S-VOL cannot be restored if Read Only)	Yes		
Failure	Yes	Yes	Yes	No	Yes		

Table D-8: When TrueCopy S-VOL/ShadowImage S-VOL Cascaded

		ShadowImage Operations							
TrueCopy Status	Pair Creation	Pair Split	Pair Resync	Pair Restore	Pair Deletion				
Paired	No	Not possible	No	No	Yes				
Synchronizing	No	Not possible	No	No	Yes				
Split	No	Yes	No	Yes	Yes				
Failure	No	Not possible	No	No	Yes				

Read/write operations, combined status supported

The tables in this section present a matrix of TrueCopy and SnapShot statuses. Read/write for a shared volume is indicated, as well as whether the combined statuses are allowed.

Because ShadowImage supports up to 8 S-VOLs per P-VOL, the status for each individual P-VOL-to-S-VOL pair can differ. The correct status must be determined.

- If the ShadowImage pairs are in Split *and* Failure status, Split status is ordinarily used (see Step 2 for exception).
- If the pairs are in Split and Failure status, and the Failure status occurred during Reverse Resynchronization, Failure status is used.

The following abbreviations are used in the tables:

A Combined status is allowed

N/A Combined status is not allowed.

R/W Read/Write is allowed.

R Read is allowed, write is not allowed.

W Write is allowed, read is not allowed.

Read and Write are not allowed.

Err Pair operation by either system causes an error.

Table D-9 shows when read/write is possible, and the status combinations that are allowed when a TrueCopy P-VOL and a ShadowImage P-VOL are cascaded.

Table D-9: I/O, Status Allowed for Local ShadowImage P-VOLs

			Shadow]	mage P-VOL	Status		
TrueCopy P-VOL Status	Paired (including Paired Internally Synchroni zing)	Synchron izing	Synchronizi R/W	ng (Restore) W	Split (Including Split Pending)	Failure	Failure (Restore)
Paired	A R/W	A R/W	N/A	N/A	A R/W	A R/W	N/A
Synchron izing	A R/W	A R/W	N/A	N/A	A R/W	A R/W	N/A
Split R/W	A R/W	A R/W	A R/W	A W	A R/W	A R/W	Err R/W
Split R	A R	A R	N/A	N/A	A R	A R	N/A
Failure R/W	A R/W	A R/W	Err R/W	Err W	A R/W	Err R/W	Err R/W
Failure R	A R	A R	N/A	N/A	A R	Err R	N/A

Table D-9: I/O, Status Allowed for Local ShadowImage P-VOLs

		ShadowImage P-VOL Status						
TrueCopy P-VOL Status	Paired (including Paired Internally Synchroni zing)	Synchron izing	Synchronizi R/W	ng (Restore) W	Split (Including Split Pending)	Failure	Failure (Restore)	
Failure R/W	A R/W	A R/W	N/A	N/A	A R/W	Err R/W	N/A	

Table D-10 shows combined statuses and read/write operations allowed when a ShadowImage P-VOL = a TrueCopy S-VOL.

Table D-10: I/O, Status Allowed for Remote ShadowImage P-VOL

		ShadowImage P-VOL Status									
TrueCopy S-VOL Status	Paired (Including Paired Internally Synchroniz ing)	Synchron izing	Synchronizin R/W	g (Restore) W	Split (Including Split Pending)	Failure	Failure (Restore)				
Paired	A R	A R	N/A	N/A	A R	A R	N/A				
Synchron izing	A R	A R	N/A	N/A	A R	A R	N/A				
Split R/W	A R/W	A R/W	A R/W	A W	A R/W	A R/W	Err R/W				
Split R	A R	A R	N/A	N/A	A R	A R	N/A				
Failure	A R	A R	N/A	N/A	A R	Err R	N/A				

Failure status excludes logical unit blockages and similar access failures.

Table D-11 shows read/write operations and combined statuses allowed when a ShadowImage S-VOL = a TrueCopy P-VOL. The ShadowImage S-VOL can be shared with the TrueCopy P-VOL only.

Table D-11: I/O, Status Allowed for Local ShadowImage S-VOL

TrueCopy S-VOL Status	ShadowImage P-VOL Status							
	Paired	Synchron izing	Synchronizin R/W	g (Restore) W	Split	Failure	Failure (Restore)	
Paired	N/A	N/A	N/A	A R/W	N/A	N/A	N/A	
Synchronizing	N/A	N/A	N/A	A R/W	N/A	N/A	N/A	

Table D-11: I/O, Status Allowed for Local ShadowImage S-VOL

TrueCopy S-VOL Status	ShadowImage P-VOL Status						
	Paired	Synchron izing	Synchronizing (Restore) R/W W		Split	Failure	Failure (Restore)
Split	Α	А	Α	Α	Err	Err	Α
Split R/W	R	R	R	R/W	R	R/W	R
Split R	Α	А	Α	Α	Err	Err	А
R	R	R	R	R	R	R/ W	R
Failure	Α	Α	Α	Α	Err	Err	Α
	R	R	R	R/W	R	R/W	R
	Α	Α	Α	А	Err	Err	Α
	R	R	R	R	R	R/W	R
	A R/W	A R/W	A R/W	A R/W	Err R/W	Err R/W	A R/W

Guidelines and restrictions

- Create the ShadowImage pair before creating the TrueCopy pair. If the TrueCopy pair already exists, it must be split before the ShadowImage pair is created.
- I/O activity on the local host can be minimized during heavy read/write periods by splitting the TrueCopy and ShadowImage pairs. This stops write activity to the S-VOLs.
- The best time to resynchronize TrueCopy pairs and ShadowImage pairs and maintain them in either one in Paired status is when local host I/O is minimal.
- ShadowImage pair volumes that include DP-VOLs created with Dynamic Provisioning cannot be used with TrueCopy.
- A TrueCopy pair cannot be created for a ShadowImage pair including a DP-VOL created by Dynamic Provisioning. The DP-VOLs cannot be set for a P-VOL or an S-VOL in TrueCopy.

Configurations not supported

The configurations shown in Figure D-7 are not supported with TrueCopy.

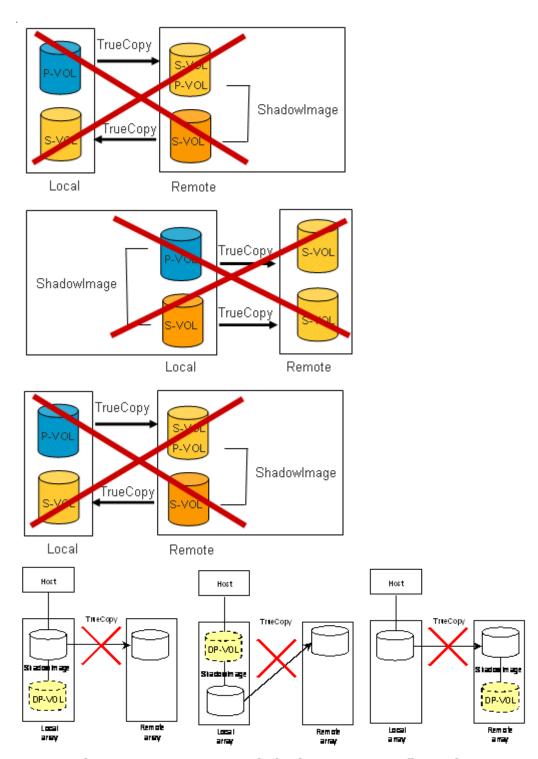


Figure D-7: Unsupported ShadowImage Configurations

Creating a backup with ShadowImage

To back up TrueCopy volume using ShadowImage

- 1. Split the TrueCopy pair. At this time, remaining data buffers on the host are flushed to the P-VOL, then copied to the S-VOL. This ensures a complete and consistent copy.
- 2. Create the ShadowImage pair.
- 3. Split the ShadowImage pair.
- 4. Re-synchronize the TrueCopy pair.

Figure D-8 shows the backup process.

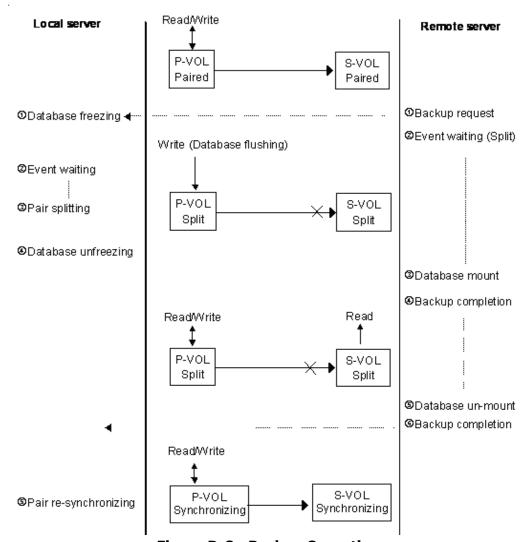


Figure D-8: Backup Operations

For disaster recovery using a ShadowImage backup, see Chapter 7, Disaster recovery.



Cascading with SnapShot

In a cascaded system, the TrueCopy P-VOL or S-VOL is shared with a SnapShot P-VOL or S-VOL. Cascading is usually done to provide a backup that can be used if the volumes are damaged or have inconsistent data. This appendix discusses the configurations, operations, and statuses that are allowed in the cascaded system.

- Cascade overview
- Cascade configurations
- Replication operations supported
- ☐ Read/write, status combinations supported
- Guidelines and restrictions
- ☐ Creating a backup with SnapShot

Cascade overview

SnapShot is cascaded with TrueCopy to:

- Make a backup of the TrueCopy S-VOL on the remote side
- Pair the SnapShot V-VOL on the local side with the TrueCopy S-VOL.
 This results in an asynchronous TrueCopy pair.
- Provide any other traditional use for SnapShot

While a SnapShot V-VOL is smaller than a ShadowImage S-VOL would be, performance when cascading with SnapShot is lower than it would be by cascading with ShadowImage.

This appendix provides the following:

- Supported cascade configurations
- TrueCopy and SnapShot operations allowed
- The combined TrueCopy and SnapShot statuses allowed
- The combined statuses that allow read/write
- Best Practices

Cascade configurations

Cascade configurations can consist of P-VOLs and S-VOLs, in both TrueCopy and SnapShot.

The following sections show supported configurations.



NOTE: Cascading TrueCopy with another TrueCopy system or with TrueCopy Extended Distance is not supported.

Configurations with SnapShot P-VOLs

The SnapShot P-VOL can be cascaded with the TrueCopy P-VOL or S-VOL. Figure E-1 shows supported configurations.

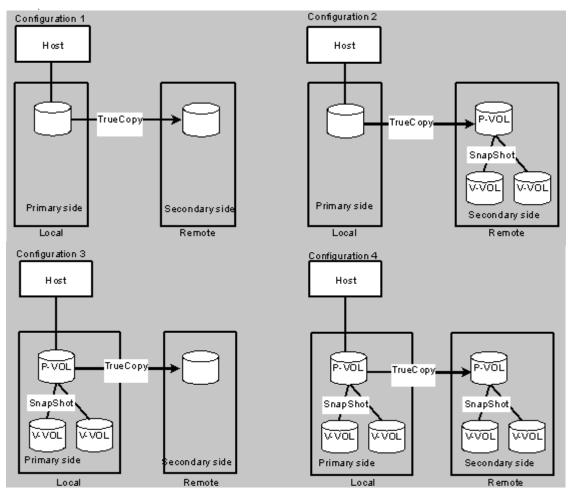


Figure E-1: Cascade with SnapShot P-VOL

Configurations with SnapShot V-VOLs

The SnapShot V-VOL can be cascaded with the TrueCopy P-VOL only. Only one V-VOL in a SnapShot pair may be cascaded.



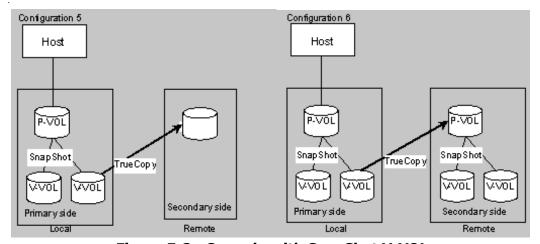


Figure E-2: Cascade with SnapShot V-VOL

Replication operations supported

Generally, a TrueCopy operation can only be performed when the pair is in the appropriate pair status. A split operation can only be performed when the TrueCopy pair is in Paired status, for example.

With cascaded volumes, SnapShot's pair status must also be taken into account before an operation can be performed.

The tables in this section show TrueCopy and SnapShot operations that may be performed.

TrueCopy operations supported

Table E-1 and Table E-2 show the TrueCopy operations that can be performed on the shared volume for the displayed SnapShot status.

Table E-1: When TrueCopy P-VOL, SnapShot V-VOL Cascaded

TrueCopy	SnapShot Pair Status						
Pair Operation	Paired Synchronizing (Restore) Split		Split	Failure	Failure (Restore)		
Creating	No	No	Yes	No	No		
Splitting	Not possible	Not possible	Yes	Not possible	Not possible		
Deleting	Yes	Yes	Yes	Yes	Yes		
Re-synchronizing	No	No	Yes	No	No		
Swapping	No	No	Yes	No	No		

Table E-2: When TrueCopy S-VOL, SnapShot V-VOL Cascaded

TrueCopy Pair Operation	SnapShot Pair Status						
	Paired	Reverse Synchronizing	Split	Failure	Failure (Restore)		
Creating	No	No	No	No	No		
Splitting	Not possible	Not possible	Yes	Not possible	Not possible		
Deleting	Yes	Yes	Yes	Yes	Yes		
Re- synchronizing	No	No	Yes	No	No		
Swapping	No	No	Yes	No	No		

SnapShot operations supported

Table E-3 and Table E-4 show the SnapShot operations that can be performed on the shared volume for the displayed TrueCopy status.

Table E-3: When SnapShot V-VOL, TrueCopy P-VOL Cascaded

SnapShot	TrueCopy Pair Status					
Pair Operation	Paired	Synchronizing	Split	Failure		
Creating	No	No	No	No		
Splitting	Not possible	Not possible	Yes	Yes		
Deleting	Yes	Yes	Yes	Yes		
Resynchronizing	No	No	Yes	Yes		
Restoring	No	No	Yes	Yes		

Table E-4: When SnapShot V-VOL, TrueCopy S-VOL Cascaded

SnapShot	TrueCopy Pair Status					
Pair Operation	Paired Synchronizing		Split	Failure		
Creating	No	No	No	No		
Splitting	Not possible	Not possible	Yes	Not possible		
Deleting	Yes	Yes	Yes	Yes		
Re-synchronizing	No	No	No	No		
Restoring	No	No	No	No		

The scenario for sharing a TrueCopy S-VOL and SnapShot V-VOL (Table E-4) becomes possible when:

- The local SnapShot V-VOL is originally used (cascaded) as the TrueCopy P-VOL.
- The TrueCopy P-VOL is paired with the remote S-VOL, as usual.
- The TrueCopy pair is swapped. Now, the local P-VOL becomes the S-VOL in the pair, though it is still cascaded with the SnapShot V-VOL.

This situation, shown in Figure E-3, is usually temporary.

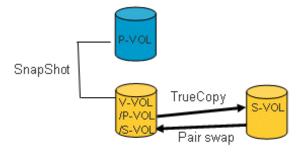


Figure E-3: SnapShot V-VOL Cascaded with TrueCopy S-VOL

Read/write, status combinations supported

The tables in this section present a matrix of TrueCopy and SnapShot statuses. Read/write for a shared volume is indicated, as well as whether the combined statuses are allowed.

- Table E-5 shows status and read/write allowed when the SnapShot P-VOL = a TrueCopy P-VOL.
- Table E-6 shows status and read/write allowed when the SnapShot V-VOL = a TrueCopy P-VOL
- Table E-7 shows status and read/write allowed when the SnapShot P-VOL = a TrueCopy S-VOL

Failure status in these tables does not include LU blockage and other access problems.

The following abbreviations and symbols are used in the tables:

Yes Combined status allowed

No Combined status not allowed.

R/W Read/Write by a host allowed.

R Read by a host allowed, write not allowed.

W Write by a host allowed, read not allowed.

Err Pair operation causes an error.

Table E-5: Read/Write, Status Allowed for Local SnapShot P-VOL

TrueCopy P-VOL Statuses	SnapShot P-VOL Statuses				
	Paired	Synchronizing (Restore)	Split	Failure	Failure (Restore)
Paired	Yes R/W	No	Yes R/W	Yes R/W	No
Synchronizing	Yes R/W	No	Yes R/W	Yes R/W	No
Split R/W	Yes R/W	Yes R/W	Yes R/W	Yes R/W	Err R/W
Split R	Yes R	No	Yes R	Yes R	No
Failure R/W	Yes R/W	Err R/W	Yes R/W	Err R/W	Err R/W
Failure R	Yes R	No	Yes R	Err R	No
Failure R/W	Yes R/W	No	Yes R/W	Err R/W	No

Table E-6: Read/Write, Status Allowed for Remote SnapShot P-VOL

TrueCopy P-VOL Statuses	SnapShot P-VOL Statuses					
	Paired	Synchronizing (Restore)	Split	Failure	Failure (Restore)	
Paired	Yes R	No	Yes R	Yes R	No	
Synchronizing	Yes R	No	Yes R	Yes R	No	
Split R/W	Yes R/W	Yes R/W	Yes R/W	Yes R/W	Err R/W	
Split R	Yes R	No	Yes R	Yes R	No	
Failure R	Yes R	No	Yes R	Err R	No	

Table E-7: Read/Write, Status Allowed for Local SnapShot V-VOL

TrueCopy P-VOL Statuses	SnapShot P-VOL Statuses				
	Paired	Synchronizing (Restore)	Split	Failure	Failure (Restore)
Paired	No	No	Yes R/W	No	No
Synchronizing	No	No	Yes R/W	No	No
Split R/W	Yes	Yes	Yes	Err	Err
	R/W	R/W	R/W	R/W	R/W
Split R	Yes	Yes	Yes	Err	Err
	R/W	R/W	R	R/W	R/W
Failure R/W	Yes	Yes	Yes	Err	Err
	R/W	R/W	R/W	R/W	R/W
Failure R	Yes	Yes	Yes	Err	Err
	R/W R	R/W	R	R/W	R/W
Failure R/W	Yes	Yes	Yes	Err	Err
	R/W	R/W	R/W	R/W	R/W

Guidelines and restrictions

The following provides basic guidelines for cascading SnapShot and TrueCopy:

- SnapShot is not required for TrueCopy.
- A cascaded SnapShot P-VOL may be paired with up to 32 SnapShot V-VOLs.
- Only one V-VOL can be cascaded.
- A SnapShot data pool cannot be cascaded.

- A SnapShot pair must be split when performing a TrueCopy pair operation.
- A TrueCopy pair must be in Split or Failure status when performing a SnapShot pair operation.
- If the TrueCopy S-VOL is restored using the SnapShot V-VOL when TrueCopy is in Synchronizing or Paired status, the TrueCopy data cannot be assured of consistency.
- When cascading a SnapShot V-VOL, create the SnapShot pair before the TrueCopy pair. If the TrueCopy pair already exists, split the pair before creating the SnapShot pair.
- When cascading a SnapShot V-VOL, the SnapShot P-VOL and TrueCopy S-VOL must be on the same RAID level and contain the same number of data disks.
- I/O performance on the local side is lowered when a TrueCopy P-VOL is cascaded with a SnapShot P-VOL.
- When host I/O activity is high, performance is maximized when the TrueCopy pair is split.
- Performance is significantly lower when the SnapShot pair is in Split status.
- Using the V-VOL on the local side results in an asynchronous backup to the remote side.

Configurations not supported

The configurations shown in Figure E-4 are not supported with TrueCopy.

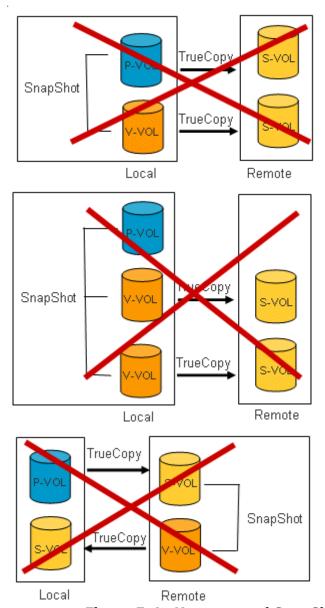


Figure E-4: Unsupported SnapShot Configurations

Creating a backup with SnapShot

To back up TrueCopy volume using SnapShot

- 1. Split the TrueCopy pair. At this time, remaining data buffers on the host are flushed to the P-VOL, then copied to the S-VOL. This ensures a complete and consistent copy.
- 2. Create the SnapShot pair.
- 3. Split the SnapShot pair.
- 4. Re-synchronize the TrueCopy pair.

Figure E-5 shows the backup process.

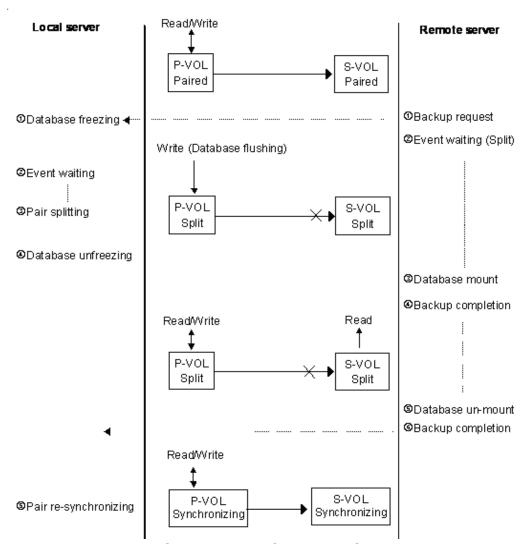


Figure E-5: Backup Operations

For disaster recovery using a SnapShot backup, see Chapter 7, Disaster recovery.

When to create a backup

When a pair is synchronizing, each write to the P-VOL is also sent to the remote S-VOL. The host does not send a write until confirmation is received for the previous write is copied to the S-VOL. This latency impacts host I/O. The best time to synchronize a pair is during off-peak hours. Figure E-6 illustrates host I/O when a pair is split and when a pair is synchronizing.

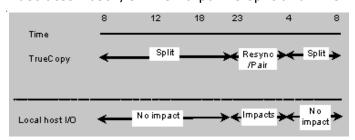


Figure E-6: I/O Performance Impact During Resynchronization



Wavelength Division Multiplexing (WDM) and dark fibre

This appendix discusses WDM and dark fibre, which are used to extend fibre channel remote paths.

■ WDM and dark fibre

WDM and dark fibre

The integrity of a light wavelength remains intact when it is combined with other light wavelengths. Light wavelengths can be combined together in a transmission by multiplexing several optical signals on a dark fiber.

Wavelength Division Multiplexing uses this technology to increase the amount of data that can be transported across distances in a dark fibre extender.

- WDM signifies the multiplexing of several channels of the optical signal.
- Dense WDM (DWDM) signifies the multiplexing of several dozen channels of the optical signal.

Figure F-1 shows an illustration of WDM.

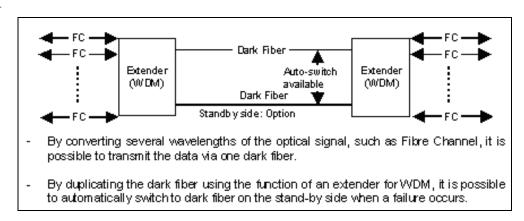


Figure F-1: Wavelength Division Multiplexing

WDM has the following characteristics:

- Response time is extended with WDM. This deterioration is made up by increasing the fibre channel BB-Credit (the number of buffer) without waiting for the response. This requires a switch.
 - If the array is connected directly to an WDM extender without a switch, BB-Credit is 4 or 8. If the array is connected with a switch (Brocade), BB-Credits are 16 and can hold up to 10 km on the standard scale. BB-Credits can be increased to a maximum of 60. By adding the Extended Fabrics option to a switch, BB-Credits can hold up to 100 km.
- For short distances (within several dozen kilometers), both signals of IN and OUT can be transmitted via one dark fiber.
- For long distances (more than several dozen kilometers), an optical amplifier is required to amplify the wavelength between two extenders to prevent attenuation through a fiber. Therefore, dark fibers are required to prepare for IN and OUT respectively. This is illustrated in Figure F-2.

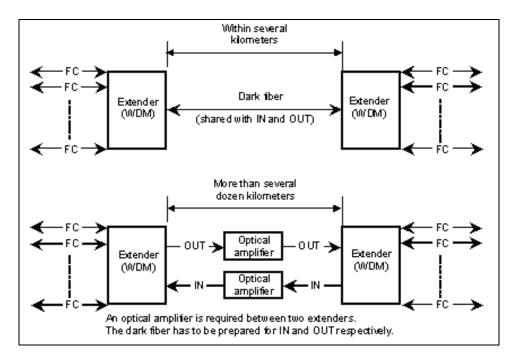


Figure F-2: Dark Fiber with WDM

- The WDM function can also be multiplexed in one dark fiber for Gbps Ethernet.
- If switching is executed during a dark fibre failure, data transfer must be moved to another path, as shown in Figure F-3.

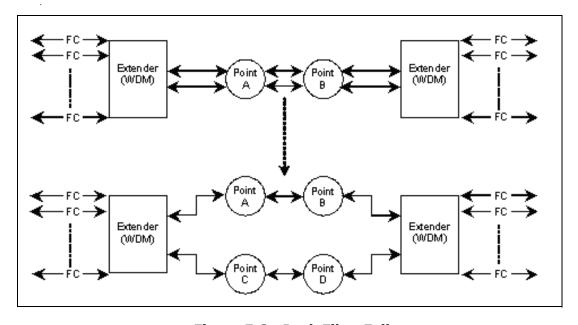


Figure F-3: Dark Fiber Failure

• It is recommend that a second line be set up for monitoring. This allows monitoring to continue if a failure occurs in the dark fiber.

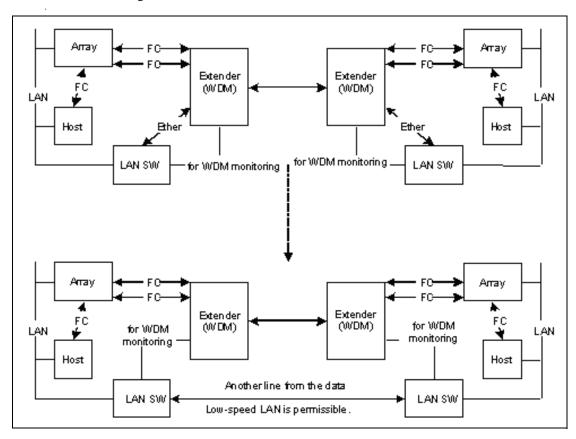


Figure F-4: Line for Monitoring

Glossary

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

Α

array

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

asynchronous

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

В

background copy

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

bps

Bits per second. The standard measure of data transmission speeds.



cache

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

capacity

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

CCI

See command control interface.

CLI

See command line interface.

cluster

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

cluster capacity

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

command control interface (CCI)

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

command devices

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



command line interface (CLI)

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

concurrency of S-VOL

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

concurrent copy

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

configuration definition file

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

consistency group (CTG)

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

consistency of S-VOL

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



CRC

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

CTG

See Consistency Group.

cycle time

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

cycle update

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

D

data pool

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

data volume

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

differential data control

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

differential data copy

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



Differential Management Logical Unit (DMLU)

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

differential-data

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

disaster recovery

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

disk array

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

DMLU

See Differential Management-Logical Unit.

dual copy

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

duplex

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

E

entire copy

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

extent

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



failover

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

fallback

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

Fault tolerance

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

FC

See fibre channel.

fibre channel

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

firmware

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

full duplex

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

G

Gbps

Gigabit per second.

granularity of differential data

Refers to the size or amount of data transferred to the S-VOL during an update cycle. Since only the differential data in the P-VOL is transferred to the S-VOL, the size of data sent to S-VOL is often the same as that of data written to the P-VOL. The amount of differential data that can be managed per write command is limited by the difference between the



number of incoming host write operations (inflow) and outgoing data transfers (outflow).

GUI

Graphical user interface.

Ι

I/O

Input/output.

initial copy

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

initiator ports

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

IOPS

I/O per second.

iSCSI

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

iSNS

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

L

LAN

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

load

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

logical

Describes a user's view of the way data or systems are organized. The opposite of logical is physical, which refers to the real organization of a system. A logical description of a file is that it is a quantity of data



collected together in one place. The file appears this way to users. Physically, the elements of the file could live in segments across a disk.

logical unit

See logical unit number.

logical unit number (LUN)

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

LU

Logical unit.

LUN

See logical unit number.

LUN Manager

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

M

metadata

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

microcode

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

Microsoft Cluster Server

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

mount

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



mount point

The location in your system where you mount your file systems or devices. For a volume that is attached to an empty folder on an NTFS file system volume, the empty folder is a mount point. In some systems a mount point is simply a directory.

P

pair

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

pair splitting

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

pair status

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

paired volume

Two volumes that are paired in a disk array.

parity

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

parity groups

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

peer-to-peer remote copy (PPRC)

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



point-in-time logical copy

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

pool volume

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

primary or local site

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

primary volume (P-VOL)

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

P-VOL

See primary volume.

R

RAID

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

Recovery Point Objective (RPO)

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

Recovery Time Objective (RTO)

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



remote or target site

Maintains mirrored data from the primary site.

remote path

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

remote volume stem

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

resynchronization

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

RPO

See Recovery Point Objective.

RTO

See Recovery Time Objective.

S

SAS

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

SATA

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

secondary volume (S VOL)

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



SMPL

Simplex.

snapshot

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

SNM₂

See Storage Navigator Modular 2.

SSD

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

Storage Navigator Modular 2

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

suspended status

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

S-VOL

See secondary volume.

S-VOL determination

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

Т

target copy

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



virtual volume (V-VOL)

In Copy-on-Write, a secondary volume in which a view of the primary volume (P-VOL) is maintained as it existed at the time of the last



snapshot. The V-VOL contains no data but is composed of pointers to data in the P-VOL and the data pool. The V-VOL appears as a full volume copy to any secondary host.

volume

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

volume copy

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

volume pair

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

V-VOL

See virtual volume.

V-VOLTL

Virtual Volume Tape Library.

W

WMS

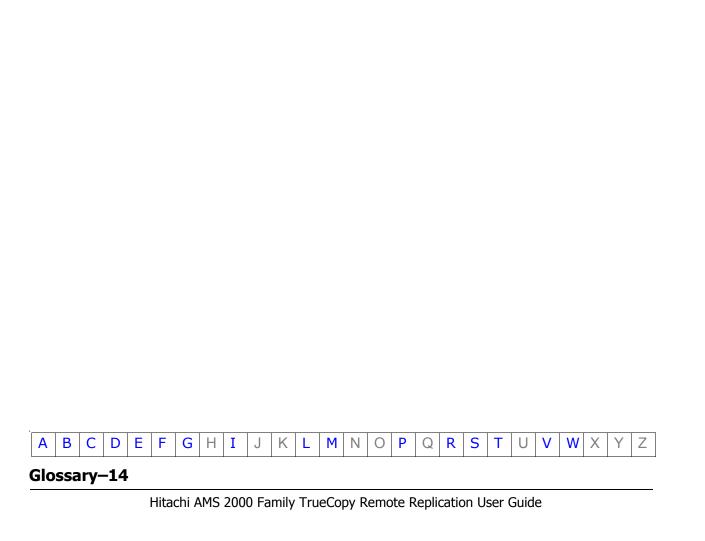
Workgroup Modular Storage.

write order guarantee

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

write workload

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



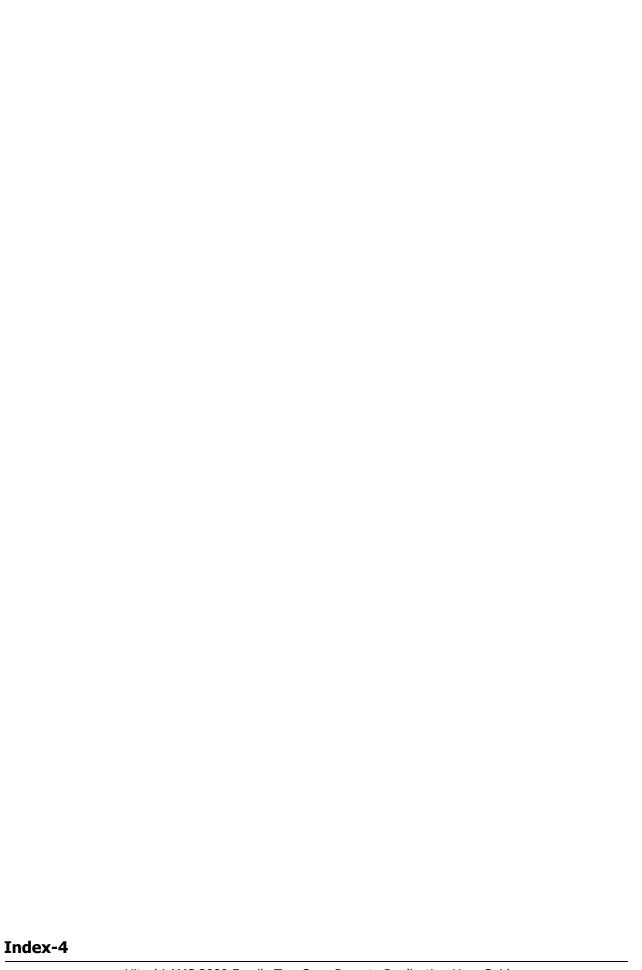
Index

adding a group name 6-7 arrays, supported combinations 3-2 assigning pairs to a consistency group 6-4 B backup script, CLI B-13 bandwidth, calculating 2-2 best practices for data paths 2-22 C Cache Residency Manager A-3 capacity, supported 3-15 cascading overview 3-16 with another TrueCopy system D-3 with ShadowImage D-2 with SnapShot E-2 CCI	delete the remote path B-8 display pair status B-9 install, enable, disable TrueCopy B-2 release DMLU B-4 resynchronize a pair B-11 split a pair B-11 swap a pair B-11 Command Control Interface. See CCI command device releasing C-3 Command Line Interface. See CLI configuration definition file C-4 consistency group creating in GUI 6-4 description 1-4 specifications A-3 copy pace 6-3 creating a pair 6-2
description 1-4 CCI, using to change a command device C-3 create pairs C-9 define config def file C-4 release a command device C-3 release pairs C-12 resynchronize pairs C-11 set environment variable C-6 set LU mapping C-3 set the command device C-2 split pairs C-11 channel extenders 2-13 CLI description 1-4 CLI, using to change pair info B-12 create a pair B-10 create multiple pairs in a group B-10 define DMLU B-4 delete a pair B-12	dark fibre F-1 data fence level 6-3 data path defining 5-5 description 1-4 failure and data recovery 7-2 data paths best practices 2-22 channel extenders 2-13 designing 2-10 preventing blockage 2-22 supported configurations 2-10 data recovery, versus performance 2-4 Data Retention Utility A-3 deleting a pair 6-7 deleting the remote path 6-8 Differential Management-Logical Unit. See DMLU direct connection 2-11 Direct connection, remote path 2-15 DMLU defining 5-4

description 1-4 drive types supported A-3 dynamic disk with Windows 2000 Server 3-7 dynamic disk with Windows Server 2003 3-9 dynamic disk with Windows Server 2008 3-8	memory, reducing A-4 monitoring remote path 8-4 N
E	never fence level 6-3
editing pair information 6-7 enabling/disabling 5-2 extenders F-1	O operating systems, restrictions with 3-6
F	P
fence level 6-3 fibre channel extenders 2-13 fibre channel, port transfer-rate 2-14	Pace field 6-3 pair operations using CCI C-8 pair-monitoring script, CLI B-14 pairs
Group Name field 6-4 group name, adding 6-7 GUI, description 1-4 GUI, using to assign pairs to a Consistency Group 6-4 create a pair 6-2 define DMLU 5-4 define remote path 5-5 delete a pair 6-7 delete remote path 6-8 edit pair information 6-7 install, enable/disable TrueCopy 5-2 monitor pair status 8-3 resynchronize a pair 6-6 split a pair 6-5 swap a pair 6-6	assigning to a consistency group 6-4 creating 6-2 deleting 6-7 resynchronizing 6-6 splitting 6-5 status monitoring, definitions 8-3 swapping 6-6 path failure, recovering the data 7-2 performance info for multiple paths 2-14 performance, with ShadowImage cascaded 2-6 planning LUN expansion 3-3 remote path 2-10 workflow 3-2 port transfer-rate 2-14 Power Saving A-4 production site failure, recovering the data 7-4 P-VOL and S-VOL, definitions 1-2
Н	R
host group, connecting to HP server 3-6 host recognition of P-VOL, S-VOL 3-6 host server failure, recovering the data 7-3 host time-out recommendation 3-6 I I/O performance, versus data recovery 2-4	RAID levels supported A-2 recovering data data path failure 7-2 host server failure 7-3 production site failure 7-4 release a command device C-3 remote array restriction, Sync Cache Ex Mode 3-
installation 5-2 interfaces for TrueCopy 1-4	remote path defining 5-5
K	deleting $6-8$ description $1-4$
key code, key file 5-2	guidelines 2-10 monitoring 8-4
L	planning 2-10
LAN requirements 2-9 logical units, pair recommendations 3-3 LUN expansion 3-3	setup with GUI 5-5 Replication Manager 1-5 Requirements LAN 2-9 resynchronizing a pair 6-6
M	resynchronizing a pair 6-6
MC/Service Guard 3-6	

```
S
scripts
   CLI backup B-13
   CLI pair-monitoring B-14
setting port transfer-rate 2-14
ShadowImage, cascading with D-2
SnapShot
   interfaces 1-5
SnapShot, cascading with E-2
specifications A-2
splitting the pair 6-5
Storage Navigator Modular 2
   description 1-5
supported capacity 3-15
supported data path configurations 2-10
S-VOL
   description 1-2
   updating 6-6
swapping pairs 6-6
switch connection 2-12
Synchronize Cache Execution Mode 3-3
Т
TCE
   remote path configurations 2-10
troubleshooting 8-5
TrueCopy
   defining the remote path (GUI) 5-5
   how it works 1-2
   installing, enabling, disabling 5-2
   interface 1-4
   operations overview 1-5
   pair status monitoring, definitions 8-3
   troubleshooting pair failure 8-5
   troubleshooting path blockage 8-5
   typical environment 1-3
   using unified LUs 3-3
U
unified LUs, in TrueCopy volumes 3-3
uninstalling 5-3
updating the S-VOL 6-6
Volume Migration A-3
W
WAN
   configurations supported 2-16
   general requirements 2-9
   types supported 2-9
WDM F-1
Windows 2000 Server, restrictions 3-7
Windows Server 2003, restrictions 3-8
Windows Server 2008, restrictions 3-8
WOCs, configurations supported 2-18
```

write-workload, measuring 2-2





Hitachi Data Systems

Corporate Headquarters 750 Central Expressway Santa Clara, California 95050-2627 U.S.A.

Phone: 1 408 970 1000

www.hds.com info@hds.com

Asia Pacific and Americas

750 Central Expressway Santa Clara, California 95050-2627 U.S.A.

Phone: 1 408 970 1000 info@hds.com

Europe Headquarters

Sefton Park Stoke Poges Buckinghamshire SL2 4HD United Kingdom Phone: + 44 (0)1753 618000

info.eu@hds.com

MK-97DF8052-08