

Hitachi AMS 2000 Family TrueCopy Extended Distance User Guide

FASTFIND LINKS

- [Document organization](#)
- [Release notes and readme](#)
- [Getting help](#)
- [Table of Contents](#)

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

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

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

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

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

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

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



Table of Contents

| | |
|--|------------|
| Preface | ix |
| Intended audience | x |
| Product version | x |
| Release notes and readme | x |
| Document revision level | x |
| Changes in this release | x |
| Document organization | xi |
| Document conventions | xii |
| Convention for storage capacity values | xiii |
| Accessing product documentation | xiii |
| Getting help | xix |
| Comments | xix |
| | |
| 1 Overview | 1-1 |
| How TCE works | 1-2 |
| Typical environment | 1-2 |
| Volume pairs | 1-3 |
| Data pools | 1-4 |
| Guaranteed write order and the update cycle | 1-4 |
| Extended update cycles | 1-5 |
| Consistency groups | 1-6 |
| Differential Management LUs (DMLU) | 1-7 |
| TCE interfaces | 1-7 |
| | |
| 2 Plan and design — sizing data pools and bandwidth | 2-1 |
| Plan and design workflow | 2-2 |
| Assessing business needs — RPO and the update cycle | 2-2 |
| Measuring write-workload | 2-3 |
| Collecting write-workload data | 2-3 |
| Calculating data pool size | 2-4 |
| Data pool key points | 2-7 |

| | |
|---|------------|
| Determining bandwidth | 2-7 |
| 3 Plan and design — remote path | 3-1 |
| Remote path requirements | 3-2 |
| Management LAN requirements | 3-3 |
| Remote data path requirements | 3-3 |
| WAN optimization controller (WOC) requirements | 3-4 |
| Remote path configurations | 3-5 |
| Fibre channel | 3-5 |
| Direct connection | 3-6 |
| Single FC switch, network connection | 3-7 |
| Double FC switch, network connection | 3-8 |
| Fibre channel extender connection | 3-9 |
| Port transfer rate for Fibre channel | 3-10 |
| iSCSI | 3-11 |
| Direct connection | 3-11 |
| Single LAN switch, WAN connection | 3-12 |
| Multiple LAN switch, WAN connection | 3-13 |
| Single LAN switch, WOC, WAN connection | 3-14 |
| Multiple LAN switch, WOC, WAN connection | 3-15 |
| Multiple array, LAN switch, WOC connection with single WAN | 3-16 |
| Multiple array, LAN switch, WOC connection with two WANs | 3-17 |
| Using the remote path — best practices | 3-18 |
| 4 Plan and design—arrays, volumes, and operating systems | 4-1 |
| Planning workflow | 4-2 |
| Planning arrays—moving data from earlier AMS models | 4-2 |
| Planning logical units for TCE volumes | 4-3 |
| Volume pair, data pool recommendations | 4-3 |
| Operating system recommendations and restrictions | 4-3 |
| Host time-out | 4-3 |
| P-VOL, S-VOL recognition by same host on VxVM, AIX®, LVM | 4-4 |
| HP server | 4-4 |
| Windows Server 2000 | 4-5 |
| Windows Server 2003/2008 | 4-5 |
| Identifying P-VOL and S-VOL LUs on Windows | 4-6 |
| Dynamic Disk with Windows 2003 | 4-6 |
| Maximum supported capacity | 4-9 |
| TCE and SnapShot capacity | 4-10 |
| TCE, SnapShot, ShadowImage concurrent capacity | 4-12 |
| Cache limitations on data and data pool volumes | 4-13 |

- 5 Requirements and specifications5-1**
 - TCE system requirements 5-2
 - Displaying the hardware revision number 5-3
 - TCE system specifications 5-3

- 6 Installation and setup6-1**
 - Installation procedures 6-2
 - Installing TCE 6-2
 - Enabling or disabling TCE. 6-4
 - Uninstalling TCE 6-4
 - Setup procedures 6-5
 - Setting up DMLUs 6-5
 - Setting up data pools. 6-6
 - Setting the LU ownership 6-7
 - Adding or changing the remote port CHAP secret 6-8
 - Setting up the remote path 6-9

- 7 Pair operations7-1**
 - TCE operations 7-2
 - Checking pair status 7-2
 - Creating the initial copy 7-2
 - Prerequisites and best practices for pair creation 7-2
 - TCE setup wizard 7-3
 - Create pair procedure 7-4
 - Splitting a pair 7-5
 - Resynchronizing a pair 7-6
 - Swapping pairs 7-7

- 8 Example scenarios and procedures8-1**
 - CLI scripting procedure for S-VOL backup. 8-2
 - Scripted TCE, SnapShot procedure 8-4
 - Procedure for swapping I/O to S-VOL when maintaining local array 8-9
 - Procedure for moving data to a remote array 8-10
 - Example procedure for moving data 8-11
 - Process for disaster recovery. 8-11
 - Takeover processing 8-11
 - Swapping P-VOL and S-VOL 8-12
 - Failback to the local array 8-12

- 9 Monitoring and maintenance9-1**
 - Monitoring pair status. 9-2
 - Monitoring data pool capacity 9-4

| | |
|---|-------------|
| Monitoring data pool usage | 9-4 |
| Expanding data pool capacity | 9-5 |
| Changing data pool threshold value | 9-5 |
| Monitoring the remote path | 9-6 |
| Changing remote path bandwidth | 9-6 |
| Monitoring cycle time | 9-6 |
| Changing cycle time | 9-7 |
| Changing copy pace | 9-7 |
| Checking RPO — Monitoring P-VOL/S-VOL time difference | 9-8 |
| Routine maintenance | 9-8 |
| Deleting a volume pair | 9-8 |
| Deleting data pools | 9-9 |
| Deleting a DMLU | 9-9 |
| Deleting the remote path | 9-9 |
| TCE tasks before a planned remote array shutdown | 9-10 |
| TCE tasks before updating firmware | 9-10 |
| 10 Troubleshooting | 10-1 |
| Troubleshooting overview | 10-2 |
| Correcting data pool shortage | 10-2 |
| Correcting array problems | 10-3 |
| Delays in settling of S-VOL Data | 10-4 |
| Correcting resynchronization errors | 10-4 |
| Using the event log | 10-6 |
| Miscellaneous troubleshooting | 10-7 |
| A Operations using CLI | A-1 |
| Installation and setup | A-2 |
| Installing | A-3 |
| Enabling and disabling | A-3 |
| Un-installing | A-4 |
| Setting the Differential Management Logical Unit | A-5 |
| Release a DMLU | A-6 |
| Setting the data pool | A-6 |
| Setting the LU ownership | A-8 |
| Setting the cycle time | A-9 |
| Setting mapping information | A-9 |
| Setting the remote port CHAP secret | A-10 |
| Setting the remote path | A-11 |
| Deleting the remote path | A-13 |
| Pair operations | A-14 |
| Displaying status for all pairs | A-14 |
| Displaying detail for a specific pair | A-14 |
| Creating a pair | A-15 |

| | |
|---|------|
| Splitting a pair | A-16 |
| Resynchronizing a pair | A-16 |
| Swapping a pair | A-16 |
| Deleting a pair | A-17 |
| Changing pair information | A-18 |
| Monitoring pair status | A-18 |
| Confirming Consistency Group (CTG) status | A-19 |
| Procedures for failure recovery | A-20 |
| Displaying the event log | A-20 |
| Reconstructing the remote path | A-20 |
| Sample script | A-21 |

B Operations using CCI B-1

| | |
|--|------|
| Setup | B-2 |
| Setting the command device | B-2 |
| Setting LU mapping | B-3 |
| Defining the configuration definition file | B-4 |
| Setting the environment variable | B-6 |
| Pair operations | B-7 |
| Checking pair status | B-7 |
| Creating a pair (paircreate) | B-8 |
| Splitting a pair (pairsplit) | B-8 |
| Resynchronizing a pair (pairresync) | B-9 |
| Suspending pairs (pairsplit -R) | B-9 |
| Releasing pairs (pairsplit -S) | B-10 |
| Splitting TCE S-VOL/SnapShot V-VOL pair (pairsplit -mscas) | B-10 |
| Confirming data transfer when status is PAIR | B-12 |
| Pair creation/resynchronization for each CTG | B-12 |
| Response time of pairsplit command | B-14 |
| Pair, group name differences in CCI and Navigator 2 | B-17 |

C Cascading with SnapShot C-1

| | |
|--|-----|
| Cascade configurations | C-2 |
| Replication operations supported | C-2 |
| TCE operations supported | C-3 |
| SnapShot operations supported | C-4 |
| Status combinations, Read/write supported | C-4 |
| Guidelines and restrictions | C-6 |
| Cascading with SnapShot on the remote side | C-6 |
| TCE, SnapShot behaviors compared | C-7 |

D Installing TCE when Cache Partition Manager is in use D-1

| | |
|--|-----|
| Initializing Cache Partition when TCE and SnapShot are installed | D-2 |
|--|-----|

E Wavelength Division Multiplexing (WDM) and dark fibreE-1
WDM and dark fibreE-2

Glossary

Index



Preface

This document provides instructions for planning, setting up, and operating TrueCopy Extended Distance.

This preface includes the following information:

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

Intended audience

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

Product version

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

Release notes and readme

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

Document revision level

| Revision | Date | Description |
|----------------|---------------|---|
| MK-97DF8054-00 | October 2008 | Initial Release |
| MK-97DF8054-01 | December 2008 | Supersedes and replaces MK-97DF8054-00. |
| MK-97DF8054-02 | March 2009 | Supersedes and replaces MK-97DF8054-01. |
| MK-97DF8054-03 | June 2009 | Supersedes and replaces MK-97DF8054-02. |
| MK-97DF8054-04 | August 2009 | Supersedes and replaces MK-97DF8054-03. |
| MK-97DF8054-05 | November 2009 | Supersedes and replaces MK-97DF8054-04. |
| MK-97DF8054-06 | January 2010 | Supersedes and replaces MK-97DF8054-05. |
| MK-97DF8054-07 | April 2010 | Supersedes and replaces MK-97DF8054-06. |

Changes in this release

The following changes occur in this release:

- Updated section [Planning arrays—moving data from earlier AMS models \(page 4-2\)](#).
- Updated section [HP server \(page 4-4\)](#).
- Updated [Table 5-1 \(page 5-2\)](#).
- Updated [Table 5-2 \(page 5-3\)](#).
- Added a new section [Setting the LU ownership \(page 6-7\)](#).
- Updated section [Setting up the remote path \(page 6-9\)](#).
- Added a new section [Setting the LU ownership \(page A-8\)](#).




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 instructions for designing, planning, implementing, using, monitoring, and troubleshooting TrueCopy Extended Distance (TCE) |
| Chapter 2, Plan and design – sizing data pools and bandwidth | Provides instructions for measuring write-workload, calculating data pool size and bandwidth. |
| Chapter 3, Plan and design – remote path | Provides network and bandwidth requirements, and supported remote path configurations |
| Chapter 4, Plan and design—arrays, volumes, and operating systems | Provides the information you need to prepare your arrays and volumes for TCE operations. |
| Chapter 5, Requirements and specifications | This chapter provides TCE system requirements and specifications. Cautions and restrictions are also provided. |
| Chapter 6, Installation and setup | Provides TCE installation and setup procedures using the Navigator 2 GUI. |
| Chapter 7, Pair operations | Provides procedures for performing basic TCE operations using the Navigator 2 GUI. |
| Chapter 8, Example scenarios and procedures | Provides backup, data moving, and disaster recovery scenarios and procedures. |
| Chapter 9, Monitoring and maintenance | Provides monitoring and maintenance information. |
| Chapter 10, Troubleshooting | Provides troubleshooting information. |
| Appendix A, Operations using CLI | Provides detailed Command Line Interface instructions for configuring and using TCE. |
| Appendix B, Operations using CCI | Provides detailed Command Line Interface instructions for configuring and using TCE. |
| Appendix C, Cascading with SnapShot | Provides supported configurations, operations, etc. for cascading TCE with SnapShot. |
| Appendix D, Installing TCE when Cache Partition Manager is in use | Provides required information when using Cache Partition Manager. |
| Wavelength Division Multiplexing (WDM) and dark fibre | Provides a discussion of WDM and dark fibre for channel extender. |
| Glossary | Provides definitions for terms and acronyms found in this document. |
| Index | Provides links and locations to specific information in this document. |

Document conventions

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

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

The following typographic conventions are used in this document.

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

Convention for storage capacity values

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

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

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

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

Accessing product documentation

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

This documentation set consists of the following documents.


Release notes

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



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

Installation and getting started

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

 **AMS2100/2300 Getting Started Guide**, MK-98DF8152

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

 **AMS2500 Getting Started Guide**, MK-97DF8032

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

AMS 2000 Family Site Preparation Guide, MK-98DF8149

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

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

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

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

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

Storage and replication features

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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


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

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

Storage Navigator Modular 2 online help

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

Hardware maintenance and operation

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

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

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

AMS 2500 Storage System Hardware Guide, MK-97DF8007

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

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

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

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

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

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

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

Command and Control (CCI)

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

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

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

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

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

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

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

Command Line Interface (CLI)

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

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

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

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

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

Dynamic Replicator documentation

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

Dynamic Replicator - Scout Release Notes, RN-99DF8211

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Getting help

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

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

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

Comments

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

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

Overview

This manual provides instructions for designing, planning, implementing, using, monitoring, and troubleshooting TrueCopy Extended Distance (TCE). This chapter consists of:

- ❑ [How TCE works](#)
- ❑ [Typical environment](#)
- ❑ [TCE interfaces](#)

How TCE works

With TrueCopy Extended Distance (TCE), you create a copy of your data at a remote location. After the initial copy is created, only changed data transfers to the remote location.

You create a TCE copy when you:

- Select a volume on the production array that you want to replicate
- Create a volume on the remote array that will contain the copy
- Establish a Fibre Channel or iSCSI link between the local and remote arrays
- Make the initial copy across the link on the remote array.

During and after the initial copy, the primary volume on the local side continues to be updated with data from the host application. When the host writes data to the P-VOL, the local array immediately returns a response to the host. This completes the I/O processing. The array performs the subsequent processing independently from I/O processing.

Updates are periodically sent to the secondary volume on the remote side at the end of the “update cycle”. This is a time period established by the user. The cycle time is based on the recovery point objective (RPO), which is the amount of data in time (2-hours’ worth, 4 hour’s worth) that can be lost after a disaster, until the operation is irreparably damaged. If the RPO is two hours, the business must be able to recover all data up to two hours before the disaster occurred.

When a disaster occurs, storage operations are transferred to the remote site and the secondary volume becomes the production volume. All the original data is available in the S-VOL, from the last completed update. The update cycle is determined by your RPO and by measuring write-workload during the TCE planning and design process.

For a detailed discussion of the disaster recovery process using TCE, please refer to [Process for disaster recovery on page 8-11](#).

Typical environment

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

- 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 on the local array that is to be copied to the secondary volume on the remote side.
- A differential management LU on local and remote arrays, which hold TCE information when the array is powered down

- Interface and command software, used to perform TCE operations. Command software uses a command device (volume) to communicate with the arrays.

Figure 1-1 shows a typical TCE environment.

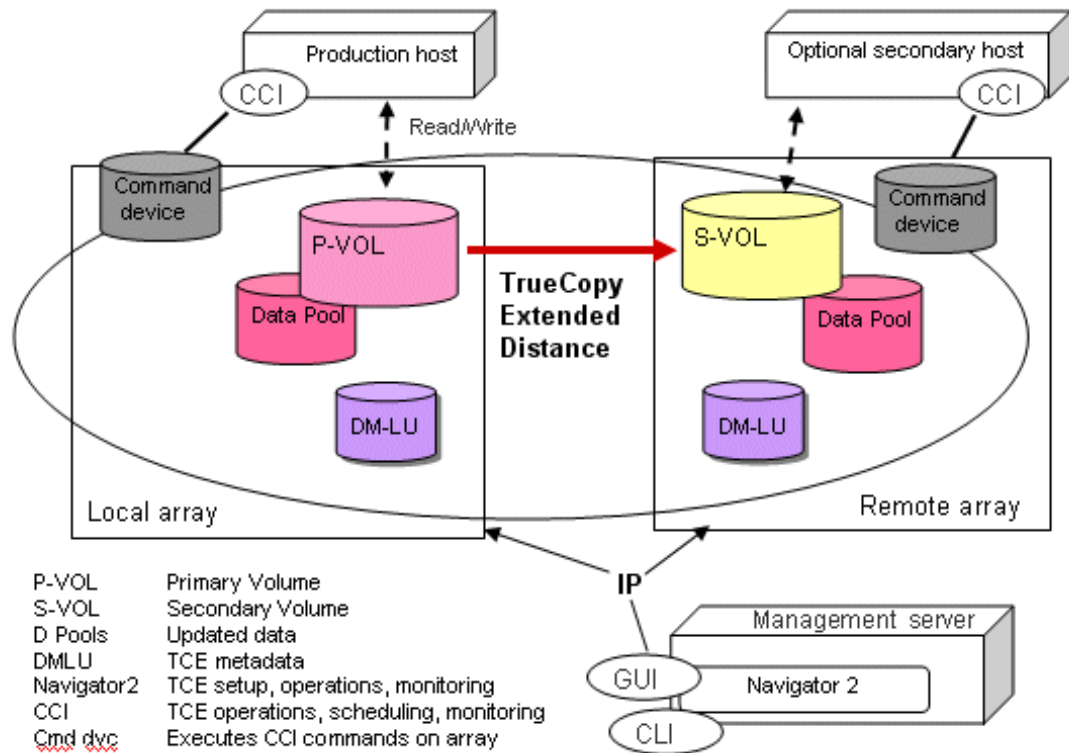


Figure 1-1: Typical TCE Environment

Volume pairs

When the initial TCE copy is completed, the production and backup volumes are said to be "Paired". The two paired volumes are referred to as the primary volume (P-VOL) and secondary volume (S-VOL). Each TCE pair consists of one P-VOL and one S-VOL. When the pair relationship is established, data flows from the P-VOL to the S-VOL.

While in the Paired status, new data is written to the P-VOL and then periodically transferred to the S-VOL, according to the user-defined update cycle.

When a pair is "split", the data flow between the volumes stops. At this time, all the differential data that has accumulated in the local array since the last update is copied to the S-VOL. This insures that its data is the same as the P-VOL's and is consistent and usable data.

During normal TCE operations, the P-VOL remains available for read/write from the host. When the pair is split, the S-VOL also is available for read/write operations from a host.

Data pools

Data from the host is continually updated to the P-VOL, as it occurs. The data pool on the local side stores the changed data that accumulates before the next the update cycle. The local data pool is used to update the S-VOL.

Data that accumulates in the data pool is referred to as differential data because it contains the difference data between the P-VOL and S-VOL.

The data in the S-VOL following an update is complete, consistent, and usable data. When the next update is to begin, this consistent data is copied to the remote data pool. This data pool is used to maintain previous point-in-time copies of the S-VOL, which are used in the event of failback.

Guaranteed write order and the update cycle

S-VOL data must have the same order in which the host updates the P-VOL. When write order is guaranteed, the S-VOL has data consistency with the P-VOL.

As explained in the previous section, data is copied from the P-VOL and local data pool to the S-VOL following the update cycle. When the update is complete, S-VOL data is identical to P-VOL data *at the end of the cycle*. Since the P-VOL continues to be updated while and after the S-VOL is being updated, S-VOL data and P-VOL data are not identical.

However, the S-VOL and P-VOL can be made identical when the pair is split. During this operation, all differential data in the local data pool is transferred to the S-VOL, as well as all cached data in host memory. This cached data is flushed to the P-VOL, then transferred to the S-VOL as part of the split operation, thus ensuring that the two are identical.

If a failure occurs during an update cycle, the data in the update is inconsistent. Write order in the S-VOL is nevertheless guaranteed — at the point-in-time of the previous update cycle, which is stored in the remote data pool.

[Figure 1-2](#) shows how S-VOL data is maintained at one update cycle back of P-VOL data.

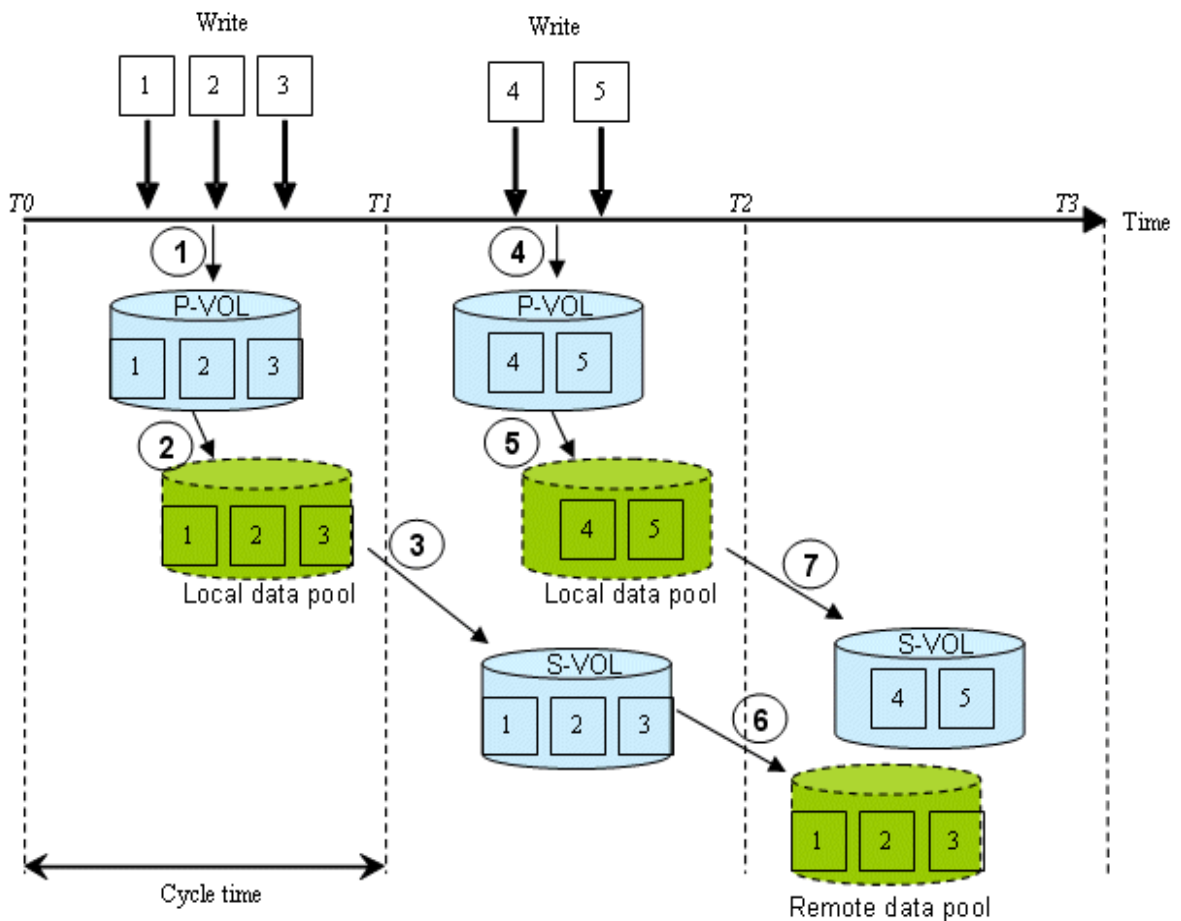


Figure 1-2: Update Cycles and Differential Data

Extended update cycles

If inflow to the P-VOL increases, all of the update data may not be sent within the cycle time. This causes the cycle to extend beyond the user-specified cycle time.

As a result, more update data in the P-VOL accumulates to be copied at the next update. Also, the time difference between the P-VOL data and S-VOL data increases, which degrades the recovery point value. In [Figure 1-2](#), if a failure occurs at the primary site immediately before time T_3 , for example, data consistency in the S-VOL during takeover is P-VOL data at time T_1 .

When inflow decreases, updates again complete within the cycle time. Cycle time should be determined according to a realistic assessment of write workload, as discussed in [Chapter 2, Plan and design — sizing data pools and bandwidth](#).

Consistency groups

Application data often spans more than one volume. With TCE, 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 TCE operations to be performed on all volumes in the group concurrently. Write order in secondary volumes is guaranteed across application logical volumes.

Figure 1-3 shows TCE operations with a consistency group.

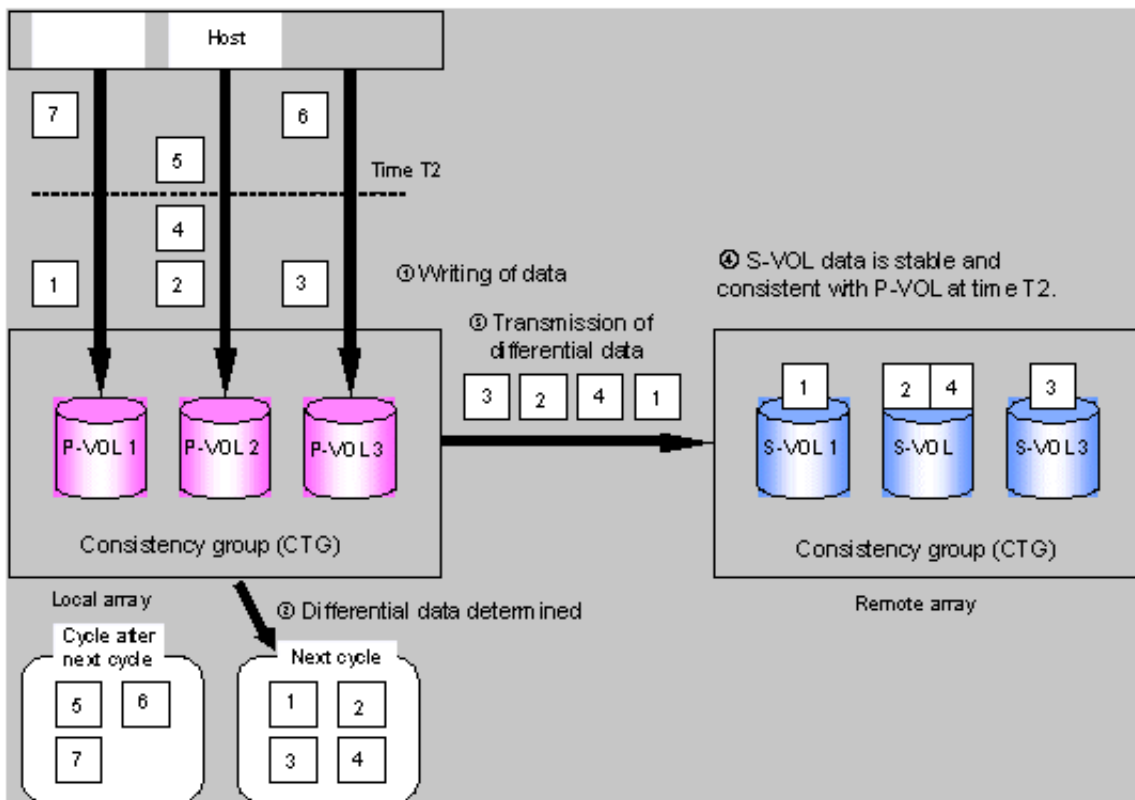


Figure 1-3: TCE Operations with Consistency Groups

In this illustration, observe the following:

- The P-VOLs belong to the same consistency group. The host updates the P-VOLs as required (1).
- The local array identifies the differential data in the P-VOLs when the cycle is started (2) in an atomic manner. The differential data of the group of the P-VOLs are determined at time T2.
- The local array transfers the differential data to the corresponding S-VOLs (3). When all differential data is transferred, each S-VOL is identical to its P-VOL at time T2 (4).
- If pairs are split or deleted, the local array stops the cycle update for the consistency group. Differential data between P-VOLs and S-VOLs is determined at that time. All differential data is sent to the S-VOLs, and the split or delete operations on the pairs completes. S-VOLs maintain data consistency across pairs in the consistency group.

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 hidden from a host. User setup is required on the local and remote arrays.

TCE interfaces

TCE can be setup, used and monitored using of the following interfaces:

- The GUI (Hitachi Storage Navigator Modular 2 Graphical User Interface), which is a browser-based interface from which TCE 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 TCE 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), which is 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, and, on Windows 2000 Server, mount/unmount operations.

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

Plan and design — sizing data pools and bandwidth

This chapter provides instructions for measuring write-workload and sizing data pools and bandwidth.

- ❑ [Plan and design workflow](#)
- ❑ [Assessing business needs — RPO and the update cycle](#)
- ❑ [Measuring write-workload](#)
- ❑ [Calculating data pool size](#)
- ❑ [Determining bandwidth](#)

Plan and design workflow

You design your TCE system around the write-workload generated by your host application. Data pools and bandwidth must be sized to accommodate write-workload. This chapter helps you perform these tasks as follows:

- Assess business requirements regarding how much data your operation must recover in the event of a disaster.
- Measure write-workload. This metric is used to ensure that data pool size and bandwidth are sufficient to hold and pass all levels of I/O.
- Calculate data pool size. Instructions are included for matching data pool capacity to the production environment.
- Calculate remote path bandwidth: This will make certain that you can copy your data to the remote site within your update cycle.

Assessing business needs — RPO and the update cycle

In a TCE system, the S-VOL will contain nearly all of the data that is in the P-VOL. The difference between them at any time will be the differential data that accumulates during the TCE update cycle.

This differential data accumulates in the local data pool until the update cycle starts, then it is transferred over the remote data path.

Update cycle time is a uniform interval of time during which differential data copies to the S-VOL. You will define the update cycle time when creating the TCE pair.

The update cycle time is based on:

- the amount of data written to your P-VOL
- the maximum amount of data loss your operation could survive during a disaster.

The data loss that your operation can survive and remain viable determines to what point in the past you must recover.

An hour's worth of data loss means that your recovery point is one hour ago. If disaster occurs at 10:00 am, upon recovery your restart will resume operations with data from 9:00 am.

Fifteen minutes worth of data loss means that your recovery point is 15 minutes prior to the disaster.

You must determine your recovery point objective (RPO). You can do this by measuring your host application's write-workload. This shows the amount of data written to the P-VOL over time. You or your organization's decision-makers can use this information to decide the number of business transactions that can be lost, the number of hours required to key in lost data and so on. The result is the RPO.

Measuring write-workload

Bandwidth and data pool size are determined by understanding the write-workload placed on the primary volume from the host application.

- After the initial copy, TCE only copies changed data to the S-VOL.
- Data is changed when the host application writes to storage.
- Write-workload is a measure of changed data over a period of time.

When you know how much data is changing, you can plan the size of your data pools and bandwidth to support your environment.

Collecting write-workload data

Workload data is collected using your operating system's performance monitoring feature. Collection should be performed during the busiest time of month, quarter, and year so you can be sure your TCE implementation will support your environment when demand is greatest. The following procedure is provided to help you collect write-workload data.

To collect workload data

1. Using your operating system's performance monitoring software, collect the following:
 - Disk-write bytes-per-second for every physical volume that will be replicated.
 - Collect this data at 10 minute intervals and over as long a period as possible. Hitachi recommends a 4-6 week period in order to accumulate data over all workload conditions including times when the demands on the system are greatest.
2. At the end of the collection period, convert the data to MB/second and import into a spreadsheet tool. In [Figure 2-1, Write-Workload Spreadsheet](#), column C shows an example of collected raw data over 10-minute segments.

| | A | B | C | D | E | F | G | H |
|----|----------|-------------------------|---|--------------------|--------------------|---------------------|---|-------------------------|
| | Sample # | Time - 10 min. segments | Raw Data - MB/sec collected per 10 min. segment | 30 Min Rolling Ave | 60 Min Rolling Ave | 24 Hour Rolling Avg | | Raw Data Project Growth |
| 1 | | | | | | | | |
| 2 | | | | | | | | 15.00% |
| 3 | 1 | 0:00 | 0.863 | | | | | 0.992 |
| 4 | 2 | 0:10 | 2.81 | | | | | 3.232 |
| 5 | 3 | 0:20 | 0.858 | | 1.24 | 2.17 | | 0.987 |
| 6 | 4 | 0:30 | 0.45 | | 1.56 | 2.17 | | 0.518 |
| 7 | 5 | 0:40 | 0.813 | | 1.36 | 2.17 | | 0.935 |
| 8 | 6 | 0:50 | 1.63 | | 1.56 | 2.17 | | 1.875 |
| 9 | 7 | 1:00 | 2.81 | | 2.24 | 2.17 | | 3.232 |
| 10 | 8 | 1:10 | 1.58 | | 2.30 | 2.18 | | 1.817 |
| 11 | 9 | 1:20 | 2.05 | | 2.36 | 2.18 | | 2.358 |
| 12 | 10 | 1:30 | 4.57 | | 2.48 | 2.17 | | 5.256 |
| 13 | 11 | 1:40 | 1.13 | | 2.66 | 2.17 | | 1.300 |
| 14 | 12 | 1:50 | 2 | | 2.72 | 2.16 | | 2.300 |
| 15 | 13 | 2:00 | 3.56 | | 2.30 | 2.16 | | 4.094 |
| 16 | 14 | 2:10 | 2.64 | | 2.26 | 2.19 | | 3.036 |
| 17 | 15 | 2:20 | 2.39 | | 2.34 | 2.19 | | 2.749 |

Figure 2-1: Write-Workload Spreadsheet

Fluctuations in write-workload can be seen from interval to interval. To calculate data pool size, the interval data will first be averaged, then used in an equation. (Your spreadsheet at this point would have only rows B and C populated.)

Calculating data pool size

In addition to write-workload data, cycle time must be known. Cycle time is the frequency that updates are sent to the remote array. This is a user-defined value that can range from 30 seconds to 1 hour. The default cycle time is 5-minutes (300 seconds). If consistency groups are used, the minimum must be 30 seconds for one CTG, increasing 30 seconds for each additional CTG, up to 16. Since the data pool stores all updated data that accumulates during the cycle time, the longer the cycle time, the larger the data pool must be. For more information on cycle time, see the discussion in [Assessing business needs — RPO and the update cycle on page 2-2](#), and also [Changing cycle time on page 9-7](#).

To calculate TCE data pool capacity

- Using write-workload data imported into a spreadsheet tool and your cycle time, calculate write rolling-averages, as follows. (Most spreadsheet tools have an average function.)
 - If cycle time is 1 hour, then calculate 60 minute rolling averages. Do this by arranging the values in six 10-minute intervals.
 - If cycle time is 30 minutes, then calculate 30 minute rolling averages, arranging the values in three 10-minute intervals.

Example rolling-average procedure for cycle time in Microsoft Excel

Cycle time in the following example is 1 hour; rolling averages are calculated using six 10-minute intervals.

- a. After converting workload data into the spreadsheet ([Figure 2-1, Write-Workload Spreadsheet](#)), in cell E4 type, `=average(b2:b7)`, and press Enter.

This instructs the tool to calculate the average value in cells B2 through B7 (six 10-minute intervals) and populate cell E4 with that data. (*The calculations used here are for example purposes only. Base your calculations on your cycle time.*)

- b. Copy the value that displays in E4.
- c. Highlight cells E5 to the E cell in the last row of workload data in the spreadsheet.
- d. Right-click the highlighted cells and select the Paste option.

Excel maintains the logic and increments the formula values initially entered in E4. It then calculates all the 60-minute averages for every 10-minute increment, and populates the E cells, as shown in [Figure 2-2](#).

| | A | B | C | D | E |
|----|----------|-------------------------|---|--------------------|--------------------|
| | Sample # | Time - 10 min. segments | Raw Data - MB/sec collected per 10 min. segment | 30 Min Rolling Ave | 60 Min Rolling Ave |
| 1 | | | | | |
| 2 | | | | | |
| 3 | 1 | 0:00 | 0.863 | | |
| 4 | 2 | 0:10 | 2.81 | | |
| 5 | 3 | 0:20 | 0.858 | | 1.24 |
| 6 | 4 | 0:30 | 0.45 | | 1.56 |
| 7 | 5 | 0:40 | 0.813 | | 1.36 |
| 8 | 6 | 0:50 | 1.63 | | 1.56 |
| 9 | 7 | 1:00 | 2.81 | | 2.24 |
| 10 | 8 | 1:10 | 1.58 | | 2.30 |
| 11 | 9 | 1:20 | 2.05 | | 2.36 |
| 12 | 10 | 1:30 | 4.57 | | 2.48 |
| 13 | 11 | 1:40 | 1.13 | | 2.66 |
| 14 | 12 | 1:50 | 2 | | 2.72 |
| 15 | 13 | 2:00 | 3.56 | | 2.30 |
| 16 | 14 | 2:10 | 2.64 | | 2.26 |
| 17 | 15 | 2:20 | 2.39 | | 2.34 |
| 18 | 16 | 2:30 | 2.06 | | 2.17 |

Figure 2-2: Rolling Averages Calculated Using 60 Minute Cycle Time

For another perspective, you can graph the data, as shown in [Figure 2-3](#).

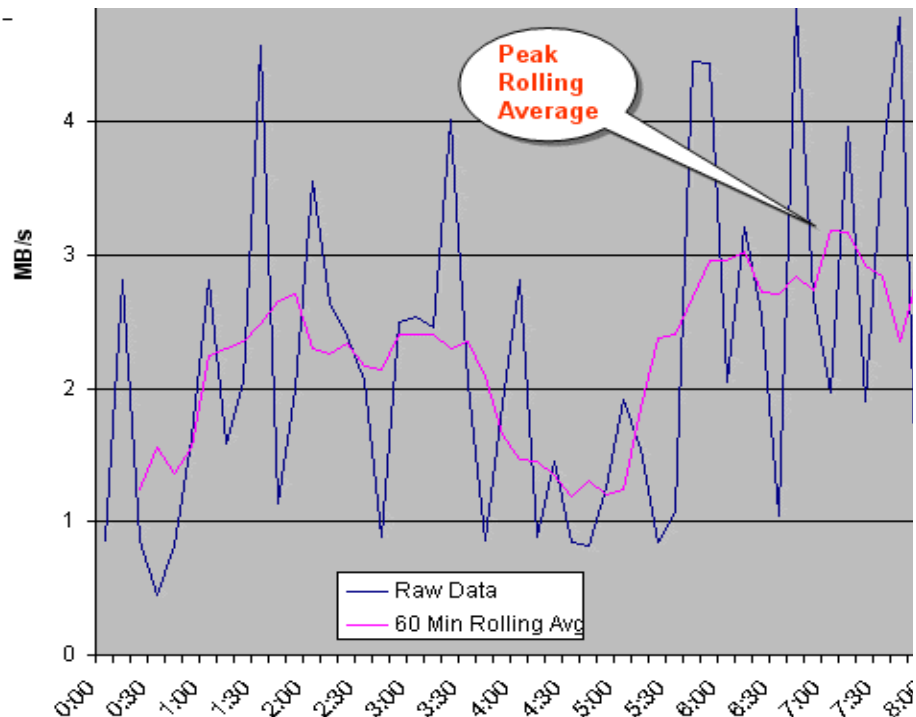


Figure 2-3: 60-Minute Rolling Averages Graphed Over Raw Data

- From the spreadsheet or graph, locate the largest value in the E column. This is your Peak Rolling Average (PRA) value. Use the PRA to calculate the *cumulative peak data change* over cycle time. The following formula calculates the largest expected data change over the cycle time. This will ensure that you do not overflow your data pool.

(PRA in MB/sec) x (cycle time seconds) = (Cumulative peak data change)

For example, if the PRA is 3 MB/sec, and the cycle time is 3600 seconds (1 hour), then:

$$3\text{MB/sec} \times 3600 \text{ seconds} = 10,800 \text{ MB}$$

This shows the maximum amount of changed data (pool data) that you can expect in a 60 minute time period. This is the base data pool size required for TCE.

- Hitachi recommends a 20-percent safety factor for data pools. Calculate a safety factor with the following formula:

(Combined base data pool size) x 1.2. For example:

$$529,200 \text{ MB} \times 1.2 = 635,040 \text{ MB}$$

- It is also recommended that annual increases in data transactions be factored into data pool sizing. This is done to minimize reconfiguration in the future. Do this by multiplying the pool size with safety factor by the percentage of expected annual growth. For example:

$$635,040 \text{ MB} \times 1.2 \text{ (20 percent growth rate for per year)} \\ = 762,048 \text{ MB}$$

Repeat this step for each year the solution will be in place.

- Convert to gigabytes, dividing by 1,000. For example:

762,048 MB / 1,000 = 762 GB

This is the size of the example data pool with safety and growth (2nd year) factored in.

Data pool key points

- Data pools must be set up on the local array and the remote array.
- The data pool must be on the same controller as the P-VOL and V-VOL(s).
- Up to 64 LUs can be assigned to a data pool.
- Plan for highest workload and multi-year growth.
- For set up information, see [Setting up data pools on page 6-6](#).

Determining bandwidth

The purpose of this section is to ensure that you have sufficient bandwidth between the local and remote arrays to copy all your write data in the time-frame you prescribe. The goal is to size the network so that it is capable of transferring estimated future write workloads.

TCE requires two remote paths, each with a minimum bandwidth of 1.5 Mbs.

To determine the bandwidth

1. Graph the data in column "C" in the [Write-Workload Spreadsheet on page 2-4](#).
2. Locate the highest peak. Based on your write-workload measurements, this is the greatest amount of data that will need to be transferred to the remote array. Bandwidth must accommodate maximum possible workload to insure that the system does not become subject to its capacity being exceeded. This would cause further problems, such as the new write data backing up in the data pool, update cycles becoming extended, and so on.
3. Though the highest peak in your workload data should be used for determining bandwidth, you should also take notice 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 high-workload processes. Changing the timing of a process may lower workload.
4. Although bandwidth can be increased, Hitachi recommends that projected growth rate be factored over a 1, 2, or 3 year period.

[Table 2-1](#) shows TCE bandwidth requirements.

Table 2-1: Bandwidth Requirements

| Average Inflow | Bandwidth Requirements | WAN Types |
|-----------------------|-------------------------------|------------------|
| .08 - .149 MB/s | 1.5 Mb/s or more | T1 |
| .15 - .299 MB/s | 3 Mb/s or more | T1 x two lines |
| .3 - .599 MB/s | 6 Mb/s or more | T2 |
| .6 - 1.199 MB/s | 12 Mb/s or more | T2 x two lines |
| 1.2 - 4.499 MB/s | 45 Mb/s or more | T3 |
| 4.500 - 9.999 MB/s | 100 Mb/s or more | Fast Ethernet |

Plan and design — remote path

A remote path is required for transferring data from the local array to the remote array. This chapter provides network and bandwidth requirements, and supported remote path configurations.

- ❑ [Remote path requirements](#)
- ❑ [Remote path configurations](#)
- ❑ [Using the remote path — best practices](#)

Remote path requirements

The remote path is the connection used to transfer data between the local array and remote array. TCE supports Fibre Channel and iSCSI port connectors and connections. The connections you use must be either one or the other: they cannot be mixed.

The following kinds of networks are used with TCE:

- 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 3-1 shows the basic TCE configuration with a LAN and WAN.

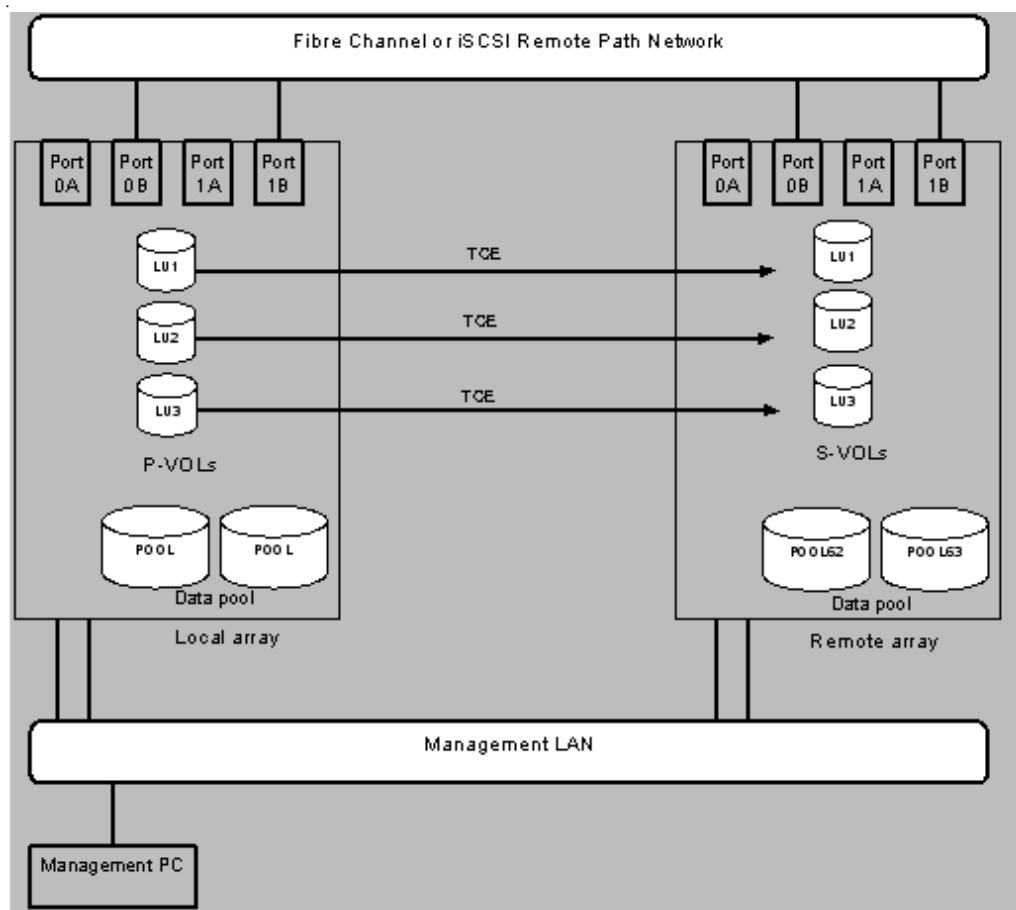


Figure 3-1: Remote Path Configuration

Requirements are provided in the following:

- [Management LAN requirements on page 3-3](#)
- [Remote data path requirements on page 3-3](#)
- [WAN optimization controller \(WOC\) requirements on page 3-4](#)
- [Fibre channel extender connection on page 3-9.](#)

Management LAN requirements

Fast Ethernet is required for an IP LAN.

Remote data path requirements

This section discusses the TCE 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 2-3](#)
- [Determining bandwidth on page 2-7](#)

[Table 3-1](#) provides remote path requirements for TCE. A WOC may also be required, depending on the distance between the local and remote sites and other factors listed in [Table 3-3](#).

Table 3-1: Remote Data Path Requirements

| Item | Requirements |
|---------------------|--|
| Bandwidth | <ul style="list-style-type: none">• 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.• See Table 2-1 on page 2-8 for bandwidth requirements. |
| Remote Path Sharing | <ul style="list-style-type: none">• The remote path must be dedicated for TCE pairs.• When two or more pairs share the same path, a WOC is recommended for each pair. |

[Table 3-2](#) shows types of WAN cabling and protocols supported by TCE and those not supported.

Table 3-2: Supported, Not Supported WAN Types

| | WAN Types |
|---------------|---|
| Supported | <ul style="list-style-type: none">• Dedicated Line (T1, T2, T3 etc) |
| Not-supported | <ul style="list-style-type: none">• 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. TCE copy performance over longer distances is significantly increased when WOC is used. A WOC guarantees bandwidth for each line.

- Use [Table 3-3](#) to determine whether your TCE system requires the addition of a WOC.
- [Table 3-4](#) shows the requirements for WOCs.

Table 3-3: Conditions Requiring a WOC

| Item | Condition |
|-------------------|---|
| Latency, Distance | <ul style="list-style-type: none">• 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 | <ul style="list-style-type: none">• If two or more pairs share the same WAN, A WOC is recommended for each pair. |

Table 3-4: WOC Requirements

| Item | Requirements |
|---------------|---|
| LAN Interface | <ul style="list-style-type: none">• Gigabit Ethernet or fast Ethernet must be supported. |
| Performance | <ul style="list-style-type: none">• Data transfer capability must be equal to or more than bandwidth of WAN. |
| Functions | <ul style="list-style-type: none">• 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

TCE supports both Fibre Channel and iSCSI connections for the remote path.

- Two remote paths must be set up between, one per controller. This ensures that an alternate path is available in the event of link failure during copy operations.
- Paths can be configured 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. Hitachi recommends making connections between the same controller/port, such as port 0B to 0B, and 1 B to 1 B, for simplicity. Ports can be used for both host I/O and replication data.

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

Fibre channel

The Fibre Channel remote data path can be set up in the following configurations:

- Direct connection
- Single Fibre Channel switch and network connection
- Double FC switch and network connection
- Wavelength Division Multiplexing (WDM) and dark fibre extender

The array supports direct or switch connection only. Hub connections are not supported.

General recommendations

The following is recommended for all supported configurations:

- TCE requires one path between the host and local array. However, two paths are recommended; the second path can be used in the event of a path failure.

Direct connection

Figure 3-2 illustrates two remote paths directly connecting the local and remote arrays. This configuration can be used when distance is very short, as when creating the initial copy or performing data recovery while both arrays are installed at the local site.

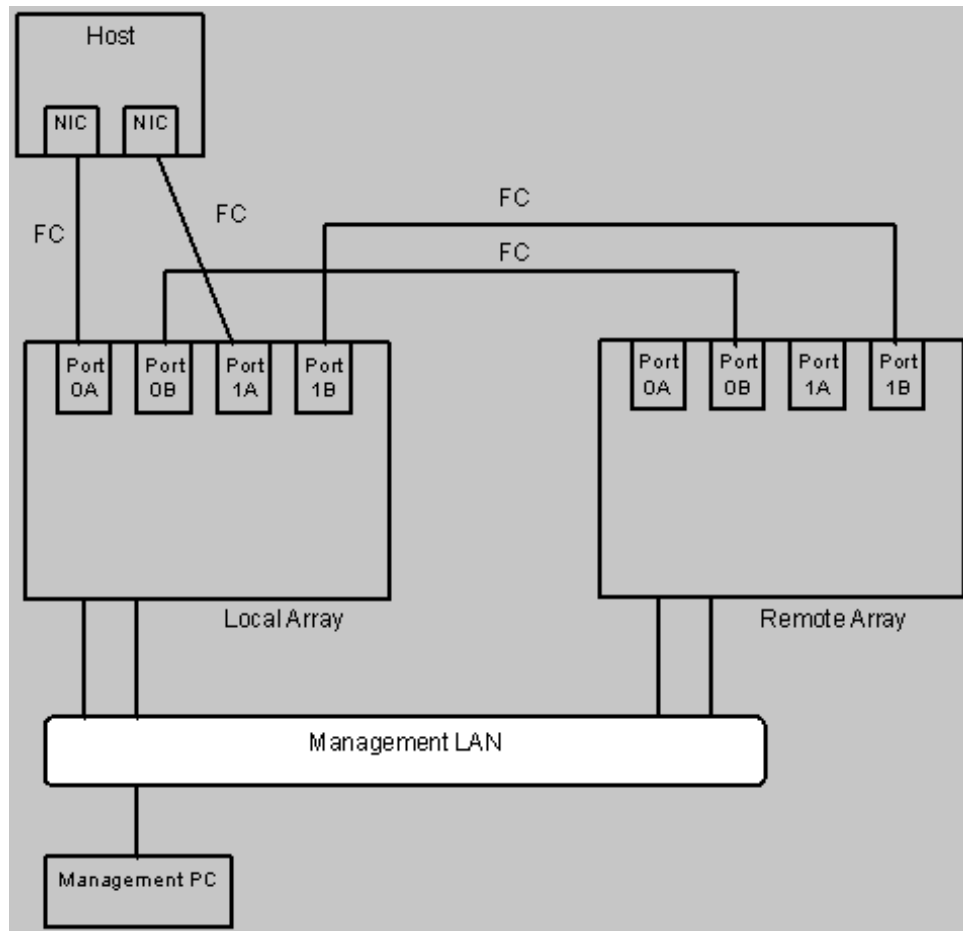


Figure 3-2: Direct FC Connection

Single FC switch, network connection

Switch connections increase throughput between the arrays. [Figure 3-3](#) illustrates two remote paths routed through one FC switch and one FC network to make the connection to the remote site.

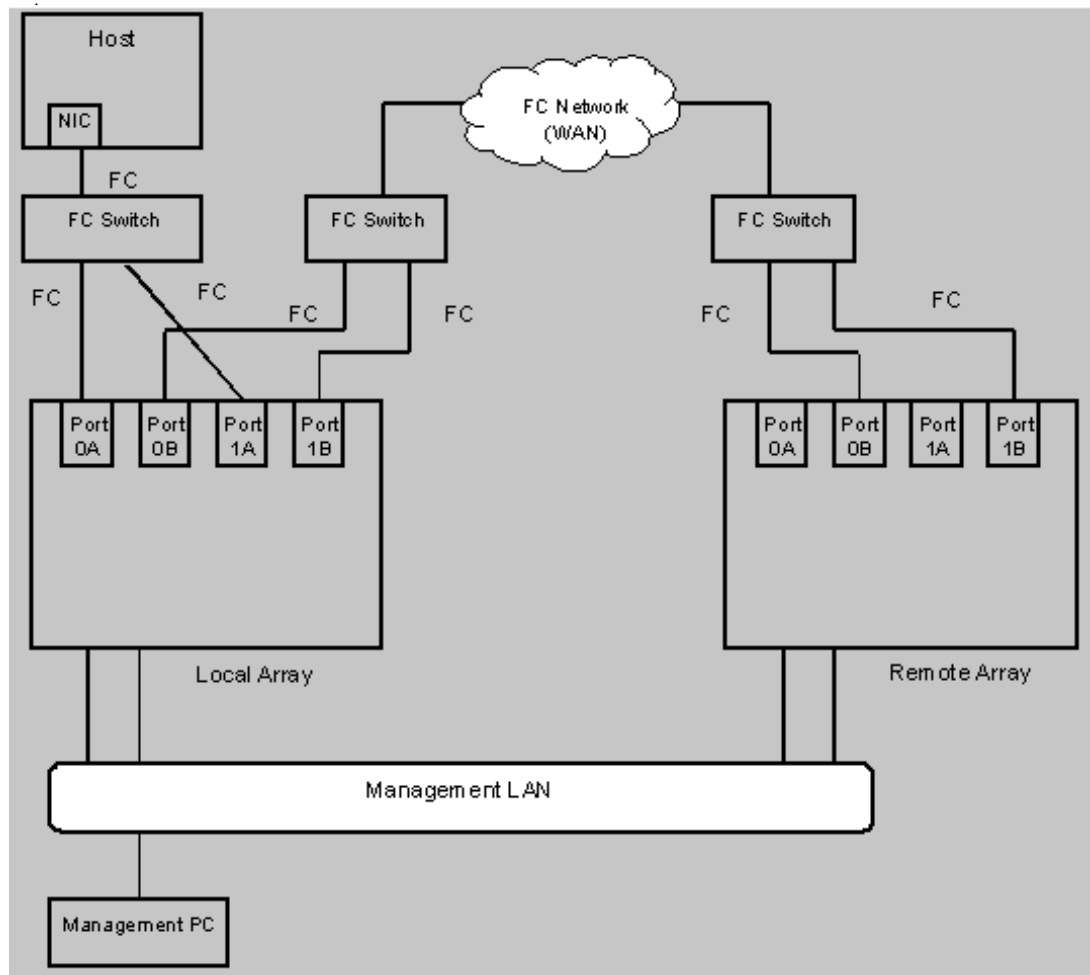


Figure 3-3: Single FC Switch, Network Connection

Recommendations

- While this configuration may be used, it is not recommended since failure in an FC switch or the network would halt copy operations.
- Separate switches should be set up for host I/O to the local array and for data transfer between arrays. Using one switch for both functions results in deteriorated performance.

Double FC switch, network connection

Figure 3-4 illustrates two remote paths using two FC switches and two FC networks to make the connection to the remote site.

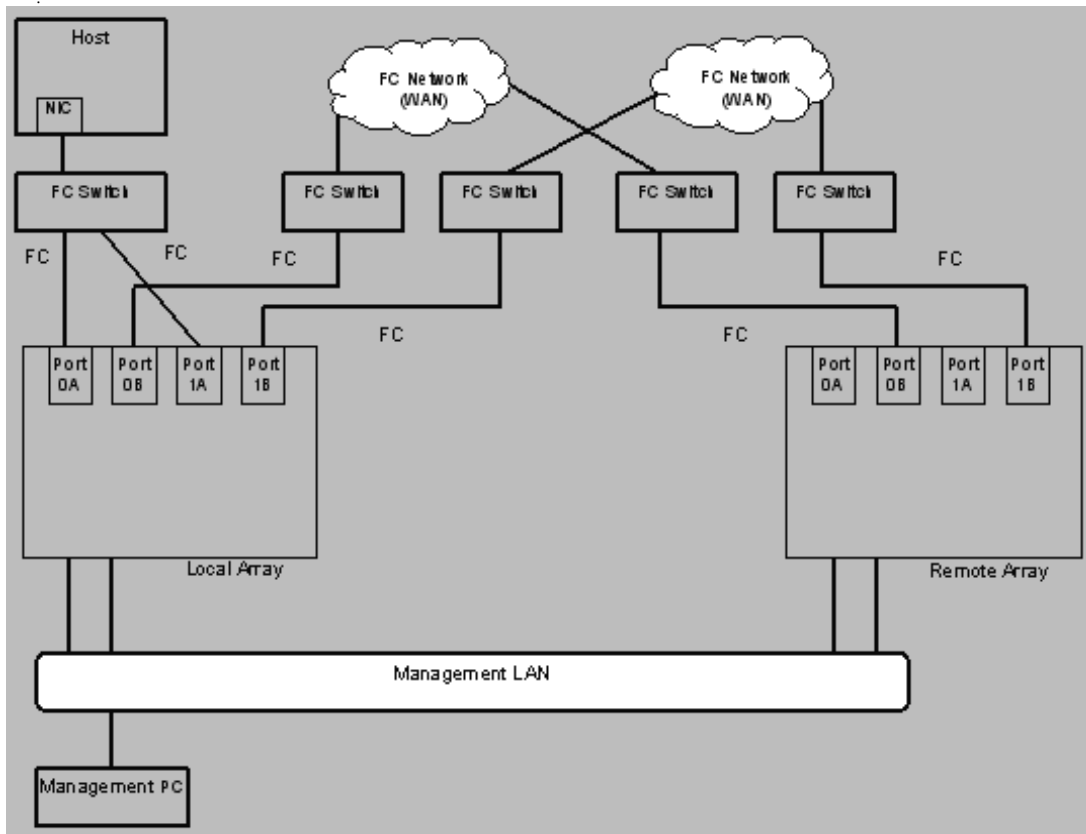


Figure 3-4: Double FC Switches, Networks Connection

Recommendations

- Separate switches should be set up for host I/O to the local array and for data transfer between arrays. Using one switch for both functions results in deteriorated performance.

Fibre channel extender connection

Channel extenders convert Fibre Channel to FCIP or iFCP, which allows you to use IP networks and significantly improve performance over longer distances.

Figure 3-5 illustrates two remote paths using two FC switches, Wavelength Division Multiplexor (WDM) extender, and dark fibre to make the connection to the remote site.

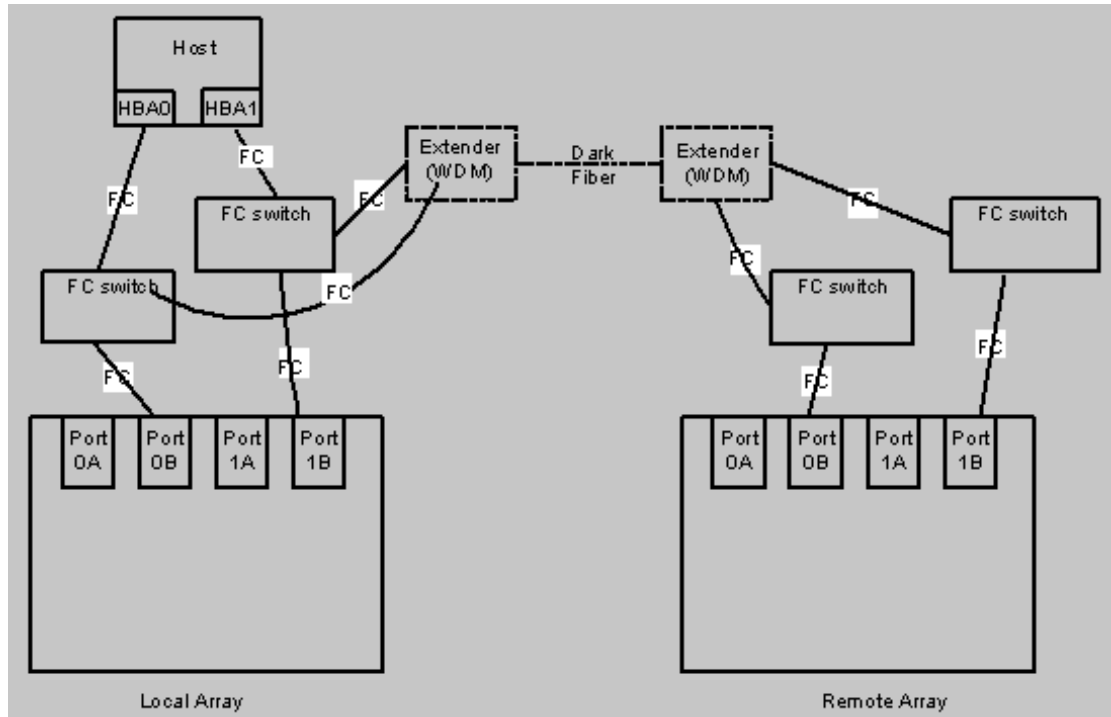


Figure 3-5: Fibre Channel Switches, WDM, Dark Fibre Connection

Recommendations

- Only qualified components are supported.

For more information on WDM, see [Appendix E, Wavelength Division Multiplexing \(WDM\) and dark fibre](#).

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 the Fibre Channel cable. Each port on the array must be set separately.

Table 3-5: Setting Port Transfer Rates

| | If the host port is set to | Set the remote array port to |
|--------------------|----------------------------|------------------------------|
| Manual mode | 1 Gbps | 1 Gbps |
| | 2 Gbps | 2 Gbps |
| | 4 Gbps | 4 Gbps |
| | 8 Gbps | 8 Gbps |
| Auto mode | 2 Gbps | Auto, with max of 2 Gbps |
| | 4 Gbps | Auto, with max of 4 Gbps |
| | 8 Gbps | Auto, with max of 8 Gbps |

Maximum speed is ensured using the manual settings.

You can specify the port transfer rate using the Navigator 2 GUI, on the **Edit FC Port** screen (Settings/FC Settings/port/Edit Port button).



NOTE: If your remote path is a direct connection, make sure that the array power is off when modifying the transfer rate to prevent remote path blockage.

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

iSCSI

The iSCSI remote data 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

Recommendations

The following is recommended for all supported configurations:

- Two paths should be configured from the host to the array. This provides a backup path in the event of path failure.

Direct connection

Figure 3-6, 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, category 5e or 6 copper LAN cable is recommended.

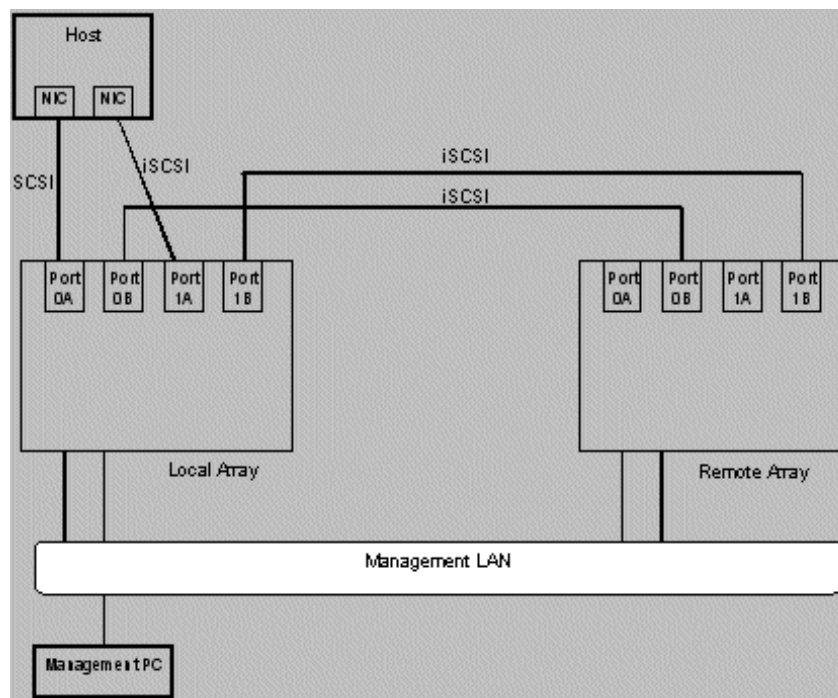


Figure 3-6: Direct iSCSI Connection

Recommendations

- 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 5e or 6 copper LAN cable is recommended.

Single LAN switch, WAN connection

Figure 3-7, illustrates two remote paths using one LAN switch and network to the remote array.

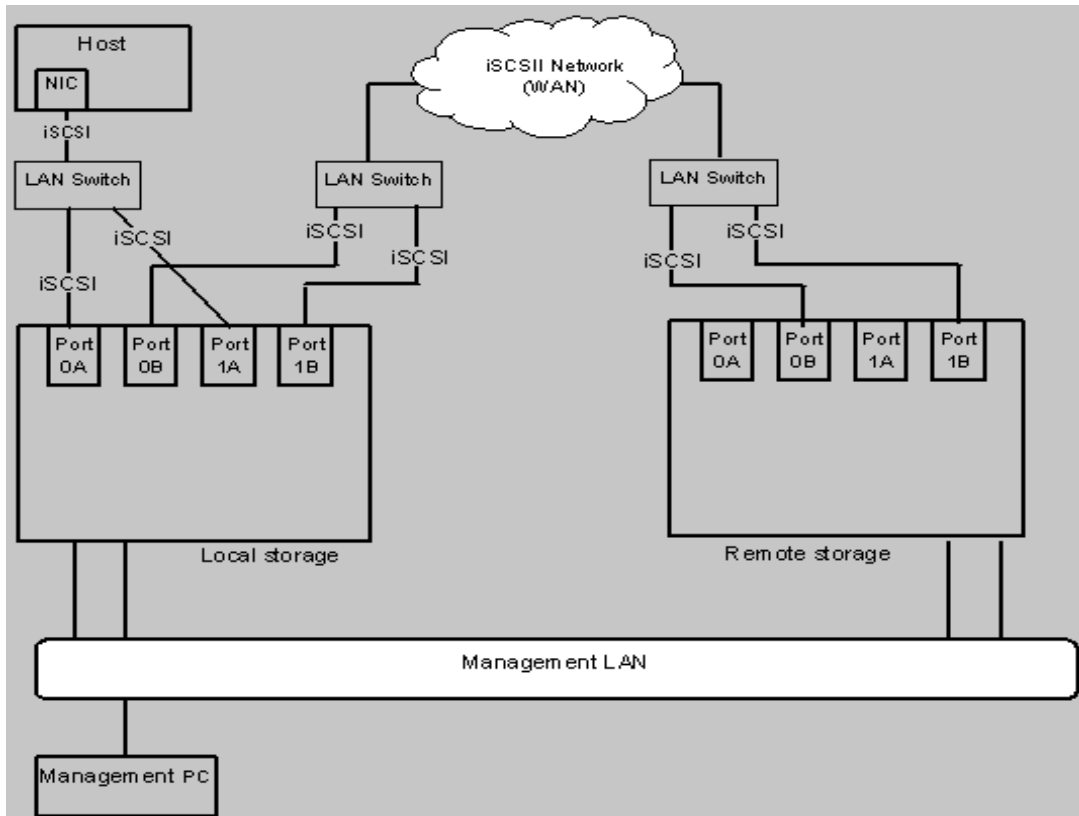


Figure 3-7: Single-Switch Connection

Recommendations

- 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 3-8, illustrates two remote paths using multiple LAN switches and WANs to make the connection to the remote site.

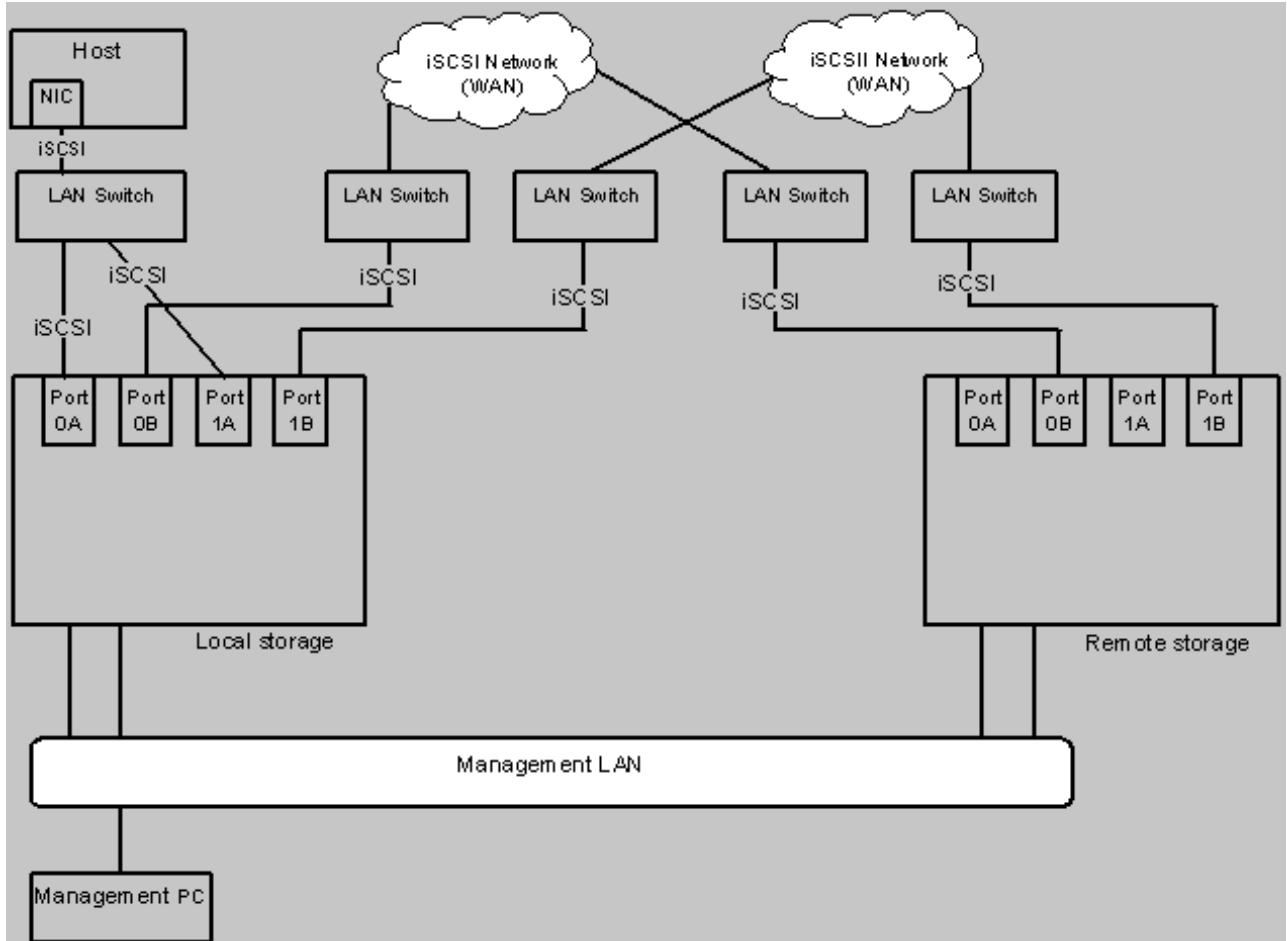


Figure 3-8: Multiple-Switch and WAN Connection

Recommendations

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

Single LAN switch, WOC, WAN connection

WOCs may be required for TCE, depending on your system's bandwidth, latency, and so on. Use of a WOC improves performance. See [WAN optimization controller \(WOC\) requirements on page 3-4](#) for more information.

Figure 3-9, illustrates two remote paths using a single LAN switch, WOC, and WAN to make the connection to the remote site.

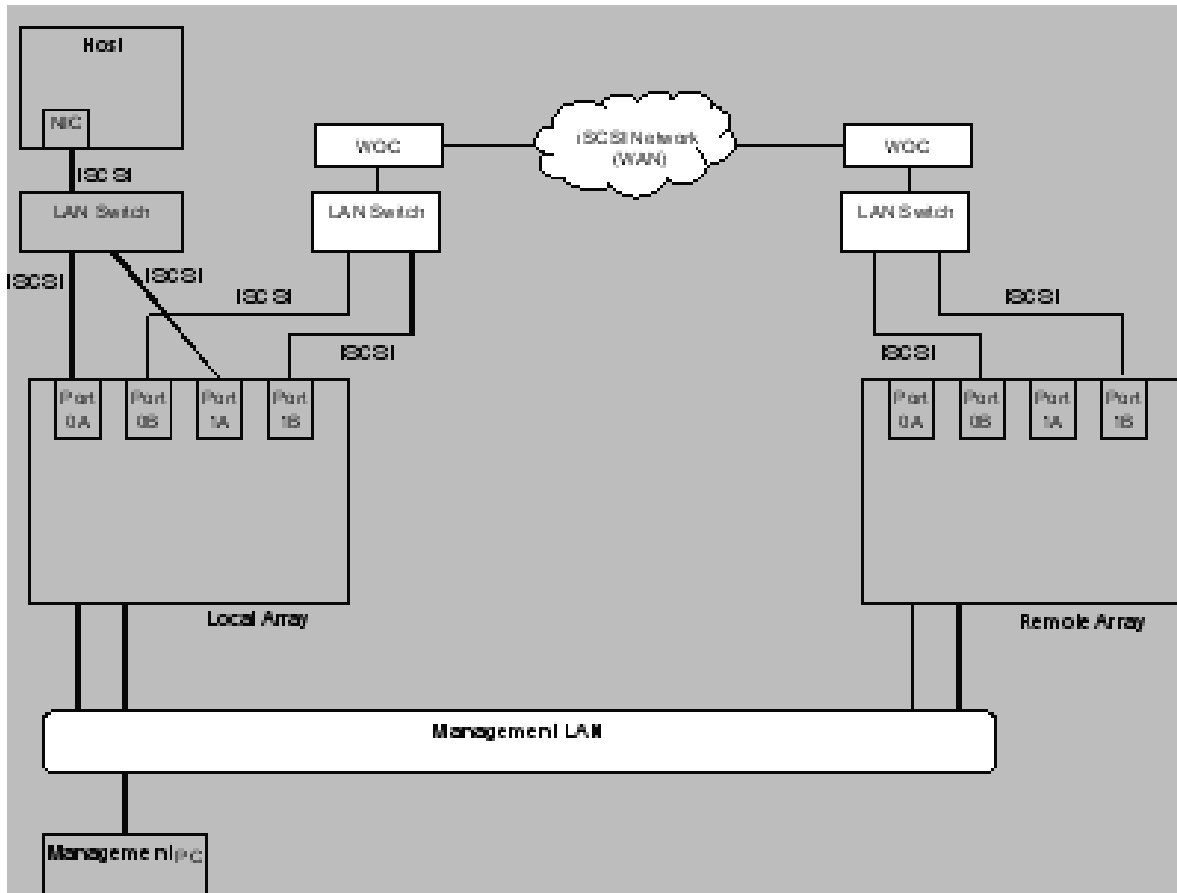


Figure 3-9: Single Switch, WOC, and WAN Connection

Multiple LAN switch, WOC, WAN connection

Figure 3-10, illustrates two remote connections using multiple LAN switches, WOCs, and WANs to make the connection to the remote site.

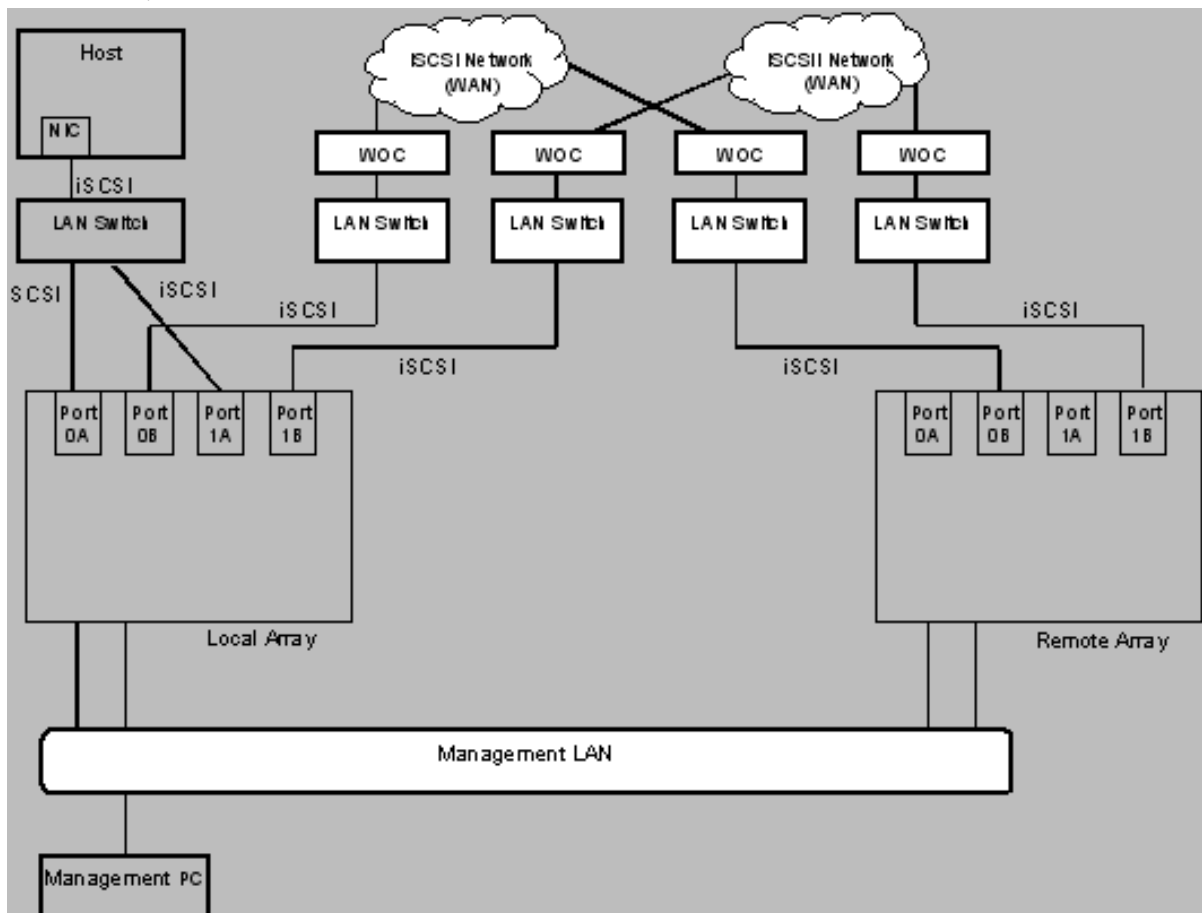


Figure 3-10: Connection Using Multiple Switch, WOC, WAN

Recommendations

- 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.
- 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 3-11, shows two local arrays connected to two remote arrays, each via a LAN switch and WOC.

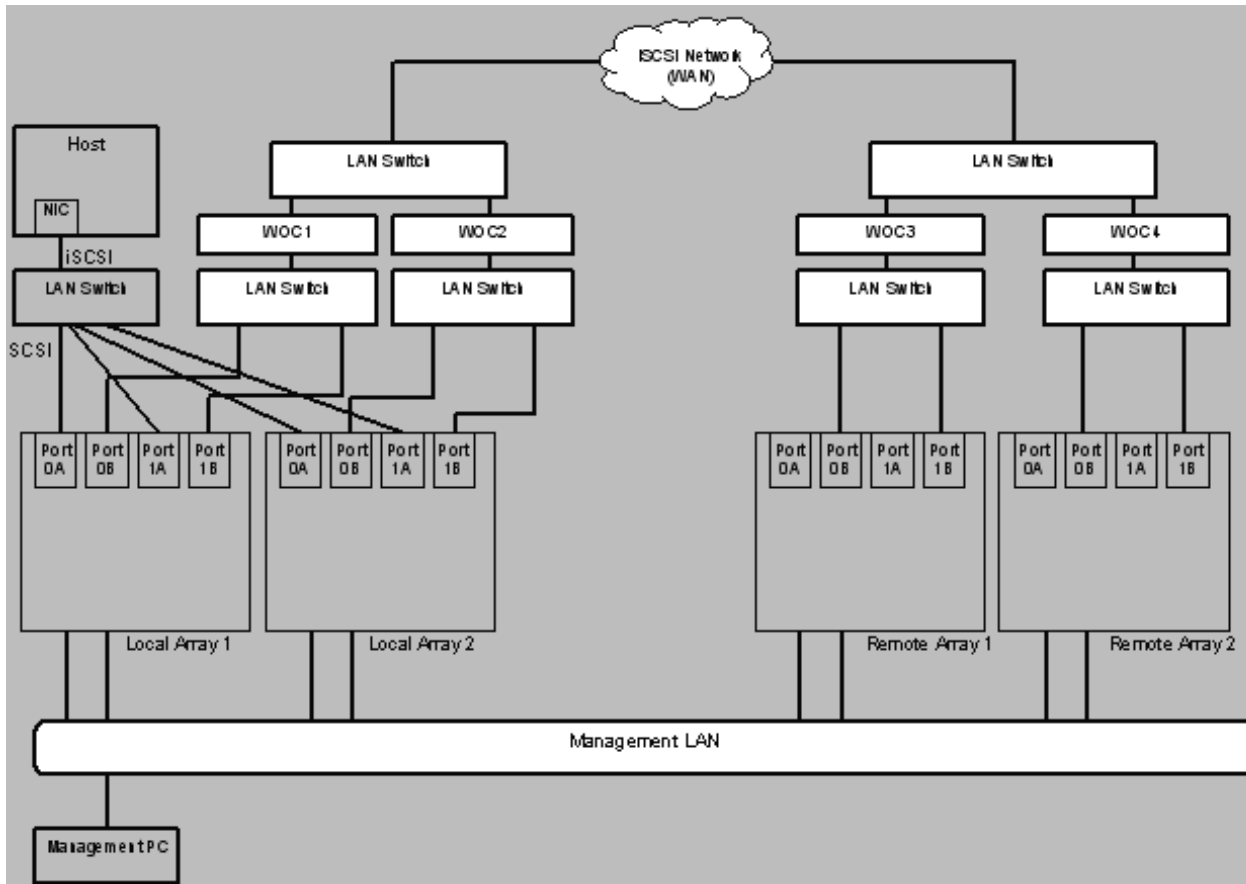


Figure 3-11: Multiple Array Connection Using Single WAN

Recommendations

- 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 the 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 3-12, shows two local arrays connected to two remote arrays, each via two LAN switches, WANs, and WOCs.

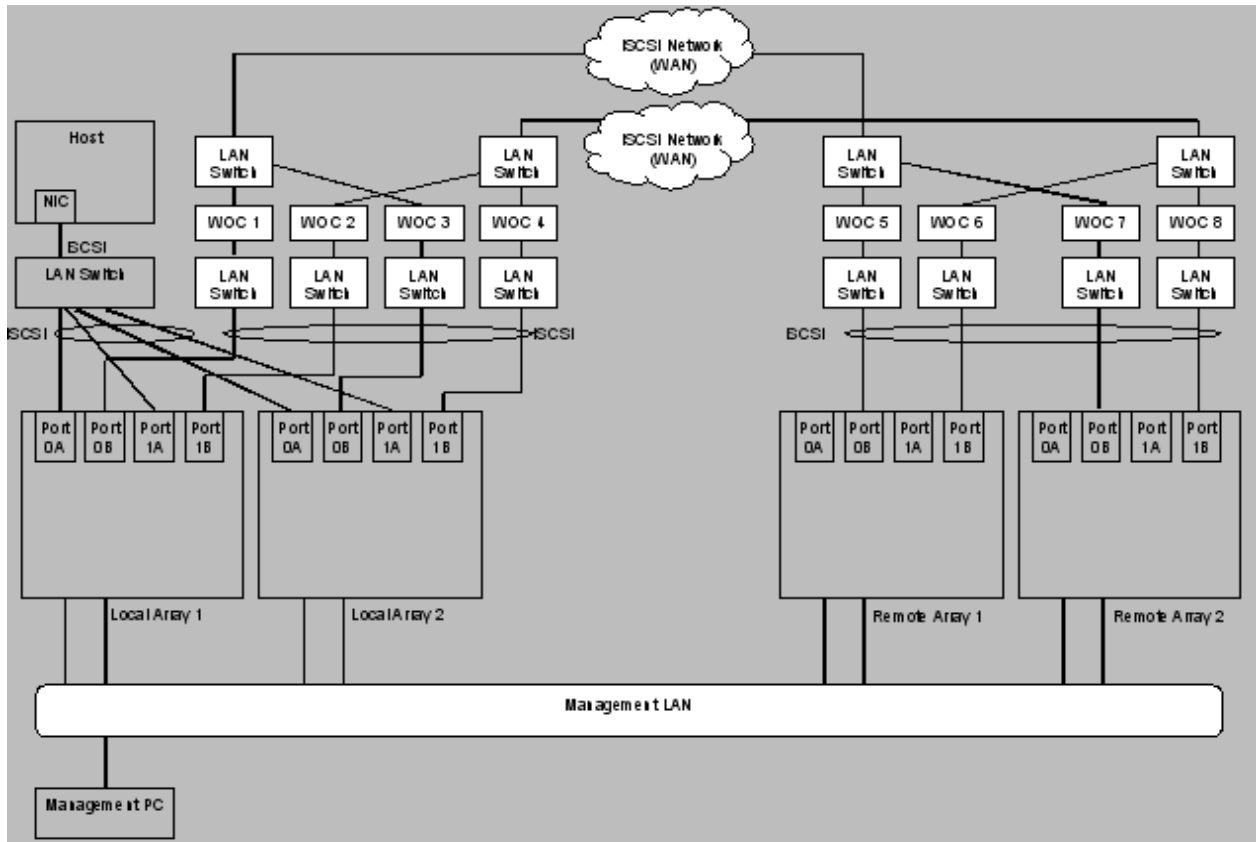


Figure 3-12: Multiple Array Connection Using Two WANs

Recommendations

- 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

Plan and design—arrays, volumes, and operating systems

This chapter provides the information you need to prepare your arrays and volumes for TCE operations

- ❑ [Planning arrays—moving data from earlier AMS models](#)
- ❑ [Planning logical units for TCE volumes](#)
- ❑ [Operating system recommendations and restrictions](#)
- ❑ [Maximum supported capacity](#)

Planning workflow

Planning a TCE system consists of determining business requirements for recovering data, measuring production write-workload and sizing data pools and bandwidth, designing the remote path, and planning your arrays and volumes. This chapter discusses arrays and volumes as follows:

- Requirements and recommendations for using previous versions of AMS with the AMS 2000 Family.
- Logical unit set up: LUs must be set up on the arrays before TCE is implemented. Volume requirements and specifications are provided.
- Operating system considerations: Operating systems have specific restrictions for replication volumes pairs. These restrictions plus recommendations are provided.
- Maximum Capacity Calculations: Required to make certain that your array has enough capacity to support TCE. Instructions are provided for calculating your volumes' maximum capacity.

Planning arrays—moving data from earlier AMS models

Logical units on AMS 2100, 2300, and 2500 systems can be paired with logical units on AMS 500 and AMS 1000 systems. Any combination of these arrays may be used on the local and remote sides.

TCE pairs with WMS 100 and AMS 200 are not supported with AMS2100, 2300, or 2500.

When using the earlier model arrays, please observe the following:

- The bandwidth of the remote path to AMS 500 or AMS 1000 must be 20 Mbps or more.
- 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 of AMS 500 or AMS 1000 must be 0780/A or later when pairing with an AMS 2100, 2300, or 2500 where the hardware Rev is 0100.
- The firmware version of AMS 500 or AMS 1000 must be 0786/A or later when pairing with an AMS 2010, 2100, 2300, or 2500 where the hardware Rev is 0200.
- Pair operations for AMS 500 and AMS 1000 cannot be performed using the Navigator 2 GUI.
- AMS500 and AMS1000 cannot use functions that are newly supported by AMS2100 or AMS2300.
- Because AMS 500 or AMS 1000 can have only one data pool per controller, you are not able to specify which data pool to use. Because of this, the data pool that is used is determined as follows:
 - When AMS 500 or AMS 1000 is the local array, data pool 0 is used if the S-VOL LUN is even; data pool 1 is used if the S-VOL LUN is odd.

- When an AMS 2100, 2300, or 2500 is the local array, the data pool number is ignored even if specified. Data pool 0 is used if the S-VOL owner-controller is 0, and data pool 1 is selected if the S-VOL owner-controller is 1.
- The AMS 500 or AMS 1000 cannot use the functions that are newly supported by AMS 2010, 2100, 2300, or 2500.

Planning logical units for TCE volumes

Please review the recommendations in the following sections before setting up TrueCopy volumes. Also, review [Requirements and specifications on page 5-1](#).

Volume pair, data pool recommendations

- The P-VOL and S-VOL must be identical in size, with matching block count. To check block size, in the Navigator 2 GUI, navigate to Groups/RAID Groups/Logical Units tab. Click the desired LUN. On the popup window that appears, review the Capacity field. This shows block size.
- 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 and secondary volumes and data pools to a RAID group consisting of SAS or SSD drives to achieve best possible performance. SATA drives can be used, however.
- Assign an LU consisting of four or more data disks, otherwise host and copying performance may be lowered.
- Limit the I/O load on both local and remote arrays to maximize performance. Performance on each array also affects performance on the other array, as well as data pool capacity and the synchronization of volumes.

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 x 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 a different host.

HP server

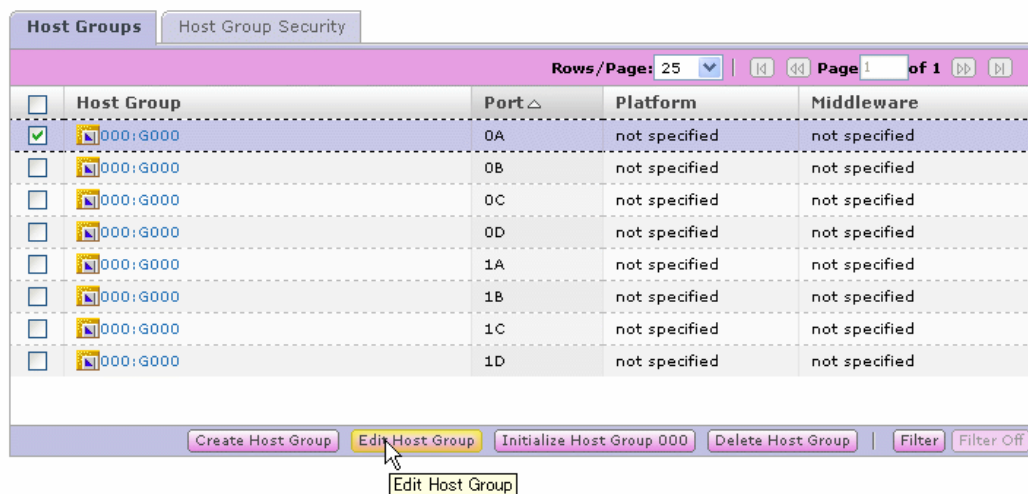
When MC/Service Guard is used on a HP server, connect the host group (Fibre Channel) or the iSCSI Target to HP server as follows:

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**.

Host Groups

DF800M_85000026 > Groups > Host Groups



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**, **Enable PSUE Read Reject 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 Server 2000

- A P-VOL and S-VOL *cannot* be made into a dynamic disk on Windows Server 2000 and Windows Server™ 2008.
- Native OS mount/dismount commands can be used for all platforms, except Windows Server 2000. 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 CCI to perform volume mount/unmount operations. For more information on the CCI mount/unmount commands, see the *Hitachi AMS Command Control Interface (CCI) Reference Guide*.

Windows Server 2003/2008

- A P-VOL and S-VOL *can* be made into a dynamic disk on Windows Server 2003.
- 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 01-13-03/00 and later.
- In Windows Server™ 2008, refer to the Hitachi Adaptable Modular Storage Command Control Interface (CCI) Reference Guide for the restrictions when the mount/unmount command is used.
- Windows® may write for the un-mounted volume. If a pair is resynchronized while retaining data on the S-VOL on the server memory, the compatible backup cannot be collected. Therefore, execute the CCI sync command immediately before re-synchronizing the pair for the un-mounted S-VOL.
- In Windows Server™ 2008, set only the P-VOL of TCE to be recognized by the host and let another host recognize the S-VOL.
- (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 2003 Server, LUs are identified by HLUN. To map LUN to HLUN on Windows, proceed as follows. These instructions provide procedures for iSCSI and Fibre Channel interfaces.

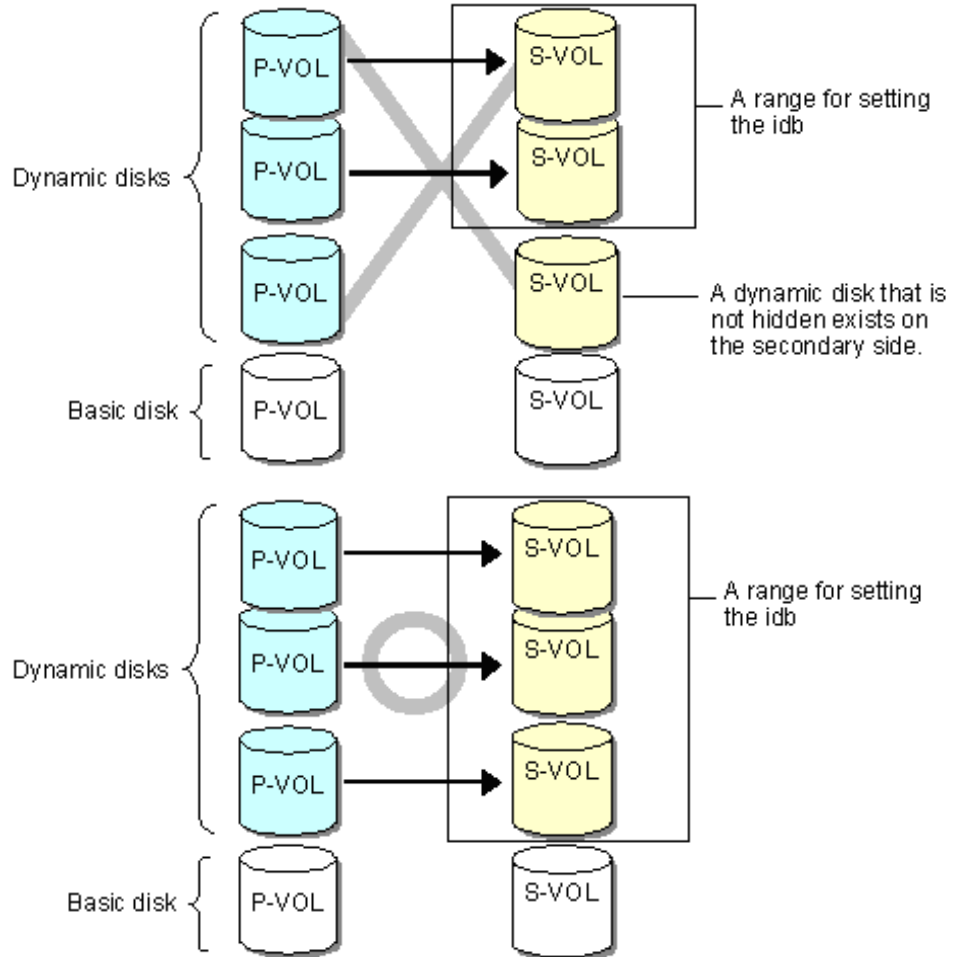
1. Identify the HLUN of your Windows disk.
 - a. From the Windows Server 2003 Control Panel, select Computer Management>Disk Administrator.
 - b. Right-click the disk whose HLUN you want to know, then select **Properties**. The number displayed to the right of "LUN" in the dialog window is the HLUN.
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 Target** 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 2d.
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 3d.

Dynamic Disk with Windows 2003

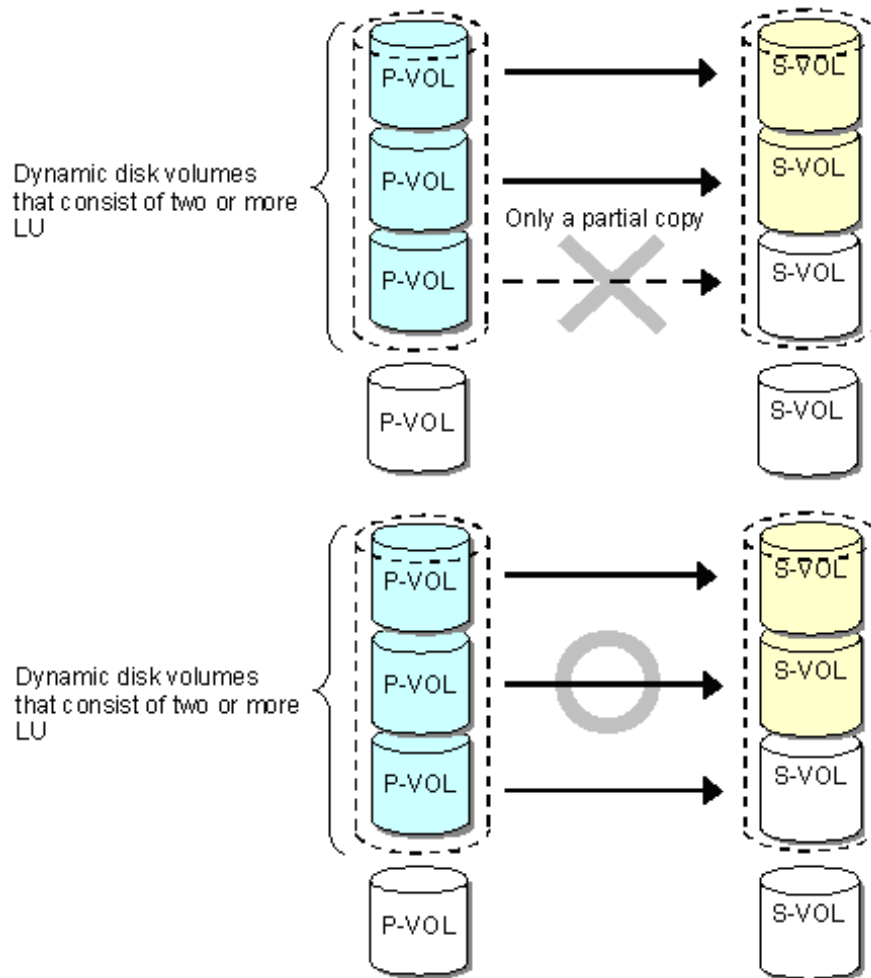
- A P-VOL and an S-VOL can be made into a dynamic disk on Windows Server 2003.
- When using an S-VOL with a secondary host, insure that the pair status is Split.
- A host cannot recognize both a P-VOL and its S-VOL at the same time. Map the P-VOL and S-VOL to separate hosts.
- An LU in which two or more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

When copying, hide all the dynamic disks that exist on the primary side using the `raidvchksset -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.



- 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.



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

Maximum supported capacity

The capacity you can assign to replication volumes per controller is limited, for the following reasons:

- The TCE P-VOL and S-VOL, and the SnapShot P-VOL if used, share common data pool resources. Therefore, data pool capacity is limited.
- The maximum capacity supported by a TCE pair depends on the P-VOL capacity of SnapShot (if used), data pool capacity, and cache memory capacity.
- When using other copy systems and TCE together, the maximum supported capacity of the P-VOL may be restricted further.

In addition to this, capacity is managed by the AMS array in blocks of 15.75 KB for data volumes and 3.2 KB for data pools. For example, when a P-VOL block's actual size is 16 KB, the array manages it as two blocks of 15.75 KB, or 31.5 KB. Data pool capacity is managed in the same way but at 3.2 KB per block.

This section provides formulas for calculating your existing or planned TCE volume capacity and comparing it to the maximum supported capacity for your particular controller and its cache memory size.

TCE capacity must be calculated for both of the following:

1. The ratio of TCE and SnapShot (if used) capacity to data pool capacity. Capacity is calculated using the following volumes:
 - TCE P-VOLs and S-VOLs
 - SnapShot P-VOLs (if used)
 - All data pools
2. Concurrent use of TCE and ShadowImage. If SnapShot is used concurrently also, it is included in this calculation. Capacity is calculated using the following volumes:
 - TCE P-VOLs
 - SnapShot P-VOLs
 - ShadowImage S-VOLs



NOTE: When SnapShot is enabled, a portion of cache memory is assigned to it for internal operations. Hitachi recommends that you review the appendix on SnapShot and Cache Partition Manager in the *Hitachi AMS 2000 Family Copy-on-Write SnapShot User Guide*.

TCE and SnapShot capacity

Because capacity is managed by the array in blocks of 15.75 KB for of data volumes and 3.2 KB for data pools, the capacity of your array's TCE and SnapShot volumes must be specially calculated.

All formulas, tables, graphs and examples pertain to one controller. On dual controller arrays, you must perform calculations for both controllers.

Managed capacity is calculated here, per controller, using the following formula:

$$\begin{aligned} & \text{Size of all TCE P-VOLs} + \\ & \text{Size of all TCE S-VOLs} + \\ & \text{Size all SnapShot P-VOLs (if used)} \\ & / 5 \\ & + \text{size of all data pool volumes} \\ & < \text{Max. Sup. Capacity} \end{aligned}$$

Maximum supported capacity is shown in [Table 4-1](#).

Table 4-1: Maximum Supported Capacities, per Controller Cache Size

| Cache Memory per Controller | Maximum Capacity (TCE P-VOLs and S-VOLs, SnapShot P-VOLs, Data Pools) | | |
|-----------------------------|--|---------------|---------------|
| | AMS2100 | AMS2300 | AMS2500 |
| 2 GB per CTL | 1.4 TB | Not supported | Not supported |
| 4 GB per CTL | 6.2 TB | 6.2 TB | Not supported |
| 8 GB per CTL | Not supported | 12.0 TB | 12.0 TB |
| 16 GB per CTL | Not supported | Not supported | 24.0 TB |



NOTE: In a dual-controller array, the calculations must be performed for both controllers.

Example:

In this example, the array and cache memory per controller is AMS 2300/4 GB.

1. List the size of each TCE P-VOL and S-VOL on the array, and of each SnapShot P-VOL (if present) in the array. For example:

TCE P-VOL 1 = 100 GB
TCE S-VOL 1 = 100 GB
SnapShot P-VOL 1 = 50 GB

2. Calculate managed P-VOL and S-VOL capacity, using the formula:

$$\text{ROUNDUP (P-VOL/S-VOL / 15.75) * 15.75}$$

For example:

TCE P-VOL1: $\text{ROUNDUP (100 / 15.75) = 7}$

$7 * 15.75 = 110.25 \text{ GB}$, the managed P-VOL Capacity

TCE S-VOL1: ROUNDUP (100 / 15.75) = 7

7 * 15.75 = 110.25 GB, the managed S-VOL Capacity

SnapShot P-VOL1: ROUNDUP (50 / 15.75) = 4

4 * 15.75 = 63 GB, the managed P-VOL Capacity

3. Add the total managed capacity of P-VOLs and S-VOLs. For example:

Total TCE P-VOL and S-VOL managed capacity = 221 GB

Total SnapShot P-VOL capacity = 63 GB

221 GB + 63 GB = 284 GB

4. For each P-VOL and S-VOL, list the data pools and their sizes. For example:

TCE P-VOL1 has 1 data pool whose capacity = 70 GB

TCE S-VOL1 has 1 data pool whose capacity = 70 GB

SnapShot P-VOL1 has 1 data pool whose capacity = 30 GB

5. Calculate managed data pool capacity, using the formula:

ROUNDUP (data pool capacity / 3.2) * 3.2

For example:

TCE P-VOL 1 data pool: ROUNDUP (70 / 3.2 = 22) * 3.2 = 71 GB

TCE S-VOL 1 data pool: ROUNDUP (70 / 3.2 = 22) * 3.2 = 71 GB

SS P-VOL 1 data pool: ROUNDUP (30 / 3.2 = 10) * 3.2 = 32 GB

6. Add total data pool managed capacity. For example:

71 GB + 71 GB + 32 GB = 175 GB

7. Calculate total managed capacity using the following equation:

ROUNDUP Total TCE/SnapShot managed capacity / 5 + total data pool managed capacity < maximum supported capacity

For example:

Divide the total TCE/SnapShot capacity by 5.

284 GB / 5 = 57 GB

8. Add the quotient to data pool managed capacity. For example:

57 GB + 176 GB = 234 GB

9. Compare managed capacity to maximum supported capacity for the 4 GB cache per controller for AMS 2300, which is 6.2 TB. The managed capacity is well below maximum supported capacity.

[Table 4-3 on page 4-13](#) through [Table 4-6 on page 4-14](#) show how closely capacity between data volumes and data pool volumes must be managed. These tables are provided for your information. Also, [Figure 4-1 on page 4-15](#) shows a graph of how data volume-to-data pool volume relates to maximum supported capacity.

TCE, SnapShot, ShadowImage concurrent capacity

If ShadowImage is used on the same controller as TCE, capacity for concurrent use must also be calculated and compared to maximum supported capacity. If SnapShot is used also, it is included in concurrent-use calculations.

Concurrent-use capacity is calculated using the following formula:

$$\begin{aligned} & \text{Maximum TCE supported capacity of P-VOL and S-VOL (TB)} \\ & = \text{TCE maximum single capacity} \\ & - (\text{Total ShadowImage S-VOL capacity} / 51) \\ & - (\text{Total SnapShot P-VOL capacity} / 3) \end{aligned}$$

TCE maximum single capacity is shown in [Table 4-2](#).

Table 4-2: TCE Maximum Single Capacity per Controller

| Equipment Type | Mounted Memory Capacity | Single Maximum Capacity (TB) |
|----------------|-------------------------|------------------------------|
| AMS2100 | 1 GB per CTL | - |
| | 2 GB per CTL | 15 |
| | 4 GB per CTL | 18 |
| AMS2300 | 1 GB per CTL | - |
| | 2 GB per CTL | 14 |
| | 4 GB per CTL | 38 |
| | 8 GB per CTL | 77 |
| AMS2500 | 2 GB per CTL | 10 |
| | 4 GB per CTL | 38 |
| | 6 GB per CTL | 54 |
| | 8 GB per CTL | 70 |
| | 10 GB per CTL | 93 |
| | 12 GB per CTL | 116 |
| | 16 GB per CTL | 140 |

Example

In this example, the array and cache memory per controller is AMS2100 and 2 GB per CTL.

$$\begin{aligned} & \text{Maximum TCE supported capacity of P-VOL and S-VOL (TB)} \\ & = \text{TCE maximum single capacity} \\ & - (\text{Total ShadowImage S-VOL capacity} / 51) \\ & - (\text{Total SnapShot P-VOL capacity} / 3) \end{aligned}$$

1. TCE Maximum single capacity = 15 TB
2. Calculate ShadowImage S-VOL managed capacity ($\text{ROUNDUP}(\text{S-VOL} / 15.75) * 15.75$). Then divide by 51. For example:
All ShadowImage S-VOLs: $\text{ROUNDUP}(4000 \text{ GB} / 15.75) = 254 \text{ GB}$
 $254 * 15.75 = 4001 \text{ GB}$, the managed S-VOL Capacity

$$4001 \text{ GB} / 51 = 79$$

3. Subtract the quotient from the TCE maximum single capacity. For example:

$$15 \text{ TB (15,000 GB)} - 79 \text{ GB} = 14,921 \text{ GB}$$

4. Calculate SnapShot P-VOL managed capacity, then divide by 3. For example:

$$\text{All SnapShot P-VOLs: ROUNDUP } 800 \text{ GB} / 15.75 = 51 \text{ GB}$$

$$51 \text{ GB} * 15.75 = 803 \text{ GB}$$

$$803 \text{ GB} / 3 = 268 \text{ GB}$$

5. Subtract the quotient from the remaining TCE maximum single capacity. For example:

$$14,921 \text{ GB} - 268 \text{ GB} = 14,653 \text{ GB, the capacity left for TCE P-VOLs and S-VOLs on the controller.}$$

If your system's managed capacity exceeds the maximum supported capacity, you can do one or more of the following:

- Change the P-VOL size
- Reduce the number of P-VOLs
- Change the data pool size
- Reduce ShapShot and ShadowImage P-VOL/S-VOL size

Cache limitations on data and data pool volumes

This section provides comparisons in capacity between the data volumes and the data pool volumes under the limitations of the AMS controllers' cache memory. The values in the tables and graph in this section are calculated from the formulas and maximum supported capacity in [TCE and SnapShot capacity on page 4-10](#).



NOTE: "Data volumes" in this section consist of TCE P-VOLs and S-VOLs and SnapShot P-VOLs (if used).

Table 4-3: P-VOL to Data Pool Capacity Ratio on AMS 2100 when Cache Memory is 2 GB/Controller

| Ratio of All P-VOL Capacity to All Data Pool Capacity | All P-VOL Capacity to All Data Pool Capacity (TB) |
|---|---|
| 1:0.1 | 4.6 : 0.4 |
| 1:0.3 | 2.8 : 0.8 |
| 1:0.5 | 2.0 : 1.0 |

Table 4-4: P-VOL to Data Pool Capacity Ratio on AMS 2300/2100 when Cache Memory is 4 GB per CTL

| Ratio of All P-VOL Capacity to All Data Pool Capacity | All P-VOL Capacity to All Data Pool Capacity (TB) |
|---|---|
| AMS 2100/2300/2500 | AMS 2100/2300 |
| 1:0.5 | 20.6 : 2.0 |
| 1:1 | 12.4 : 3.7 |
| 1:3 | 8.8 : 4.4 |

Table 4-5: P-VOL to Data Pool Capacity Ratio on AMS 2500/2300 when Cache Memory is 8 GB per CTL

| Ratio of All P-VOL Capacity to All Data Pool Capacity | All P-VOL Capacity to All Data Pool Capacity (TB) |
|---|---|
| 1:0.1 | 40.0 : 4.0 |
| 1:0.3 | 24.0 : 7.2 |
| 1:0.5 | 17.1 : 8.5 |

Table 4-6: P-VOL to Data Pool Capacity Ratio on AMS 2500 when Cache Memory is 16 GB per CTL

| Ratio of All P-VOL Capacity to All Data Pool Capacity | All P-VOL Capacity to All Data Pool Capacity (TB) |
|---|---|
| 1:0.1 | 80.0 : 8.0 |
| 1:0.3 | 48.0 : 14.4 |
| 1:0.5 | 34.2 : 17.1 |

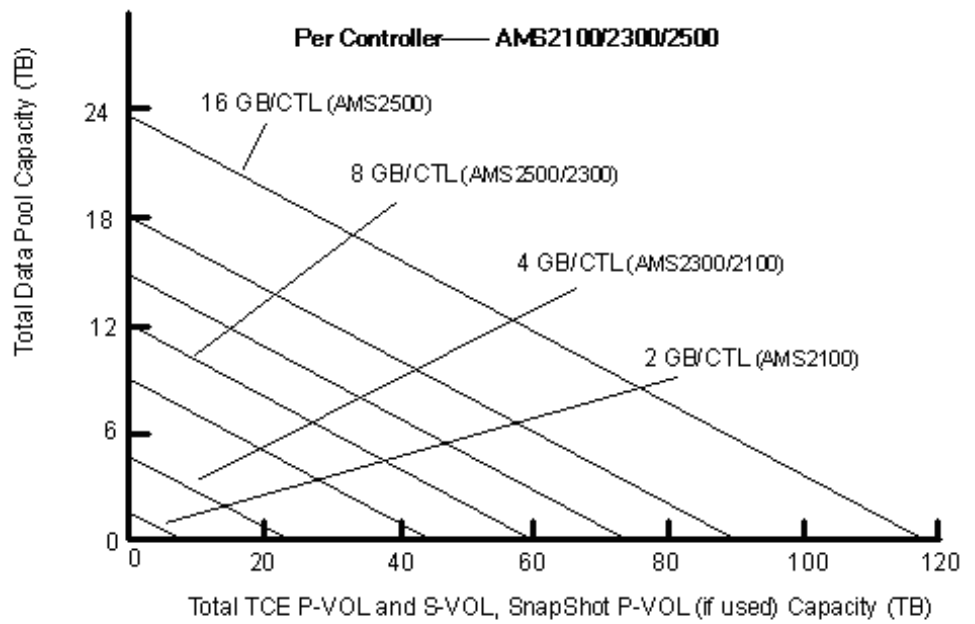


Figure 4-1: Relation of Data Volume, Data Pool Capacities to Cache Size per Controller

Requirements and specifications

This chapter provides TCE system requirements and specifications. Cautions and restrictions are also provided.

- [TCE system requirements](#)
- [TCE system specifications](#)

TCE system requirements

Table 5-1 describes the minimum TCE requirements.

Table 5-1: TCE Requirements

| Item | Minimum Requirements |
|--------------------------------------|---|
| AMS firmware version | <p>Version 0832/B or later is required for AMS 2100 or AMS 2300 arrays with hardware Rev. 0100.</p> <p>Version 0840/A or later is required for AMS2500 arrays with hardware Rev. 0100.</p> <p>Version 0890/A or later is required for AMS2100, 2300, or 2500 arrays with hardware Rev. 0200.</p> <p>Firmware version 0890/A or more is required on both local side and remote side arrays when connecting Rev. 0200 hardware.</p> |
| Navigator 2 version | <p>Version 3.21 or higher is required for the management PC for AMS 2100 or 2300 arrays where the hardware Rev. is 0100.</p> <p>Version 4.00 or higher is required for the management PC for an AMS2500 array where the hardware Rev. is 0100.</p> <p>Version 9.00 or higher is required for the management PC for AMS 2100, 2300, or 2500 where the hardware Rev. is 0200.</p> |
| CCI version | 01-21-03/06 or later is required for Windows host only |
| Number of AMS arrays | 2 |
| Supported array AMS models | AMS2100/2300/2500 |
| TCE license keys | One per array. |
| Number of controllers: | 2 (dual configuration) |
| Volume size | S-VOL block count = P-VOL block count. |
| Command devices per array (CCI only) | 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. |

Displaying the hardware revision number

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.

The screenshot shows the configuration page for array AMS2300_85000026. The 'Summary' table is as follows:

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

The 'Common Array Tasks' section includes:

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

TCE system specifications

Table 5-2 describes the TCE specifications.

Table 5-2: TCE Specifications

| Parameter | TCE Specification |
|---|--|
| User interface | <ul style="list-style-type: none"> Navigator 2 GUI Navigator 2 CLI CCI |
| Controller configuration | Configuration of dual controller is required. |
| Cache memory | <ul style="list-style-type: none"> AMS2100: 2 GB/controller AMS2300: 2, 4 GB/controller AMS2500: 2, 4, 6, 8 GB/controller |
| Host interface | AMS 2100, 2300, and 2500: Fibre channel or iSCSI (cannot mix) |
| Remote path | <ul style="list-style-type: none"> One remote path per controller is required—totaling two for a pair. The interface type of multiple remote paths between local and remote arrays must be the same. |
| Number of hosts when remote path is iSCSI | Maximum number of connectable hosts per port: 239. |

Table 5-2: TCE Specifications (Continued)

| Parameter | TCE Specification |
|---|---|
| Data pool | <ul style="list-style-type: none"> • Recommended minimum size: 20 GB • Maximum # of data pools per array: 64 • Maximum # of LUs that can be assigned to one data pool: 64 • Maximum # of LUs that can be used as data pools: 128. • When the array firmware version is less than 0852/A, a unified LU cannot be assigned to a data pool. If 0852/A or higher, a unified LU can be assigned to a data pool. • Data pools must set up for both the P-VOL and S-VOL. |
| Port modes | Initiator and target intermix mode. One port may be used for host I/O and TCE at the same time. |
| Bandwidth | <ul style="list-style-type: none"> • Minimum: 1.5 M. • Recommended: 100M or more. • When low bandwidth is used: <ul style="list-style-type: none"> - The time limit for execution of CCI commands and host I/O must be extended. - Response time for CCI commands may take several seconds. |
| License | Key is required. |
| Command device (CCI only) | <ul style="list-style-type: none"> • Required for CCI. • Minimum size: (33 MB; 65,538 blocks (1 block = 512 bytes) • Must be set up on local and remote arrays. • Maximum # allowed per array: 128 |
| DMLU | <ul style="list-style-type: none"> • Required. • Must be set up on local and remote arrays. • Minimum capacity per DMLU: 10 GB • Maximum number allowed per array: 2 • If setting up two DMLUs on an array, they should belong to different RAID groups. |
| Maximum # of LUs that can be used for TCE pairs | <ul style="list-style-type: none"> • AMS2100: 1,022 • AMS2300: 2,046 • AMS2500: 2,046 <p>The maximum when different types of arrays are used for TCE (i.e. AMS500 and AMS2100) is the array with the smallest maximum.</p> |
| Pair structure | One S-VOL per P-VOL. |
| Supported RAID level | <ul style="list-style-type: none"> • 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 | Local RAID level can be different than remote level. The number of data disks does not have to be the same. |
| Size of pair volumes | LU size of the P-VOL and S-VOL must be equal—identical block counts. |
| Types of drive for P-VOL, S-VOL, and data pool | If the drive types are supported by the array, they can be set for a P-VOL, an S-VOL, and data pools. SAS, SSD and SATA drives are supported for all volumes. SAS or SSD drives are recommended. Set all configured LUs using the same drive type. |
| Supported capacity value of P-VOL and S-VOL | Capacity is limited. See Maximum supported capacity on page 4-9 . |
| Copy pace | User adjustable rate that data is copied to remote array. See the copy pace step on page 7-5 for more information. |

Table 5-2: TCE Specifications (Continued)

| Parameter | TCE Specification |
|---|---|
| Consistency Group (CTG) | <ul style="list-style-type: none"> • Maximum allowed: 16 • Maximum # of pairs allowed per consistency group: <ul style="list-style-type: none"> - AMS2100: 1,022 - AMS2300: 2,046 - AMS2500: 2,046 |
| Management of LUs while using TCE | <p>A TCE pair must be deleted before the following operations:</p> <ul style="list-style-type: none"> • Deleting the pair's RAID group, LU, or data pool • Formatting an LU in the pair • Growing or shrinking an LU in the pair |
| Pair creation using unified LUs | <ul style="list-style-type: none"> • A TCE pair can be created using a unified LU. <ul style="list-style-type: none"> - When array firmware is less than 0852/A, the size of each LU making up the unified LU must be 1 GB or larger. - When the array firmware is 0852/A or later, there are no restrictions on the LUs making up the unified LU. • LUs that are already in a P-VOL or S-VOL cannot be unified. • Unified LUs that are in a P-VOL or S-VOL cannot be released. |
| Restriction during RAID group expansion | <p>A RAID group in which a TCE P-VOL or data pool exists can be expanded only when pair status is Simplex or Split. If the TCE data pool is shared with SnapShot, the SnapShot pairs must be in Simplex or Paired status.</p> |
| Unified LU for data pool | Not allowed. |
| Differential data | When pair status is Split, data sent to the P-VOL and S-VOL are managed as differential data. |
| Host access to a data pool | A data pool LU is hidden from a host. |
| Expansion of data pool capacity | <ul style="list-style-type: none"> • Data pools can be expanded by adding an LU. • Mixing SAS/SSD and SATA drives in a data pool is not supported. Set all configured LUs using the same drive type. |
| Reduction of data pool capacity | Yes. The pairs associated with a data pool must be deleted before the data pool can be reduced. |
| Failures | <ul style="list-style-type: none"> • When the copy operation from P-VOL to S-VOL fails, TCE suspends the pair (Failure). Because TCE copies data to the remote S-VOL regularly, data is restored to the S-VOL from the update immediately before the occurrence of the failure. • A drive failure does not affect TCE pair status because of the RAID architecture. |
| Data pool usage at 100% | When data pool usage is 100%, the status of any pair using the pool becomes Pool Full. P-VOL data cannot be updated to the S-VOL. |
| Array restart at TCE installation | The array is restarted after installation to set the data pool, unless it is also used by SnapShot. Then there is no restart. |
| TCE use with TrueCopy | Not Allowed. |
| TCE use with SnapShot | SnapShot can be cascaded with TCE or used separately. Only a SnapShot P-VOL can be cascaded with TCE. |
| TCE use with ShadowImage | Although TCE can be used at the same time as a ShadowImage system, it cannot be cascaded with ShadowImage. |
| TCE use with LUN Expansion | When firmware version is less than 0852/A, it is not allowed to create a TCE pair using unified LUs, which unify the LU with 1 GB or less capacity. |

Table 5-2: TCE Specifications (Continued)

| Parameter | TCE Specification |
|--------------------------------------|---|
| TCE use with Data Retention Utility | <p>Allowed.</p> <ul style="list-style-type: none"> • When S-VOL Disable is set for an LU, a pair cannot be created using the LU as the S-VOL. • S-VOL Disable can be set for an LU that is currently an S-VOL, if pair status is Split. |
| TCE use with Cache Residency Manager | <p>Allowed. However, an LU specified by Cache Residency Manager cannot be used as a P-VOL, S-VOL, or data pool.</p> |
| TCE use with Cache Partition Manager | <ul style="list-style-type: none"> • TCE can be used together with Cache Partition Manager. • Make the segment size of LUs to be used as a TCE data pool no larger than the default, (16 KB). • See Appendix D, Installing TCE when Cache Partition Manager is in use for details on initialization. |
| TCE use with SNMP Agent | <p>Allowed. A trap is transmitted for the following:</p> <ul style="list-style-type: none"> • Remote path failure. • Threshold value of the data pool is exceeded. • Actual cycle time exceeds the default or user-specified value. • Pair status changes to: <ul style="list-style-type: none"> - Pool Full - Failure. - Inconsistent because the data pool is full or because of a failure. |
| TCE use with Volume Migration | <p>Allowed. However, a Volume Migration P-VOL, S-VOL, or Reserved LU cannot be used as a TCE P-VOL or S-VOL.</p> |
| TCE use with Power Saving | <p>Allowed, however, pair operations are limited to split and delete.</p> |
| Reduction of memory | <p>Reduce memory only after disabling TCE.</p> |
| Load balancing | <p>Not supported.</p> |

Installation and setup

This chapter provides TCE installation and setup procedures using the Navigator 2 GUI. Instructions for CLI and CCI can be found in the appendixes.

- ❑ [Installation procedures](#)
- ❑ [Setup procedures](#)

Installation procedures

The following sections provide instructions for installing, enabling/disabling, and uninstalling TCE. Please note the following:

- TCE must be installed 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.

In cases where the DKN-200-NGW1 (NAS unit) is connected to the disk array, check the following items in advance.

1. Prior to this operation, execute Correspondence when connecting the NAS unit if each of the following three items apply to the disk array.
 - NAS unit is connected to the disk array. Ask the disk array administrator to confirm whether the NAS unit is connected or not.
 - NAS unit is in operation. Ask the NAS unit administrator to confirm whether the NAS service is operating or not.
 - A failure has not occurred on the NAS unit. Ask the NAS unit administrator to check whether failure has occurred or not by checking with the NAS administration software, NAS Manager GUI, List of RAS Information, etc. In case of failure, execute the maintenance operation together with the NAS maintenance personnel.
2. Correspondence when connecting the NAS unit.
 - If the NAS unit is connected, ask the NAS unit administrator for termination of NAS OS and planned shutdown of the NAS unit.
3. Points to be checked after completing this operation:
 - Ask the NAS unit administrator to reboot the NAS unit. After rebooting, ask the NAS unit administrator to refer to "Recovering from FC path errors" in the *Hitachi NAS Manager User's Guide* and check the status of the Fibre Channel path (FC path in short) and to recover the FC path if it is in a failure status.
 - In addition, if there are any personnel for the NAS unit maintenance, ask the NAS unit maintenance personnel to reboot the NAS unit.

Installing TCE

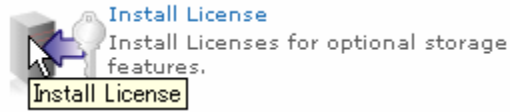
Prerequisites

- A key code or key file is required to install or uninstall TCE. 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>.
- The array may require a restart at the end of the installation procedure. If SnapShot is enabled at the time, no restart is necessary.
- If restart is required, it can be done either when prompted or at a later time.

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

To install TCE

1. In the Navigator 2 GUI, click the check box for the array where you want to install TCE, 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.
6. On the Reboot Array screen, click the **Reboot Array** button to reboot, or click **Close** to finish the installation without rebooting.
7. When the reboot is complete, click **Close**.

Enabling or disabling TCE

TCE is automatically enabled when it is installed. You can disable or re-enable it.

Prerequisites

- To enable when using TCE with iSCSI, there must be fewer than 240 connected to a port on the array.
- When disabling TCE:
 - pairs must be deleted and the status of the logical units must be Simplex.
 - Data pools must be deleted, unless SnapShot will continue to be used.
 - The remote path must be deleted.

To enable or disable TCE

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 **TC-Extended** in the Licenses list.
4. Click **Change Status**. The Change License screen displays.
5. 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 TCE is disabled. Click **Close**.

Uninstalling TCE

Prerequisite

- TCE pairs must be deleted. Volume status must be Simplex.
- Data pools must be deleted, unless SnapShot will continue to be used.
- The remote path must be deleted.
- A key code or key file is required. 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>.

To uninstall TCE

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 **TC-Extended** 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, data pools, CHAP secret (iSCSI only), and remote path.

Setting up DMLUs

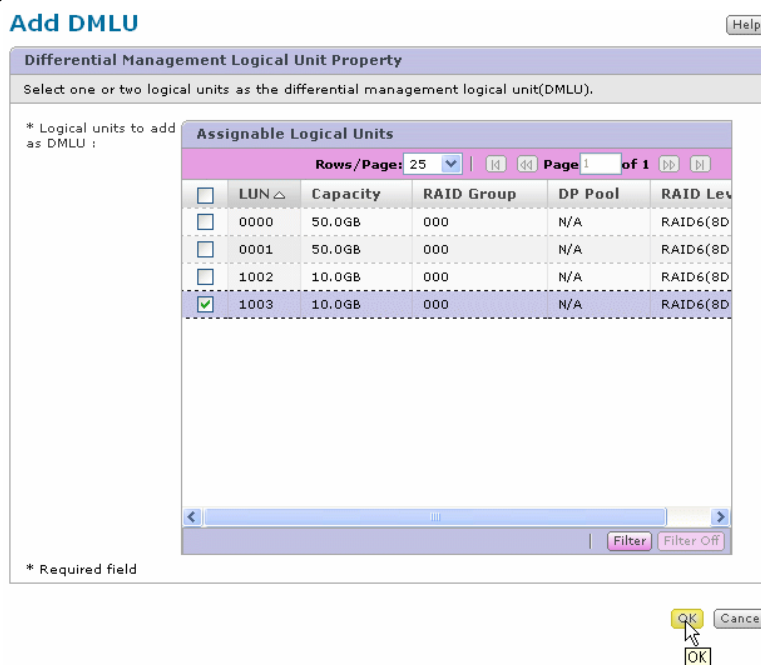
The DMLU (differential management-logical unit) must be set up prior to using TCE. The DMLU is used by the system for storing TCE status information when the array is powered down.

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 (recommended size).
- DMLUs must be set up on both the local and remote arrays.
- One DMLU is required on each array; two are recommended, the second used as backup. However, no more than two DMLUs can be installed per array.
- When setting up more than one DMLU, assign them to different RAID groups to provide a backup in the event of a drive failure.
- Specifications for DMLUs should also be reviewed. See [TCE system specifications on page 5-3](#).

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**.

Setting up data pools

On the local array, the data pool stores differential data before it is updated to the S-VOL. On the remote array, the data pool stores the S-VOL's previous update as a data-consistent backup when the current update is occurring. See [Data pools on page 1-4](#) for more descriptive information.

Prerequisites

- To review the data pool sizing procedure, see [Calculating data pool size on page 2-4](#).
- Up to 64 LUs can be assigned to a data pool.
- Hitachi recommends a minimum of 20 GB for data pool size.
- A logical unit consisting of SAS/SSD drives and one consisting of SATA drives cannot be used for a data pool. Set all configured LUs using the same drive type.
- When Cache Partition Manager is used with TCE, the segment size of LUs belonging to a data pool must be the default size (16 kB) or less. See *Hitachi Storage Navigator Modular 2 Storage Features Reference Guide* for more information.

To create and assign volumes for data pools

1. In Navigator 2 GUI, select the desired array, then click the **Show & Configure Array** button.
2. From the Replication tree, select the **Replication** icon, then select the **Setup** icon. The Setup screen displays.
3. Select **Data Pools**. View screen instructions by clicking the Help button.



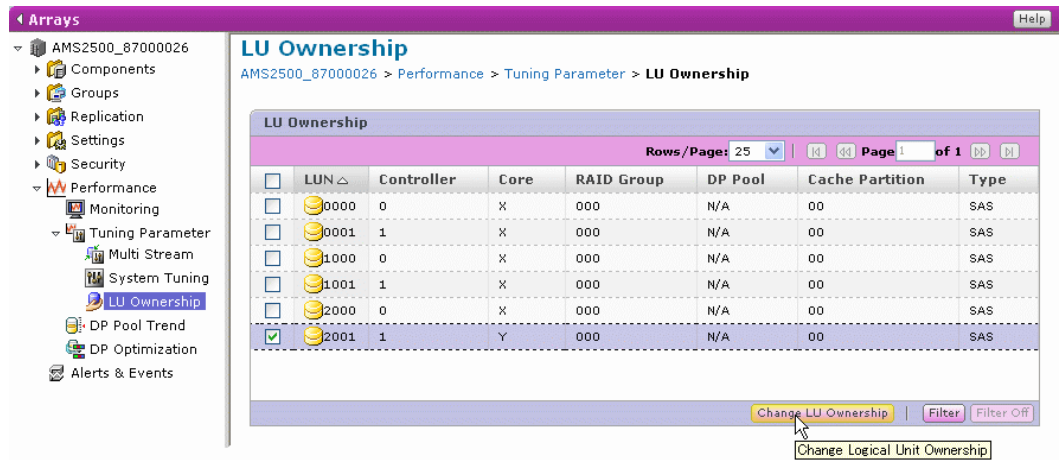
NOTE: The default **Threshold** value is 70%. When capacity reaches the Threshold plus 1 percent, both data pool and pair status change to "Threshold over", and the array issues a warning. If capacity reaches 100 percent, the pair fails and all data in the S-VOL is lost.

Setting the LU ownership

The load balancing function is not applied to the LUs specified as a TCE pair. Since the ownership of the LUs specified as a TCE pair is the same as the ownership of the LUs specified as a data pool, perform the setting so that the ownership of LUs specified as a data pool is balanced in advance.

Perform the following steps to set the LU ownership:

1. Select the **LU Ownership** icon in the Tuning Parameter tree view of the Performance tree view.



The Change Logical Unit Ownership screen appears.

2. Select Controller 0 or Controller 1 and X Core or Y Core, then click OK.

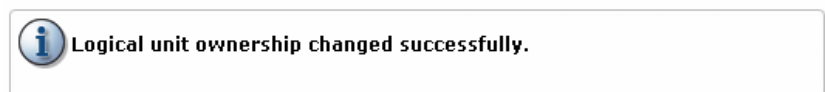
Change Logical Unit Ownership - 2001



The Core item is not displayed: AMS 2010/2100/2300

3. The LU ownership change message appears. Click **Close**.

Change Logical Unit Ownership - 2001



Adding or changing the remote port CHAP secret

(For arrays with iSCSI connectors 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.

- CHAP authentication is automatically configured with a default CHAP secret when the TCE Setup Wizard is used. You can change the default secret if desired.
- CHAP authentication is not configured when the Create Pair procedure is used, but it can be added.

Prerequisites

- 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.

1. 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 on-screen instructions.
6. Click **OK** when finished.
7. The confirmation message appears. Click **Close**.

Add Remote Port CHAP



To change a CHAP secret

1. Split the TCE pairs, after confirming first that the status of all pairs is Paired.
 - To confirm pair status, see [Monitoring pair status on page 9-2](#).
 - To split pairs, see [Splitting a pair on page 7-5](#).

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 9-9](#).
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 [Setting up the remote path on page 6-9](#).
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 a pair on page 7-6](#).

Setting up the remote path

A remote path is the data transfer connection between the local and remote arrays.

- 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.

The Navigator 2 GUI allows you to create the remote path in two ways:

- Use the TCE Backup Wizard, in which the remote path and the initial TCE pair are created. This is the simplest and quickest method for setting the remote path. See [TCE setup wizard on page 7-3](#).
- Use the Create Remote Path procedure, described below. Use this method when you will create the initial pair using the Create Pair procedure, rather than the TCE wizard. The Create Remote Path and Create Pair procedures allow for more customizing.

Prerequisites

- Both local and remote arrays must be connected to the network for the remote path.
- The remote array ID will be required. This 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 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 6-8](#) 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 **Create Path**. The **Create Remote Path** screen appears.

Create Remote Path

Remote Path Properties

Define the settings for a remote path to the remote array. If CHAP authentication for iSCSI is enabled, enter secrets into both the secret and secret (retype) fields. A secret can consist of 12 to 32 characters (alphanumeric characters or special symbols: ',', '-', '+', '@', '_', '=', '!', '/', '[', ']', '\', or space).

* Interface Type: Fibre iSCSI

* Remote Array ID: 85000027

* Bandwidth: 1000 Mbps
From 0.2 to 8000.0

Remote Path 0:

* Local Port: 0A

* Remote Port: 0A

Remote Path 1:

* Local Port: 1A

* Remote Port: 1A

* Required field

OK Cancel

OK

3. For Interface Type, select **Fibre** or **iSCSI**.
4. Enter the **Remote Array ID**.
5. Enter **Bandwidth**. When connecting the array directly to the other array, set the bandwidth according to the transfer rate.
6. (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. [Table 6-1](#) shows port selections by array type

Table 6-1: Port selections by AMS array type

| Path | Local array type | | | | Remote array type | | | |
|--------------------|-------------------------|-------------------------|-------------------------|--------------------------------|-------------------------|--------------------------------|--------------------------------|--|
| | 2100 with H/W Rev. 0100 | 2300 with H/W Rev. 0100 | 2100 with H/W Rev. 0200 | 2500/2300 with H/W Rev. 0200 | 2100 with H/W Rev. 0100 | 2300 with H/W Rev. 0100 | 2100 with H/W Rev. 0200 | 2500/2300 with H/W Rev. 0200 |
| Path 0 Port | 0A or 0B | 0A, 0B, 0C, 0D | 0A, 0B, 0E, 0F | 0A, 0B, 0C, 0D, 0E, 0F, 0G, 0H | 0A, 0B, 1A, 1B | 0A, 0B, 0C, 0D, 1A, 1B, 1C, 1D | 0A, 0B, 0E, 0F, 1A, 1B, 1E, 1F | 0A, 0B, 0C, 0D, 0E, 0F, 0G, 0H, 1A, 1B, 1C, 1D, 1E, 1F, 1G, 1H |
| Path 1 Port | 1A or 1B | 1A, 1B, 1C, 1D | 1A, 1B, 1E, 1F | 1A, 1B, 1C, 1D, 1E, 1F, 1G, 1H | 0A, 0B, 1A, 1B | 0A, 0B, 0C, 0D, 1A, 1B, 1C, 1D | 0A, 0B, 0E, 0F, 1A, 1B, 1E, 1F | 0A, 0B, 0C, 0D, 0E, 0F, 0G, 0H, 1A, 1B, 1C, 1D, 1E, 1F, 1G, 1H |

8. Click **Ok**.

Pair operations

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

- ❑ [TCE operations](#)
- ❑ [Checking pair status](#)
- ❑ [Creating the initial copy](#)
- ❑ [Splitting a pair](#)
- ❑ [Resynchronizing a pair](#)
- ❑ [Swapping pairs](#)

TCE operations

Basic TCE operations consist of the following:

- Checking pair status. Each operation requires the pair to be in a specific status.
- Creating the pair, in which the S-VOL becomes a duplicate of the P-VOL.
- Splitting the pair, which stops updates from the P-VOL to the S-VOL and allows read/write of the S-VOL.
- Re-synchronizing the pair, in which the S-VOL again mirrors the on-going, current data in the P-VOL.
- Swapping pairs, which reverses pair roles.
- Deleting a pair, data pool, DMLU, or remote path.
- Editing pair information.

These operations are described in the following sections. All procedures relate to the Navigator2 GUI.

Checking pair status

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

- Find an operation's status requirement in the Prerequisites sections below.
- To monitor pair status, refer to [Monitoring pair status on page 9-2](#).

Creating the initial copy

Two methods are used for creating the initial TCE copy:

- The GUI setup wizard, which is the simplest and quickest method. Includes remote path setup.
- The GUI Create Pair procedure, which requires more setup but allows for more customizing.

Both procedures are described in this section.

During pair creation:

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

Prerequisites and best practices for pair creation

- Both arrays must be able to communicate with each other via their respective controller 0 and controller 1 ports.
- Bandwidth for the remote path must be known.

- Local and remote arrays must be able to communicate with the Hitachi Storage Navigator 2 server, which manages the arrays.
- Logical units must be set up and formatted on the remote array for the secondary volume or volumes.
 - You will be required to enter the LUN for the S-VOL if using the Create Pair procedure (S-VOL LUs are automatically selected by the setup wizard).
 - In the Create Pair procedure, the LUN for the S-VOL must be the same as the corresponding P-VOL's LUN.
 - Block size of the S-VOL must be the same as the P-VOL.
- Two data pools must be set up on the local array and two on the remote array. You will be required to enter the LUN for the remote data pools.
- DMLUs must be set up on both arrays.
- The remote array ID is required during both initial copy procedures. This is listed on the highest-level GUI screen for the array.
- The create pair and resynchronize operations affect performance on the host. Best practice is to perform the operation when I/O load is light.
- For bi-directional pairs (host applications at the local and remote sites write to P-VOLs on the respective arrays), creating or resynchronizing pairs may be performed at the same time. However, best practice is to perform the operations one at a time to lower performance impact.

TCE setup wizard

To create a pair using the backup wizard

1. In Navigator 2 GUI, select the local array then click the **Show & Configure Array** button.
2. If Password Protection is installed and enabled, log in with the registered user ID and password for the array.
3. On the array page under Common Array Tasks, click the **TrueCopy Extended Distance** link. The TrueCopy Extended Distance Setup Wizard opens.
4. Review the Introduction screen, then click **Next**. The Setup Remote Path screen displays.
5. Enter the **Remote Array ID**.
6. Select the LUN whose data you want to copy to the remote array. Click **Next**.
7. On the Confirm screen, review pair information and click **Confirm**.
8. On the completion screen, click **Finish**.

If you are using iSCSI, you may want to change the CHAP secret, which was created during the wizard procedure automatically using a default. See [Adding or changing the remote port CHAP secret on page 6-8](#) for details.

Create pair procedure

With the Create Pair procedure, you create a TCE pair and specify copy pace, consistency groups, and other options. Please review the prerequisites on page 7-2 before starting.

To create a pair using the Create Pair procedure

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 **Create Pair** button at the bottom of the screen.
4. On the Create Pair screen displays, confirm that the **Copy Type** is TCE and enter a name in the **Pair Name** box following on-screen guidelines. If omitted, the pair is assigned a default name. In either case, the pair is named in the local array, but not in the remote array. On the remote array, the pair appears with no name. Add a name using Edit Pair.
5. In the **Primary Volume** field, select the local primary volumes that you want to copy to the remote side.



NOTE: In Windows 2003 Server, LUs are identified by HLUN. The LUN and H-LUN may be different. See [Identifying P-VOL and S-VOL LUs on Windows on page 4-6](#) to map LUN to HLUN.

6. In the **Secondary Volume** box, enter the S-VOL LUN(s) on the remote array that the primary volume(s) will be copied to. Remote LUNs must be:
 - the same as local LUNs.
 - the same size as local LUNs.
7. From the **Pool Number of Local Array** drop-down list, select the previously set up data pool.
8. In the **Pool Number of Remote Array** box, enter the LUN set up for the data pool on the remote array.
9. For **Group Assignment**, you assign the new pair to a consistency group.
 - 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.
 - If you do not want to assign the pair to a consistency group, they will be assigned automatically. Leave the **New or existing Group Number** button selected with no number entered in the box.



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

- a. After completing the create pair procedure, on the **Pairs** screen, check the box for the pair belonging to the group.
 - b. Click the **Edit Pair** button.
 - c. On the Edit Pair screen, enter the **Group Name** and click **OK**.
-
10. Select the **Advanced** tab.
 11. From the **Copy Pace** drop-down list, select a pace. Copy pace is the rate at which a pair is created or resynchronized. The time required to complete this task depends on the I/O load, the amount of data to be copied, cycle time, and bandwidth. Select one of the following:
 - Slow — The option takes longer when host I/O activity is high. The time to copy may be quite lengthy.
 - Medium — (Recommended) The process is performed continuously, but copying does not have priority and the time to completion is not guaranteed.
 - Fast — The copy/resync process is performed continuously and has priority. Host I/O performance will be degraded. The time to copy can be guaranteed because it has priority.
 12. In the **Do initial copy from the primary volume...** field, leave **Yes** checked to copy the primary to the secondary volume.

Clear the check box to create a pair without copying the P-VOL at this time, and thus reduce the time it takes to set up the configuration for the pair. Use this option also when data in the primary and secondary volumes already match. The system treats the two volumes as paired even though no data is presently transferred. Resync can be selected manually at a later time when it is appropriate.
 13. Click **OK**, then click **Close** on the confirmation screen that appears. The pair has been created.

Splitting a pair

Data is copied to the S-VOL at every update cycle until the pair is split.

- When the split is executed, all differential data accumulated in the local array is updated to the S-VOL.
- After the split operation, write updates continue to the P-VOL but not to the S-VOL.

After the Split Pair operation:

- S-VOL data is consistent to P-VOL data at the time of the split. The S-VOL can receive read/write instructions.
- The TCE pair can be made identical again by re-synchronizing from primary-to-secondary or secondary-to-primary.

The pair must be in Paired status. The time required to split the pair depends on the amount of data that must be copied to the S-VOL so that the data is current with the P-VOL's data.

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. Select the pair you want to split.
4. Click the **Split Pair** button at the bottom of the screen. View further instructions by clicking the Help button, as needed.

Resynchronizing a pair

Re-synchronizing a pair 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. It must be deleted and created again.
- Best practice is to perform a resynchronization when I/O load is low, to reduce impact on host activities.

Prerequisites

- The pair must be in Split, Failure, or Pool Full status.

To resync the pair

1. In the 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. The direction of data flow is also reversed.

This is done when host operations are switched to the S-VOL, and when host-storage operations are again functional on the local array.

Prerequisites and Notes

- The pair must be in Paired, Split, or Pool Full 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 swap is attempted when a remote path failure exists, the primary/secondary roles are not reversed and data is restored from the remote data pool to the S-VOL.
- When a pair that is in a consistency group is swapped, all pairs in the group are swapped.

To swap TCE pairs

1. 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.
7. If the message, "DMER090094: The LU whose pair status is Busy exists in the target group" displays, proceed as follows:
 - a. Check the pair status for each LU in the target group. Pair status will change to Takeover. Confirm this before proceeding. Click the **Refresh Information** button to see the latest status.
 - b. When the pairs have changed to Takeover status. execute the Swap command again,

Example scenarios and procedures

This chapter describes four use-cases and the processes for handling them.

- ❑ [CLI scripting procedure for S-VOL backup](#)
- ❑ [Procedure for swapping I/O to S-VOL when maintaining local array](#)
- ❑ [Procedure for moving data to a remote array](#)
- ❑ [Process for disaster recovery](#)

CLI scripting procedure for S-VOL backup

SnapShot can be used with TCE to maintain timed backups of S-VOL data. The following illustrates and explains how to perform TCE and SnapShot operations.

An example scenario is used in which three hosts, A, B, and C, write data to logical units on the local array, as shown in [Figure 8-1](#).

- A database application on host A writes to LU1 and LU2.
- A file system application on host B, and a mail server application on host C, store their data as indicated in the graphic.
- Database storage, LU1 and LU2 on the local array, is backed up every night at 11 o'clock to LU1 and LU2 (TCE S-VOLs) on the remote array.
- The TCE S-VOLs are backed up daily, using SnapShot. Each SnapShot backup is held for seven days. There are seven SnapShot backups (LU101 through LU161 shown in the graphic on the remote side).
- The LUs for the other applications are also backed up with SnapShot on the remote array, as indicated. These snapshots are made at different times than the database snapshots to avoid performance problems.
- Each host is connected by a LAN to the arrays.
- CLI scripts are used for TCE and SnapShot operations.

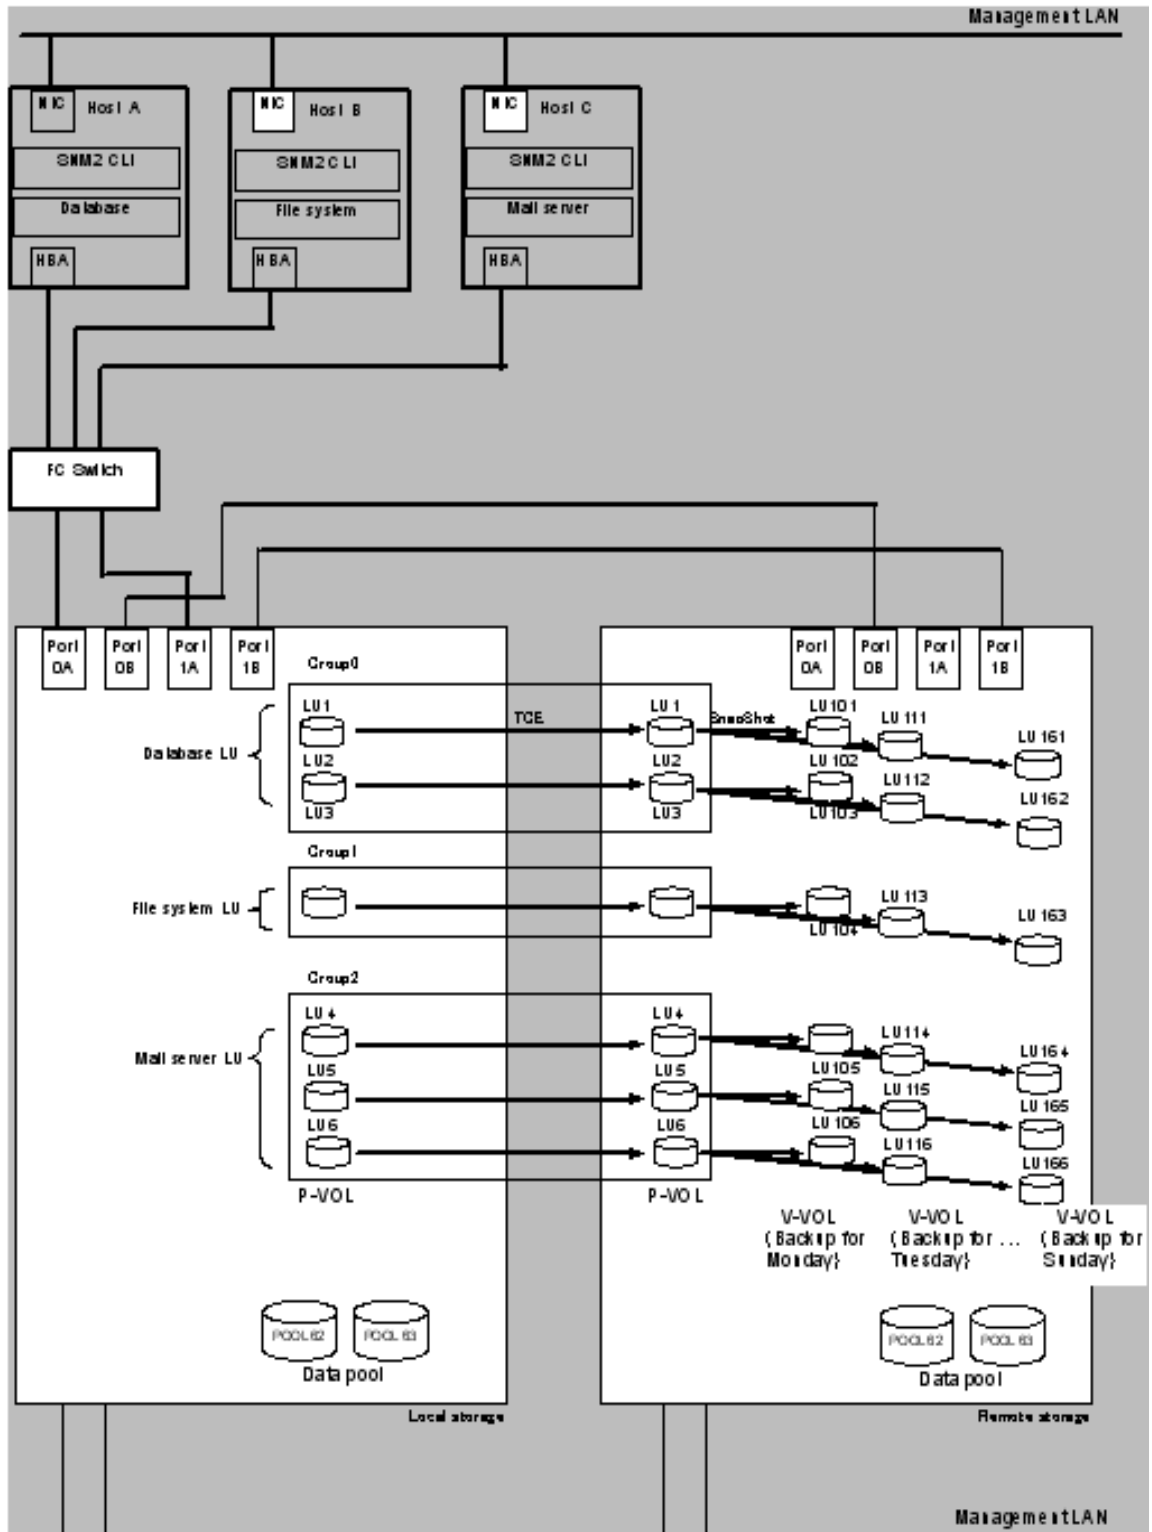


Figure 8-1: Configuration Example for a Remote Backup System

Scripted TCE, SnapShot procedure

The TCE/SnapShot system shown in [Figure 8-1](#) is set up using the Navigator 2 GUI. Day-to-day operations are handled using CCI or CLI scripts. In this example, CLI scripts are used.

To manage the various operations, they are broken down by host, LU, and backup schedule, as shown in [Table 8-1](#).

Table 8-1: TCE Logical Units by Host Application

| Host Name | Application | LU to Use Backup Target LU) | Backup in Remote Array (SnapShot Logical Unit) | | | |
|-----------|-------------|-----------------------------|--|-------------|-----|------------|
| | | | For Monday | For Tuesday | ... | For Sunday |
| Host A | Database | LU1 (D drive) | LU101 | LU111 | ... | LU161 |
| | | LU2 (E drive) | LU102 | LU112 | ... | LU162 |
| Host B | File system | LU3 (D drive) | LU103 | LU113 | ... | LU163 |
| Host C | Mail server | LU4 (M drive) | LU104 | LU114 | ... | LU164 |
| | | LU5 (N drive) | LU105 | LU115 | ... | LU165 |
| | | LU6 (O drive) | LU106 | LU116 | ... | LU166 |

In the procedure example that follows, scripts are executed for host A on Monday at 11 p.m. The following assumptions are made:

- The system is completed.
- The TCE pairs are in Paired status.
- The SnapShot pairs are in Split status.
- Host A uses a Windows operating system.

The variables used in the script are shown in [Table 8-2](#). The procedure and scripts follow.

Table 8-2: CLI Script Variables and Descriptions

| # | Variable Name | Content | Remarks |
|---|------------------------------------|--|--|
| 1 | STONAVM_HOME | Specify the directory in which SNM2 CLI was installed. | When the script is in the directory in which SNM2 CLI was installed, specify ".". |
| 2 | STONAVM_RSP_PASS | Be sure to specify "on" when executing SNM2 CLI in the script. | This is the environment variable to enter "Yes" automatically for the inquiry of SNM2 CLI command. |
| 3 | LOCAL | Name of the local array registered in SNM2 CLI | |
| 4 | REMOTE | Name of the remote array registered in SNM2 CLI | |
| 5 | TCE_PAIR_DB1, TCE_PAIR_DB2 | Name of the TCE pair generated at the setup | The default names are as follows. TCE_LUxxxx_LUyyyy xxxx: LUN of P-VOL yyyy: LUN of S-VOL |
| 6 | SS_PAIR_DB1_MON SS_PAIR_DB2_MON | Name of the SnapShot pair when creating the backup in the remote array on Monday | The default names are as follows. TCE_LUxxxx_LUyyyy xxxx: LUN of P-VOL yyyy: LUN of S-VOL |
| 7 | DB1_DIR DB2_DIR | Directory on the host where the LU is mounted | |
| 8 | LU1_GUID, LU2_GUID | GUID of the backup target LU recognized by the host | You can search it by the mountvol command of Windows®. |
| 9 | TIME | Time-out value of the aureplicationmon command | Make it longer than the time taken for the resynchronization of TCE. |

1. Specify the variables to be used in the script, as shown below.

```

set STONAVM_HOME=.
set STONAVM_RSP_PASS=on
set LOCAL=LocalArray
set REMOTE=RemoteArray
set TCE_PAIR_DB1=TCE_LU0001_LU0001
set TCE_PAIR_DB2=TCE_LU0002_LU0002
set SS_PAIR_DB1_MON=SS_LU0001_LU0101
set SS_PAIR_DB2_MON=SS_LU0002_LU0102

set DB1_DIR=D:\
set DB2_DIR=E:\
set LU1_GUID=xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxxxx
set LU2_GUID=yyyyyyyy-YYYY-YYYY-YYYY-YYYYYYYYYYYYY
set TIME=18000
(To be continued)

```

2. Stop the database application, then unmount the P-VOL, as shown below. Doing this stabilizes the data in the P-VOL.

```
(Continued from the previous section)
<Stop the access to C:\sms100\DB1 and C:\sms100\DB2>

REM Unmount of P-VOL
raidqry -x umount %DB1_DIR%
raidqry -x umount %DB2_DIR%
(To be continued)
```

Note that `raidqry` is a CCI command.

3. Split the TCE pair, then check that the pair status becomes Split, as shown below. This updates data in the S-VOL and makes it available for secondary uses, including SnapShot operations.

```
(Continued from the previous section)
REM pair split
aureplicationremote -unit %LOCAL% -split -tce -pairname
%TCE_PAIR_DB1% -gno 0
aureplicationremote -unit %LOCAL% -split -tce -pairname
%TCE_PAIR_DB2% -gno 0
REM Wait until the TCE pair status becomes Split.
aureplicationmon -unit %LOCAL% -evwait -tce -pairname
%TCE_PAIR_DB1% -gno 0 -st split -pvol -timeout %TIME%
aureplicationmon -unit %LOCAL% -evwait -tce -pairname
%TCE_PAIR_DB1% -gno 0 -nowait
IF NOT %ERRORLEVEL% == 13 GOTO ERROR_TCE_Split
aureplicationmon -unit %LOCAL% -evwait -tce -pairname
%TCE_PAIR_DB2% -gno 0 -st split -pvol -timeout %TIME%
aureplicationmon -unit %LOCAL% -evwait -tce -pairname
%TCE_PAIR_DB2% -gno 0 -nowait
IF NOT %ERRORLEVEL% == 13 GOTO ERROR_TCE_Split
(To be continued)
```

4. Mount the P-VOL, and restart the database application, as shown below.

```
(Continued from the previous section)
REM Mount of P-VOL
raidqry -x mount %DB1_DIR% Volume{%LU1_GUID%}
raidqry -x mount %DB2_DIR% Volume{%LU2_GUID%}

<Restart access to C:\sms100\DB1 and C:\sms100\DB2>
(To be continued)
```

5. Resynchronize the SnapShot backup. Then split the SnapShot backup. These operations are shown in the example below.

```
(Continued from the previous section)
REM Resynchronization of the SnapShot pair which is cascaded
aureplicationlocal -unit %REMOTE% -resync -ss -pairname
%SS_PAIR_DB1_MON% -gno 0
aureplicationlocal -unit %REMOTE% -resync -ss -pairname
%SS_PAIR_DB2_MON% -gno 0
REM Wait until the SnapShot pair status becomes Paired.
aureplicationmon -unit %REMOTE% -evwait -ss -pairname
%SS_PAIR_DB1_MON% -gno 0 -st paired -pvol -timeout %TIME%
aureplicationmon -unit %REMOTE% -evwait -ss -pairname
%SS_PAIR_DB1_MON% -gno 0 -nowait
IF NOT %ERRORLEVEL% == 12 GOTO ERROR_SS_Resync
aureplicationmon -unit %REMOTE% -evwait -ss -pairname
%SS_PAIR_DB2_MON% -gno 0 -st paired -pvol -timeout %TIME%
aureplicationmon -unit %REMOTE% -evwait -ss -pairname
%SS_PAIR_DB2_MON% -gno 0 -nowait
IF NOT %ERRORLEVEL% == 12 GOTO ERROR_SS_Resync

REM Pair split of the SnapShot pair which is cascaded
aureplicationlocal -unit %REMOTE% -split -ss -pairname
%SS_PAIR_DB1_MON% -gno 0
aureplicationlocal -unit %REMOTE% -split -ss -pairname
%SS_PAIR_DB2_MON% -gno 0
REM Wait until the SnapShot pair status becomes Split.
aureplicationmon -unit %REMOTE% -evwait -ss -pairname
%SS_PAIR_DB1_MON% -gno 0 -st split -pvol -timeout %TIME%
aureplicationmon -unit %REMOTE% -evwait -ss -pairname
%SS_PAIR_DB1_MON% -gno 0 -nowait
IF NOT %ERRORLEVEL% == 13 GOTO ERROR_SS_Split
aureplicationmon -unit %REMOTE% -evwait -ss -pairname
%SS_PAIR_DB2_MON% -gno 0 -st split -pvol -timeout %TIME%
aureplicationmon -unit %REMOTE% -evwait -ss -pairname
%SS_PAIR_DB2_MON% -gno 0 -nowait
IF NOT %ERRORLEVEL% == 13 GOTO ERROR_SS_Split
(To be continued)
```

- When the SnapShot backup operations are completed, re-synchronize the TCE pair, as shown below. When the TCE pair status becomes Paired, the backup procedure is completed.

```
(Continued from the previous section)
REM Return the pair status to Paired (Pair resynchronization)
aureplicationremote -unit %LOCAL% -resync -tce -pairname
%TCE_PAIR_DB1% -gno 0
aureplicationremote -unit %LOCAL% -resync -tce -pairname
%TCE_PAIR_DB2% -gno 0
REM Wait until the TCE pair status becomes Paired.
aureplicationmon -unit %LOCAL% -evwait -tce -pairname
%TCE_PAIR_DB1% -gno 0 -st paired -pvol -timeout %TIME%
aureplicationmon -unit %LOCAL% -evwait -tce -pairname
%TCE_PAIR_DB1% -gno 0 -nowait
IF NOT %ERRORLEVEL% == 12 GOTO ERROR_TCE_Resync
aureplicationmon -unit %LOCAL% -evwait -tce -pairname
%TCE_PAIR_DB2% -gno 0 -st paired -pvol -timeout %TIME%
aureplicationmon -unit %LOCAL% -evwait -tce -pairname
%TCE_PAIR_DB2% -gno 0 -nowait
IF NOT %ERRORLEVEL% == 12 GOTO ERROR_TCE_Resync
echo The backup is completed.
GOTO END
(To be continued)
```

- If pair status does not become Paired within the aureplicationmon command time-out period, perform error processing, as shown below.

```
(Continued from the previous section)
REM Error processing
:ERROR_TCE_Split
< Processing when the S-VOL data of TCE is not determined within
the specified time>
GOTO END
:ERROR_SS_Resync
< Processing when SnapShot pair resynchronization fails and the
SnapShot pair status does not become Paired>
GOTO END
:ERROR_SS_Split
< Processing when SnapShot pair split fails and the SnapShot pair
status does not become Split>
GOTO END
:ERROR_TCE_Resync
< Processing when TCE pair resynchronization does not terminate
within the specified time>
GOTO END

:END
```

Procedure for swapping I/O to S-VOL when maintaining local array

The following shows a procedure for temporarily shifting I/O to the S-VOL in order to perform maintenance on the local array. In the procedure, host server duties are switched to a standby server.

1. On the local array, stop the I/O to the P-VOL.
2. Split the pair, which makes P-VOL and S-VOL data identical.
3. On the remote site, execute the swap pair command. Since no data is transferred, the status is changed to Paired after one cycle time.
4. Split the pair.
5. Restart I/O, using the S-VOL on the remote array.
6. On the local site, perform maintenance on the local array.
7. When maintenance on the local array is completed, resynchronize the pair from the remote array. This copies the data that has been updated on the S-VOL during the maintenance period.
8. On the remote array, when pair status is Paired, stop I/O to the remote array, and un-mount the S-VOL.
9. Split the pair, which makes data on the P-VOL and S-VOL identical.
10. On the local site, issue the pair swap command. When this is completed, the S-VOL in the local array becomes the P-VOL again.
11. Business can restart at the local site. Mount the new P-VOL on the local array to local host server and restart I/O.

Procedure for moving data to a remote array

This section provides a procedure in which application data is copied to a remote array, and the copied data is analyzed. An example scenario is used in which:

- A database application on host A writes to the P-VOL logical units LU1 and LU2, as shown in [Figure 8-2](#).
- The P-VOL LUs are in the same consistency group (CTG).
- The P-VOL LUs are paired with the S-VOL LUs, LU1 and LU2 on the remote array.
- A data-analyzing application on host D analyzes the data in the S-VOL. Analysis processing is performed once every hour.

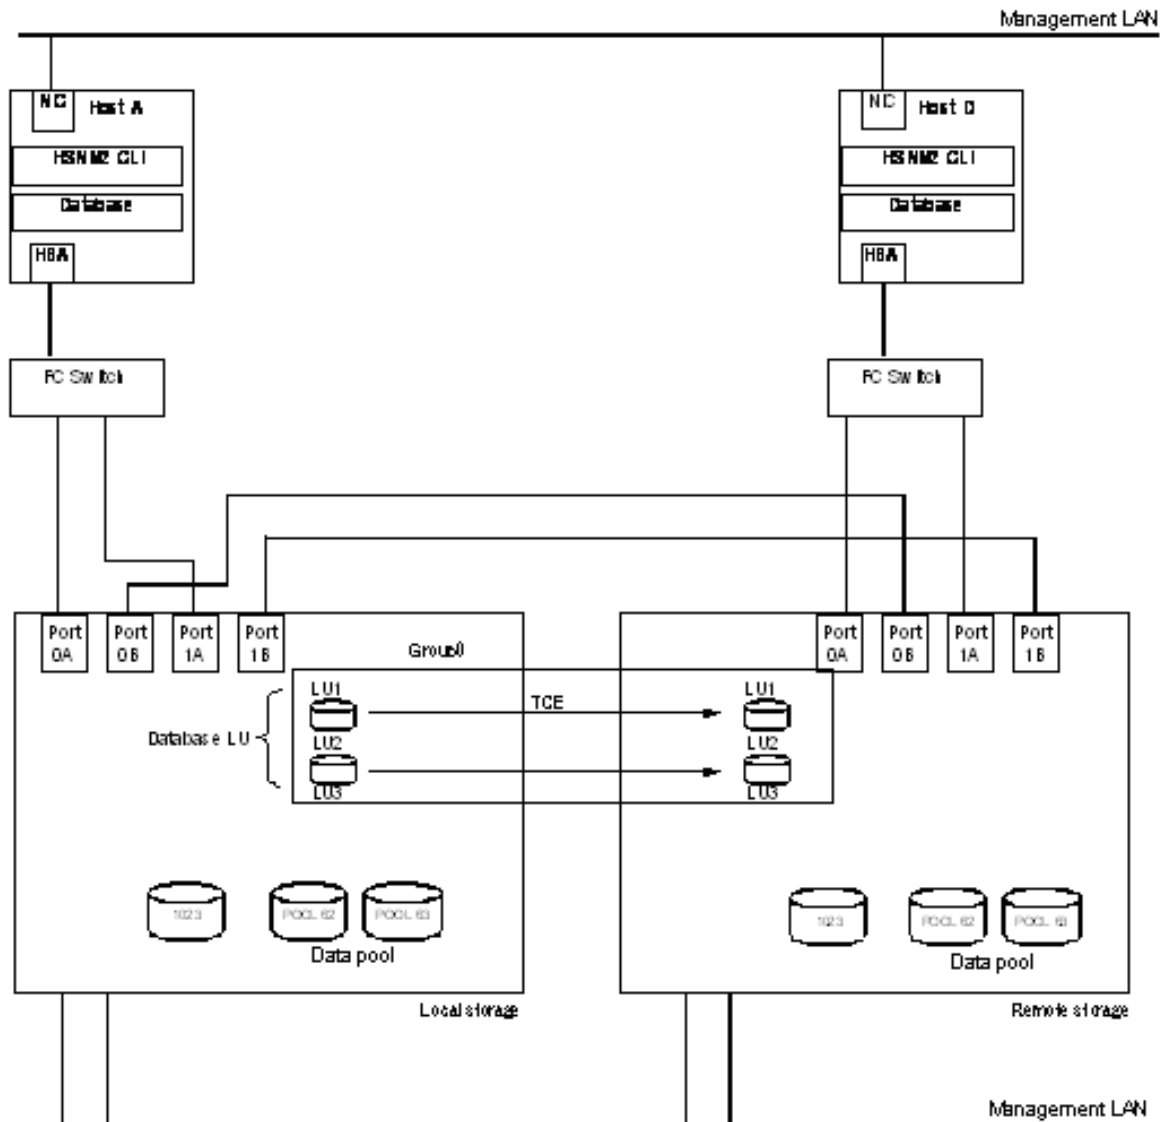


Figure 8-2: Configuration Example for Moving Data

Example procedure for moving data

1. Stop the applications that are writing to the P-VOL, then un-mount the P-VOL.
2. Split the TCE pair. Updated differential data on the P-VOL is transferred to the S-VOL. Data in the S-VOL is stabilized and usable after the split is completed.
3. Mount the P-VOL and then resume writing to the P-VOL.
4. Mount the S-VOL.
5. Read and analyze the S-VOL data. The S-VOL data can be updated, but updated data will be lost when the TCE pair is resynchronized. If updated data is necessary, be sure to backup the S-VOL on the remote side.
6. Un-mount the S-VOL.
7. Re-synchronize the TCE pair.

Process for disaster recovery

This section explains behaviors and the general process for continuing operations on the S-VOL and then failing back to the P-VOL, when the primary site has been disabled.

In the event of a disaster at the primary site, the cycle update process is suspended and updating of the S-VOL stops. If the host requests an S-VOL takeover (CCI horctakeover), the remote array restores the S-VOL using data in the data pool from the previous cycle.

The AMS version of TCE does not support mirroring consistency of S-VOL data, even if the local array and remote path are functional. P-VOL and S-VOL data are therefore not identical when takeover is executed. Any P-VOL data updates made during the time the takeover command was issued cannot be salvaged.

Takeover processing

S-VOL takeover is performed when the horctakeover operation is issued by the secondary array. The TCE pair is split and system operation can be continued with the S-VOL only. In order to settle the S-VOL data being copied cyclically, it is restored using the data that was pre-determined in the preceding cycle and saved to the data pool, as mentioned above. The S-VOL is immediately enabled to receive the I/O instruction.

When the SVOL_Takeover is executed, data restoration processing from the data pool of the secondary site to the S-VOL is performed in the background. During the period from the execution of the SVOL_Takeover until the completion of the data restoration processing, performance of the host I/O for the S-VOL is deteriorated. P-VOL and S-VOL data are not the same after this operation is performed.

For details on the horctakeover command, see *Hitachi AMS Command Control Interface (CCI) Reference Guide* (MK-97DF8121).

Swapping P-VOL and S-VOL

SWAP Takeover ensures that system operation continues by reversing the characteristics of the P-VOL and the S-VOL and swapping the relationship between the P-VOL and S-VOL. After S-VOL takeover, host operations continue on the S-VOL, and S-VOL data becomes updated as a result of I/O operations. When continuing application processing using the S-VOL or when restoring application processing to the P-VOL, the swap function makes the P-VOL up-to-date, by reflecting updated data on the S-VOL to the P-VOL.

Failback to the local array

The fallback process involves restarting business operations at the local site. The following shows the procedure after the pair swap is performed.

1. On the remote array, after S-VOL takeover and the TCE pair swap command are executed, the S-VOL is mounted, and data restoration is executed (fsck for UNIX and chkdsk for Windows).
2. I/O is restarted using the S-VOL.
3. When the local site/array is restored, the TCE pair is created from the remote array. At this time, the S-VOL is located on the local array.
4. After the initial copy is completed and status is Paired, I/O to the remote TCE volume is stopped and it is unmounted.
5. The TCE pair is split. This completes transfer of data from the remote volume to the local volume.
6. At the local site, the pair swap command is issued. When this is completed, the S-VOL in the local array becomes the P-VOL.
7. Mount the new P-VOL on the local array is mounted and I/O is restarted.

Monitoring and maintenance

This chapter provides information and instructions for monitoring and maintaining the TCE system.

- ❑ [Monitoring pair status](#)
- ❑ [Monitoring data pool capacity](#)
- ❑ [Monitoring the remote path](#)
- ❑ [Monitoring cycle time](#)
- ❑ [Changing copy pace](#)
- ❑ [Checking RPO — Monitoring P-VOL/S-VOL time difference](#)
- ❑ [Routine maintenance](#)

Monitoring pair status

Pair status should be checked periodically to insure that TCE pairs are operating correctly. If the pair status becomes Failure or Pool Full, data cannot be copied from the local array to the remote array.

Also, status should be checked before performing a TCE operation. Specific operations require specific pair statuses.

Monitoring using the GUI is done at the user's discretion. Monitoring should be performed frequently. Email notifications can be set up to inform you when failure and other events occur.

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 9-1](#).

| Pairs | | | | | | | Rows/Page |
|--------------------------|-------|------------------------|-----------|-----------------|------------|--------------|-----------|
| <input type="checkbox"/> | Name | Local LUN [△] | Attribute | Remote Array ID | Remote LUN | Status | |
| <input type="checkbox"/> | Pair1 | 0003 | Primary | 85010031 | 0001 | Paired(100%) | |
| <input type="checkbox"/> | Pair2 | 0004 | Primary | 85010031 | 0003 | Split(100%) | |
| <input type="checkbox"/> | Pair3 | 0005 | Primary | 85010031 | 0004 | Paired(100%) | |

Figure 9-1: Pairs Screen

3. Locate the pair whose status you want to review in the **Pair** list. Status descriptions are provided in [Table 9-1](#). 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 9-1: Pair Status Definitions

| Pair Status | Description | Access to P-VOL | Access to S-VOL |
|---------------|--|----------------------|----------------------|
| Simplex | TCE pair is not created. Simplex volumes accept Read/Write I/Os. Not displayed in TCE pair list on Navigator 2 GUI. | P-VOL does not exist | S-VOL does not exist |
| Synchronizing | Copying is in progress, initiated by Create Pair or Resynchronize Pair operations. Upon completion, pair status changes to Paired. Data written to the P-VOL during copying is transferred as differential data after the copying operation is completed. Copy progress is shown on the Pairs screen in the Navigator 2 GUI. | Read/Write | Read Only |

Table 9-1: Pair Status Definitions (Continued)

| Pair Status | Description | Access to P-VOL | Access to S-VOL |
|---------------|--|-----------------|-------------------------------------|
| Paired | Copying of data from P-VOL to S-VOL is completed. While the pair is in Paired status, data consistency in the S-VOL is guaranteed. Data written to the P-VOL is transferred periodically to the S-VOL as differential data. | Read/Write | Read Only |
| Paired:split | When a pair-split operation is initiated, the differential data accumulated in the local array is updated to the S-VOL before the status changes to Split. Paired:split is a transitional status between Paired and Split. | Read/Write | Read Only |
| Paired:delete | When a pair-delete operation is initiated, the differential data accumulated in the local array is updated to the S-VOL before the status changes to Simplex. Paired:delete is a transitional status between Paired and Simplex. | Read/Write | Read Only |
| Split | Updates to the S-VOL are suspended; S-VOL data is consistent and usable by an application for read/write. Data written to the P-VOL and to the S-VOL are managed as differential data in the local and remote arrays. | Read/Write | Read/Write (mountable) or Read only |
| Pool Full | <p>If the local data pool capacity exceeds 90%, while status is Paired, the following takes place:</p> <ul style="list-style-type: none"> • Pair status on the local array changes to Pool Full. • Pair status on the remote array at this time remains Paired. • Data updates stop from the P-VOL to S-VOL. • Data written to the P-VOL is managed as differential data <p>If the remote data pool capacity reaches 100% while the pair status is Paired, the following takes place:</p> <ul style="list-style-type: none"> • Pair status on the remote array changes to Pool Full. • Pair status on the local array changes to Failure. <p>If a pair in a group becomes Pool Full, the status all pairs in the group becomes Pool Full.</p> <p>To recover, add LUs to the data pool or reduce use of the data pool. Then resynchronize the pair.</p> | Read/Write | Read Only |
| Takeover | Takeover is a transitional status after Swap Pair is initiated. The data in the remote data pool, which is in a consistent state established at the end of the previous cycle, is restored to the S-VOL. Immediately after the pair becomes Takeover, the pair relationship is swapped and copy from the new P-VOL to the new S-VOL is started. | | Read/Write |
| Busy | Busy is a transitional status after Swap Pair is initiated. Takeover occurs after Busy. This status can be seen from the Navigator 2 GUI, though not from CCI. | | Read/Write |
| Inconsistent | This status on the remote array occurs when copying from P-VOL to S-VOL stops due to failure in the S-VOL. The failure includes failure of the HDD that constitutes the S-VOL, or the data pool for the S-VOL becomes full. To recover, resynchronize the pair, which leads to a full volume copy of the P-VOL to the S-VOL. | | No Read/Write |

Table 9-1: Pair Status Definitions (Continued)

| Pair Status | Description | Access to P-VOL | Access to S-VOL |
|-------------|---|-----------------|-----------------|
| Failure | <p>P-VOL pair status changes to Failure if copying from the P-VOL to the S-VOL can no longer continue. The failure includes HDD failure and remote path failure that disconnects the local array and the remote array.</p> <ul style="list-style-type: none"> • Data consistency is guaranteed in the group if the pair status at the local array changes from Paired to Failure. • Data consistency is not guaranteed if pair status changes from Synchronizing to Failure. • Data written to the P-VOL is managed as differential data. <p>To recover, remove the cause then resynchronize the pair.</p> | Read/Write | |

Monitoring data pool capacity

Monitoring data pool capacity is critical for the following reasons:

- Data copying from the local to remote array will halt when:
 - The local data pool's use rate reaches 90 percent
 - The remote data pool's capacity is full

Also, the local array could be damaged if data copying is stopped for these reasons.

This section provides instructions for

- Monitoring data pool usage
- Specifying the threshold value
- Adding capacity to the data pool

Monitoring data pool usage

To monitor data pool usage level

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. The Setup screen displays.
3. Select **Data Pools**. The Data Pools screen displays.
4. Locate the desired data pool and review the **% Used** column. This shows the percentage of the data pool's capacity that is being used. Click the **Refresh Information** button to make sure the data is current.

If the status becomes Threshold Over, then capacity is more than the Threshold level. Threshold is set by the user during data pool setup. It is a percentage of the data pool that, when reached, indicates to you that maximum capacity is close to being reached. The default Threshold level is 70%.

If usage reaches the Threshold level or is close to it on a regular basis, the data pool should be expanded. If SnapShot is being used on the array, it may be necessary to reduce V-VOL number or lifespan.

Expanding data pool capacity

When monitoring indicates that the data pool is frequently at or above the Threshold value, you can add new volumes to expand its size.

To expand data pool capacity

The Storage system allows a maximum of 128 volumes for data pools. One data pool may consist of up to 64 volumes.

1. Create the logical unit or units that will be assigned to the data pool. See online Help for more information.
2. In the Replication tree, select **Setup**. The Setup screen appears.
3. Select **Data Pools**. The Data Pool list appears.
4. Select the data pool to be expanded.
5. Click **Edit Data Pool**. The Edit Data Pool screen appears.
6. Select the LU to be added to the data pool.
7. Click **OK**.
8. When the confirmation message appears, click **Close**

Changing data pool threshold value

The Threshold value protects the data pool from overflowing by issuing a warning when the use rate reaches "Threshold Value + 1%". At this point, data pool status is changed from Normal to Pool Full. When this occurs, it is recommended that you increase data pool capacity or decrease the number of pairs using the data pool.

The Threshold value is specified during data pool setup. You can change the value using the following procedure.

To change the Threshold value

1. In the Replication tree, select **Setup**. The Setup screen appears.
2. Select **Data Pools**. The Data Pool list appears.
3. Select the data pool whose threshold you want to change, then click **Edit Data Pool**. The Edit Data Pool screen appears.
4. In the Threshold box, enter the new value.
5. Click **OK**.
6. When the confirmation message appears, click **Close**.

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 and cycle time 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, using the buttons at the bottom of the screen.

Changing remote path bandwidth

Increase the amount of bandwidth allocated to the remote path when data copying is slower than write-workload. Bandwidth that is slow results in untransferred data accumulating in the data pool. This in turn can result in a full data pool, which causes pair failure.

To change bandwidth

1. In the Replication tree, click **Setup**, then click **Remote Path**. The Remote Path screen displays.
2. Click the check box for the remote path, then click **Change Bandwidth**. The Change Bandwidth screen displays.
3. Enter a new bandwidth.
4. Click **OK**.
5. When the confirmation screen appears, click **Close**.

Monitoring cycle time

Cycle time is the interval between updates from the P-VOL to the S-VOL. Cycle time is set to the default of 300 seconds during pair creation.

Cycle time can range between 30-seconds to 3600 seconds. When consistency groups are used, the minimum cycle time increases. For one group the minimum cycle time is 30 seconds, for two groups minimum cycle time is 60 seconds, and so on, up to 16 groups with a minimum of 8 minutes.

Updated data is copied to the S-VOL at the cycle time intervals. Be aware that this does not guarantee that all differential data can be sent within the cycle time. If the inflow to the P-VOL increases and the differential data to be copied is larger than bandwidth and the update cycle allow, then the cycle expands until all the data is copied.

When the inflow to the P-VOL decreases, the cycle time normalizes again. If you suspect that the cycle time should be modified to improve efficiency, you can reset it.

You learn of cycle time problems through monitoring. Monitoring cycle time can be done by checking group status, using CLI. See [Confirming Consistency Group \(CTG\) status on page A-19](#) for details.

Changing cycle time

To change cycle time

1. In the Replication tree, click **Setup**, then click **Options**. The Options screen appears.
2. Click **Edit Options**. The Edit Options screen appears.
3. Enter the new **Cycle Time** in seconds. The limits are 30 seconds to 3600 seconds.
4. Click **OK**.
5. When the confirmation screen appears, click **Close**.

Changing copy pace

Copy pace is the rate that data is copied during pair creation, resynchronization, and updating. The pace can be slow, medium, and fast. The time that it takes to complete copying depends on pace, the amount of data to be copied, and bandwidth.

To change copy pace

1. Connect to the local array and select the **Remote Replication** icon in the Replication tree view.
2. Select a pair from the pair list.
3. Click **Edit Pair**. The Edit Pair screen appears.
4. Select a **copy pace** from the dropdown list.
 - Slow — Copying takes longer when host I/O activity is high. The time to complete copying may be lengthy.
 - Medium — (Recommended) Copying is performed continuously, though it does not have priority; the time to completion is not guaranteed.
 - Fast — Copying is performed continuously and has priority. Host I/O performance is degraded. Copying time is guaranteed.
5. Click **OK**.
6. When the confirmation message appears, click **Close**.

Checking RPO — Monitoring P-VOL/S-VOL time difference

You can determine whether your desired RPO is being met by the TCE system. Do this by monitoring the time difference between P-VOL data and S-VOL data. See the section on synchronous waiting command (pairsyncwait) in the *Hitachi AMS Command Control Interface (CCI) Reference Guide* for details.

Routine maintenance

You may want to delete a volume pair, data pool, DMLU, or remote path. The following sections provide prerequisites and procedures.

Deleting a volume pair

When a pair is deleted, the P-VOL and S-VOL change to Simplex status and the pair is no longer displayed in the GUI Remote Replication pair list.

Please review the following before deleting a pair:

- When a pair is deleted, the primary and secondary volumes return to the Simplex state after the differential data accumulated in the local array is updated to the S-VOL. Both are available for use in another pair. Pair status is *Paired:delete* while differential data is transferred.
- If failure occurs when the pair is *Paired:delete*, the data transfer is terminated and the pair becomes *Failure*. While pair status changes to *Failure*, it cannot be resynchronized.
- Deleting a pair whose status is *Synchronizing* causes the status to become *Simplex* immediately. In this case, data consistency is not guaranteed.
- A Delete Pair operation can result in the pair deleted in the local array but not in the remote array. This can occur when there is a remote path failure or the pair status on the remote array is *Busy*. In this instance, wait for the pair status on the remote array to become *Takeover*, then delete it.
- 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. Therefore, when the P-VOL tries to send differential data to the S-VOL, it sees that the S-VOL is absent, and P-VOL pair status changes to *Failure*.
 - When a pair's status changes to *Failure*, the status of the other pairs in the group also becomes *Failure*.
- After an *SVOL_Takeover* command is issued, the pair cannot be deleted until S-VOL data is restored from the remote data pool.

To delete a TCE pair

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

Deleting data pools

Prerequisites

Before a data pool is deleted, the TCE pairs associated with it must be deleted.

To delete a data pool

1. Select the **Data Pools** icon under Setup in the tree view.
2. Select the data pool you want to delete in the **Data Pool** list.
3. Click **Delete Data Pool**.
4. A message appears. Click **Close**.

Deleting a DMLU

When TCE is enabled on the array and only one DMLU exists, it cannot be deleted. If two DMLUs exist, one can be deleted.

To delete a DMLU

1. In the Replication tree view, select **Setup** and then **DMLU**. The Differential Management Logical Units list appears.
1. Select the **LUN** you want to remove.
2. Click the **Remove DMLU** button. A success message displays.
3. Click **Close**.

Deleting the remote path

Delete the remote path from the local array.

Prerequisites

- Pairs must be in Split or Simplex status.

To delete the remote path

1. In the Storage Navigator 2 GUI, select the **Setup** icon in the Replication tree view, then select **Remote Path**.
2. On the Remote Path screen, click the box for the path that is to be deleted.
3. Click the **Delete Path** button.
4. Click **Close** on the Delete Remote Path screen.

TCE tasks before a planned remote array shutdown

Before shutting down the remote array, do the following:

- Split all TCE pairs. If you perform the shutdown without splitting the pairs, the P-VOL status changes to Failure. In this case, re-synchronize the pair after restarting the remote array.
- Delete the remote path (from local array).

TCE tasks before updating firmware

Before and after updating an array's firmware, perform the following TCE operations:

- TCE pairs must be split before updating the array firmware.
- After the firmware is updated, resynchronize TCE pairs.

Troubleshooting

This chapter provides information and instructions for troubleshooting the TCE system.

- ❑ [Troubleshooting overview](#)
- ❑ [Correcting data pool shortage](#)
- ❑ [Correcting array problems](#)
- ❑ [Correcting resynchronization errors](#)
- ❑ [Using the event log](#)
- ❑ [Miscellaneous troubleshooting](#)

Troubleshooting overview

TCE stops operating when any of the following occur:

- Pair status changes to **Failure**
- Pair status changes to **Pool Full**
- Remote path status changes to **Detached**

The following steps can help track down the cause of the problem and take corrective action.

1. Check the Event Log, which may indicate the cause of the failure. See [Using the event log on page 10-6](#).
2. Check pair status.
 - a. If pair status is Pool Full, please continue with instructions in [Correcting data pool shortage on page 10-2](#).
 - b. If pair status is Failure, check the following:
 - Check the status of the local and remote arrays. If there is a Warning, please continue with instructions in [Correcting array problems on page 10-3](#).
 - Check pair operation procedures. Resynchronize the pairs. If a problem occurs during resynchronization, please continue with instructions in [Correcting resynchronization errors on page 10-4](#).
3. Check remote path status. If status is Detached, please continue with instructions in [Correcting array problems on page 10-3](#).

Correcting data pool shortage

TCE pair are automatically split the data pool is full. The data pool becomes full when:

- The amount of I/O to the P-VOL (inflow) is chronically more than the data transferred to the S-VOL (outflow).
- The controller of the primary or secondary array is continuously overloaded. This causes data transfer to slow down.
- A remote path or the controller is switched, delaying data transfer continually, causing the pair to be placed in Pool Full status.

Data loss becomes greater than the RPO if a failure occurs under these circumstances.

TCE data copying stops and pair status changes to Pool Full when:

- The local array's data pool capacity exceeds 90 percent
- The remote array's data pool is full

In addition, when the local array's data pool capacity exceeds 90 percent, the data in the pool is deleted. Also, if SnapShot is using the local or remote array, the data is deleted when capacity is full.

To recover pairs in Pool Full status

1. Confirm the usage rates for local and remote array data pools.

- To check the local data pool, see [Monitoring data pool capacity on page 9-4](#).
 - To check the remote data pool, review the event log. (Shortages in the remote data pool prompt the array to resynchronize SnapShot pairs if they exist—thus reducing pool usage. The GUI may not immediately show the lowered rate.) Refer to [Using the event log on page 10-6](#).
2. If both local and remote data pools have sufficient space, resynchronize the pairs.
 3. To correct a data pool shortage, proceed as follows:
 - a. If there is enough disk space on the array, create and assign more LUs to the data pool. See [Expanding data pool capacity on page 9-5](#).
 - b. If LUs cannot be added to the data pool, review the importance of your TCE pairs. Delete the pairs not vital to business operations. See [Deleting a volume pair on page 9-8](#).
 - c. When corrective steps have been taken, resynchronize the pairs.
 4. For SnapShot pair recovery, review the troubleshooting chapter in *Hitachi AMS Copy-On-Write SnapShot User Guide* (MK-97DF8124).

Correcting array problems

A problem or failure in an array or remote network path can cause pairs to stop copying. Take the following action to correct array problems.

1. Review the information log to see what the hardware failure is.
2. Restore the array. Drive failures must be corrected by Hitachi maintenance personnel.
3. When the system is restored, recover TCE pairs.

For a detached remote path, parts should be replaced, then the remote path setup again.

For drive multiple failures (shown in [Table 10-1](#)) the pairs most likely need to be deleted and recreated.

Table 10-1: Array Failures Affecting TCE

| Location of Failure | Probable Result |
|---------------------------|--|
| P-VOL | Data not copied to the S-VOL may have been lost. |
| S-VOL | Remote copy cannot be continued because the S-VOL cannot be updated. |
| Data pool in local array | Remote copy cannot be continued because differential data is not available. |
| Data pool in remote array | Takeover to the S-VOL cannot be executed because internally pre-determined S-VOL data is lost. |

Delays in settling of S-VOL Data

When the amount of data that flows into the primary array from the host is larger than outflow from the secondary array, more time is required to complete the settling of the S-VOL data, because the amount of data to be transferred increases.

When the settlement of the S-VOL data is delayed, the amount of the data loss increases if a failure in the primary array occurs.

Differential data in the primary array increases a when:

- The load on the controller is heavy
- An initial or resynchronization copy is made
- SATA drives are used
- The path or controller is switched

Correcting resynchronization errors

When a failure occurs after a resynchronization has started, an error message cannot be displayed. In this case, you can check for the error detail code in the Event Log. [Figure 10-1](#) shows an example of the detail code.

The error message for pair resynchronizing is “The change of the remote pair status failed”.

```

12/18/2007 11:32:11 CO IB1900 Remote copy failed(CTG-14)
12/18/2007 11:32:11 CO IB1G00 Pair status changed by the error(CTG-14)
12/18/2007 11:32:11 CO IB1M00 The change of the remote pair status failed(LU-0005/0004, code-030A)
12/18/2007 11:28:22 CO IB1D00 Remote copy completed(CTG-14,0001/0001)
12/18/2007 11:28:15 CO IB1R00 Remote copy started(LU-0005/0004,CTG-14)
12/18/2007 11:17:25 CO I1G500 Quick Format completed(LU-0005)
  
```

Annotations in the image:

- Time when an error occurred: 12/18/2007 11:32:11
- Controller number: CO
- LU number in which error occurred: LU-0005/0004
- Error code: code-030A

Figure 10-1: Detail Code Example for Failure during Resync

[Table 10-2](#) lists error codes that can occur during a pair resync and the actions you can take to make corrections.

Table 10-2: Error Codes for Failure during Resync

| Error Code | Error Contents | Actions to be Taken |
|------------|---|---|
| 0307 | The array ID of the remote array cannot be specified. | Check the serial number of the remote array. |
| 0308 | The LU assigned to a TCE pair cannot be specified. | The resynchronization cannot be performed. Create a pair again after deleting the pair. |
| 0309 | Restoration from the Data Pool is in progress. | Retry after waiting for a while. |

Table 10-2: Error Codes for Failure during Resync (Continued)

| Error Code | Error Contents | Actions to be Taken |
|------------|---|--|
| 030A | The target S-VOL of TCE is a P-VOL of SnapShot. Besides, the SnapShot pair is being restored or reading/writing is not allowed. | When the SnapShot pair is being restored, execute it after the restoration is completed. When reading/writing is not allowed, execute it after enabling the reading/writing. |
| 030C | The TCE pair cannot be specified in the CTG. | The resynchronization cannot be performed. Create a pair again after deleting the pair. |
| 0310 | The status of the TCE pair is Takeover. | |
| 0311 | The status of the TCE pair is Simplex. | |
| 031F | The LU of the S-VOL of the TCE is S-VOL Disable. | Check the LU status of in the remote array, release the S-VOL Disable, and execute it again. |
| 0320 | The target LU in the remote array is undergoing the parity correction. | Retry after waiting for a while. |
| 0321 | The status of the target LU in the remote array is other than normal or regressed. | Execute it again after restoring the target LU status. |
| 0322 | The number of unused bits is insufficient. | Retry after waiting for a while. |
| 0323 | The LU status of the Data Pool is other than normal or regressed. | Execute it again after restoring the LU status of the Data Pool. |
| 0324 | The LU of the Data Pool is undergoing the parity correction. | Retry after waiting for a while. |
| 0325 | The expiration date of the temporary key is expired. | The resynchronization cannot be performed because the trial time limit is expired. Purchase the permanent key. |
| 0326 | The disk drives that configure a RAID group, to which a target LU in the remote array belongs have been spun down. | Perform the operation again after spinning up the disk drives that configure the RAID group. |
| 032D | The status of the RAID group that includes the S-VOL is not Normal. | Perform the same operation after the status becomes Normal. |

Using the event log

Using the event log helps in locating the reasons for a problem. The event log can be displayed using Navigator 2 GUI or CLI.

To display the Event Log using the GUI

1. Select the Alerts & Events icon. The Alerts & Events screen appears.
2. Click the Event Log tab. The Event Log displays as shown in [Figure 10-2](#).

Collect Trace Refre

Alerts & Events

DF800M_85010053 > Alerts & Events

| Summary | | | |
|--------------|------------|-------------------|----------|
| Array Status | Warning | Serial No. | 85010053 |
| Type | DF800M | Firmware Revision | 0840/A-M |
| Build Date | 2007/12/11 | | |

Alert Parts **Event Log**

Rows/Page: 25 | Page 1

Messages

| |
|---|
| 05/06/2008 17:50:45 00 IS0000 Remote copy failed(LU-0001/0001) :MANUAL/STRC |
| 05/06/2008 17:49:34 00 IB1D00 Remote copy completed(CTG-00,0001/0002) |
| 05/06/2008 17:49:12 10 IB1B00 Remote copy started(LU-0004/0004,CTG-00) |
| 05/06/2008 17:49:11 00 IB1B00 Remote copy started(LU-0001/0001,CTG-00) |
| 05/06/2008 17:48:26 10 I1G500 Quick Format completed(LU-0006) |
| 05/06/2008 17:48:20 10 I1G400 Quick Format started(LU-0006) |
| 05/06/2008 17:47:46 00 I1G500 Quick Format completed(LU-0005) |
| 05/06/2008 17:47:45 00 IS0000 Remote copy failed(LU-0001/0001) :MANUAL/STRC |
| 05/06/2008 17:47:39 00 I1G400 Quick Format started(LU-0005) |
| 05/06/2008 17:45:50 00 IB1B00 Remote copy started(LU-0001/0001,CTG-00) |
| 05/06/2008 17:45:46 00 I12100 LU deleted(LU-0006) |
| 05/06/2008 17:45:46 00 I12100 LU deleted(LU-0005) |

Figure 10-2: Event Log in Navigator 2 GUI

Event Log messages show the time when an error occurred, the message, and an error detail code, as shown in [Figure 10-3](#). If the data pool is full, the error message is "I6D000 Data pool does not have free space (Data pool-xx)", where xx is the data pool number.

Time when an error occurred

```
05/18/2007 11:32:11 00 I6D000 Data pool does not have free space (Data pool-52)
```

Data pool number

Figure 10-3: Detail Code Example for Data Pool Error

Miscellaneous troubleshooting

Table 10-3 contains details on pair and takeover operations that may help when troubleshooting. Review these restrictions to see if they apply to your problem.

Table 10-3: Miscellaneous Troubleshooting

| Restriction | Description |
|---|---|
| Restrictions for pair splitting | When a pair split operation is begun, data is first copied from the P-VOL to the S-VOL. This causes a time delay before the status of the pair becomes Split. |
| | The splitting of the TCE pair cannot be done when the pairsplit -mscas processing is being executed for the CTG. |
| | When a command to split pairs in each CTG is issued while the pairsplit -mscas processing is being executed for the cascaded SnapShot pair, the splitting cannot be executed for all the pairs in the CTG. |
| | When a command to split each pair is issued and the target pair is under the completion processing, it cannot be accepted if the Paired to be split is undergoing the end operation. |
| | When a command to split each pair is issued and the target pair is under the completion processing, it cannot be accepted if the Paired to be split is undergoing the splitting operation. |
| | When a command to split pairs in each group is issued, it cannot be executed if even a single pair that is being split exists in the CTG concerned. |
| | When a command to terminate pairs in each group is issued, it cannot be executed if even a single pair that is being split exists in the CTG concerned. |
| | The pairsplit -P command is not supported. |
| Restrictions on execution of the horctakeover (SVOL_Takeover) command | When the SVOL_Takeover operation is performed for a pair by the horctakeover command, the S-VOL is first restored from the data pool. This causes a time delay before the status of the pair changes. |
| | The restoration of up to four LUs can be done in parallel for each controller. When restoration of four or more LUs is required, the first four LUs are selected according to an order given in the requirement, but the following LUs are selected in ascending order of the LU numbers. |
| | Because the SVOL_Takeover operation is performed on the secondary side only, the differential data of the P-VOL that has not been transferred is not reflected on the S-VOL data even when the TCE pair is operating normally. |
| | When the S-VOL of the pair, to which the instruction to perform the SVOL_Takeover operation is issued, is in the Inconsistent status that does not allow Read/Write operation, the SVOL_Takeover operation cannot be executed. Whether the Split is Inconsistent or not can be referred to using Navigator 2. |
| | When the command specifies the target as a group, it cannot be executed for all the pairs in the CTG if even a single pair in the Inconsistent status exists in the CTG. |
| | When the command specifies the target as a pair, it cannot be executed if the target pair is in the Simplex or Synchronizing status. |

Table 10-3: Miscellaneous Troubleshooting (Continued)

| Restriction | Description |
|---|---|
| Restrictions on execution of the pairsplit -macas command | The pair splitting instruction cannot be issued to the SnapShot pair cascaded with the TCE S-VOL pair in the Synchronizing or Paired status from the host on the secondary side. |
| | When even a single pair in the CTG is being split or deleted, the command cannot be executed. |
| | Pairsplit -mascas processing is continued unless it becomes Failure or Pool Full. |
| Restrictions on the performance of pair delete operation | When a delete pair operation is begun, data is first copied from the P-VOL to the S-VOL. This causes a time delay before the status of the pair changes. |
| | The end processing is continued unless it becomes Failure or Pool Full. |
| | A pair cannot be deleted it is being split. When a delete pair command is issued to a group, it will not be executed if any of the pairs in the group is being split. |
| | A pair cannot be deleted when the pairsplit -mascas command is being executed. This applies singly or by the CTG. |
| | When a delete pair command is issued to a group, it will not be executed if any of the pairs in the group is undergoing the pair split -mascas operation. |
| | Also in the execution of the pairsplit -R command that requires the secondary array to delete a pair, the differential data of the P-VOL that has not been transferred is not reflected on the S-VOL data in the same way as the case of the SVOL_Takeover operation. |
| | The pairsplit -R command cannot be executed during the restoration of the S-VOL data through the SVOL_Takeover operation. |
| Restrictions while using load balancing | The load balancing function is not applied to the LUs specified as a TCE pair. Since the ownership of the LUs specified as a TCE pair is the same as the ownership of the LUs specified as a data pool, perform the setting so that the ownership of LUs specified as a data pool is balanced in advance. |

Operations using CLI

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

- [Installation and setup](#)
- [Pair operations](#)
- [Procedures for failure recovery](#)
- [Sample script](#)



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

Installation and setup

The following sections provide installation/uninstalling, enabling/disabling, and setup procedures using CLI.

TCE 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>. See [Installation procedures on page 6-2](#) for prerequisites information.

Before installing or uninstalling TCE, verify the following:

- The array must be operating in a normal state. Installation and un-installation cannot be performed if a failure has occurred.
- Make sure that a spin-down operation is not in progress when installing or uninstalling TCE.
- The array may require a restart at the end of the installation procedure. If SnapShot is already enabled, no restart is necessary. If restart is required, it can be done when prompted, or at a later time.
- TCE cannot be installed if more than 239 hosts are connected to a port on the array.

In cases where the DKN-200-NGW1 (NAS unit) is connected to the disk array, check the following items in advance.

1. Prior to this operation, execute Correspondence when connecting the NAS unit if each of the following three items apply to the disk array.
 - NAS unit is connected to the disk array. Ask the disk array administrator to confirm whether the NAS unit is connected or not.
 - NAS unit is in operation. Ask the NAS unit administrator to confirm whether the NAS service is operating or not.
 - A failure has not occurred on the NAS unit. Ask the NAS unit administrator to check whether failure has occurred or not by checking with the NAS administration software, NAS Manager GUI, List of RAS Information, etc. In case of failure, execute the maintenance operation together with the NAS maintenance personnel.
2. Correspondence when connecting the NAS unit.
 - If the NAS unit is connected, ask the NAS unit administrator for termination of NAS OS and planned shutdown of the NAS unit.
3. Points to be checked after completing this operation:
 - Ask the NAS unit administrator to reboot the NAS unit. After rebooting, ask the NAS unit administrator to refer to "Recovering from FC path errors" in the *Hitachi NAS Manager User's Guide* and check the status of the Fibre Channel path (FC path in short) and to recover the FC path if it is in a failure status.
 - In addition, if there are any personnel for the NAS unit maintenance, ask the NAS unit maintenance personnel to reboot the NAS unit.

Installing

To install TCE

1. From the command prompt, register the array in which the TCE is to be installed, and then connect to the array.
2. Execute the `auopt` command to install TCE. 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.
In order to complete the setting, it is necessary to reboot the subsystem.
Host will be unable to access the subsystem while restarting. Host
applications that use the subsystem will terminate abnormally. Please stop
host access before you restart the subsystem.
Also, if you are logging in, the login status will be canceled when
restarting begins.
When using Remote Replication, restarting the remote subsystem will cause
both Remote Replication paths to fail.
Remote Replication pair status will be changed to "Failure(PSUE)" when pair
status is "Paired(PAIR)" or "Synchronizing(COPY)". Please change Remote
Replication pair status to "Split(PSUS)" before restart.
Do you agree with restarting? (y/n [n]): y
Are you sure you want to execute?
(y/n [n]): y
Now restarting the subsystem. Start Time hh:mm:ss Time Required 4 - 15min.
The subsystem restarted successfully.
%
```

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

```
% auopt -unit array-name -refer
Option Name      Type      Term      Status
TC-EXTENDED     Permanent ---      Enable
%
```

TCE is installed and enabled.

Enabling and disabling

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

Prerequisites for disabling

- TCE pairs must be released (the status of all LUs must be Simplex).
- The remote path must be released.
- Data pools must be deleted unless a SnapShot system exists on the array.
- Make sure a spin-down operation is not in progress when uninstalling.
- TCE cannot be enabled if more than 239 hosts are connected to a port on the array.

To enable/disable TCE

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 TCE 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 TC-EXTENDED -st disable
Are you sure you want to disable the option? (y/n[n]): y
The option has been set successfully.
In order to complete the setting, it is necessary to reboot the subsystem.
Host will be unable to access the subsystem while restarting. Host
applications that use the subsystem will terminate abnormally. Please stop
host access before you restart the subsystem.
Also, if you are logging in, the login status will be canceled when
restarting begins.
When using Remote Replication, restarting the remote subsystem will cause
both Remote Replication paths to fail.
Remote Replication pair status will be changed to "Failure(PSUE)" when pair
status is "Paired(PAIR)" or "Synchronizing(COPY)". Please change Remote
Replication
pair status to "Split(PSUS)" before restart.
Do you agree with restarting? (y/n [n]): y
Are you sure you want to execute?
(y/n [n]): y
Now restarting the subsystem. Start Time hh:mm:ss Time Required 4 - 15min.
The subsystem restarted successfully.
%
```

3. Execute the `auopt` command to confirm that the status has been changed. For example:

```
% auopt -unit array-name -refer
Option Name      Type      Term      Status
TC-EXTENDED     Permanent ---      Disable
%
```

Un-installing

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

Prerequisites for uninstalling

- TCE pairs must be released (the status of all LUs must be Simplex).
- The remote path must be released.
- Data pools must be deleted, unless a SnapShot system exists on the array.
- Make sure a spin-down operation is not in progress.

To uninstall TCE

1. From the command prompt, register the array in which the TCE is to be uninstalled, and then connect to the array.
2. Execute the `auopt` command to uninstall TCE. 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.
In order to complete the setting, it is necessary to reboot the subsystem.
Host will be unable to access the subsystem while restarting. Host
applications that use the subsystem will terminate abnormally. Please stop
host access before you restart the subsystem.
Also, if you are logging in, the login status will be canceled when restarting
begins.
When using Remote Replication, restarting the remote subsystem will cause
both Remote Replication paths to fail.
Remote Replication pair status will be changed to "Failure(PSUE)" when pair
status is "Paired(PAIR)" or "Synchronizing(COPY)". Please change Remote
Replication pair status to "Split(PSUS)" before restart.
Do you agree with restarting? (y/n [n]): y
Are you sure you want to execute?
(y/n [n]): y
Now restarting the subsystem. Start Time hh:mm:ss Time Required 4 - 15min.
The subsystem restarted successfully.
%

```

3. Execute the `auopt` command to confirm that TCE 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 TCE copies can be made. Please see the prerequisites under [Setting up DMLUs on page 6-5](#) 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 `auDM-LU` command. This command first displays LUs that can be assigned as DM-LUs and then creates a DM-LU. For example:

```

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

```

Release a DMLU

Observe the following when releasing a DMLU for TCE:

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

To release a TCE DMLU

Use the following example:

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

Setting the data pool

Create a data pool for storing differential data to be used by TCE.

- Up to 64 data pools can be designated for each array.
- A data pool should be a minimum of 20 GB.
- Logical units assigned to the data pool must be set up and formatted previously.
- Up to 64 logical units can be assigned to a data pool.
- The accurate capacity of a data pool cannot be determined immediately after a logical unit has been assigned. Data pool capacity can only be confirmed after about 3 minutes per 100 GB.
- An LU with a SAS/SSD drive and an LU with a SATA drive cannot coexist in a data pool. Set all configured LUs using the same drive type.
- When using SnapShot with Cache Partition Manager, the segment size of the LU belonging to a data pool must be the default size (16 kB) or less.

To set up the data pool

1. From the command prompt, register the array to which you want to create the Data Pool, then connect to the array.
2. Execute the `aupool` command create a Data Pool.

First, display the LUs to be assigned to a Data Pool, and then create a Data Pool.

The following is the example of specifying LU 100 for Data Pool 0.

```

% aupool -unit array-name -availablelist -poolno 0
Data Pool      : 0
Available Logical Units
  LUN  Capacity RAID Group RAID Level  Type Status
  100  30.0 GB           0  6( 9D+2P) SAS Normal
  200  35.0 GB           0  6( 9D+2P) SAS Normal
%
% aupool -unit array-name -add -poolno 0 -lu 100
Are you sure you want to add the logical unit(s) to the data
pool 0?
(y/n[n]): y
The logical unit has been successfully added.
%

```

3. Execute the `aupool` command to verify that the Data Pool has been created. Refer to the following example.

```

% aupool -unit array-name -refer -poolno 0
Data Pool      : 0
Data Pool Usage Rate: 6% (2.0/30.0 GB)
Threshold      : 70%
Usage Status   : Normal
  LUN  Capacity      RAID Group RAID Level  Type Status
  100  30.0 GB           0  6( 9D+2P) SAS Normal
%

```

4. When deleting the logical unit set as the Data Pool, it is necessary to delete all Snapshot images (V-VOLs). To delete an existing Data Pool, refer to the following example.

```

% aupool -unit array-name -rm -poolno 0
Are you sure you want to delete all logical units from the data
pool 0?
(y/n[n]): y
The logical units have been successfully deleted.
%

```

5. To change an existing threshold value for a Data Pool, refer to the following example.

```

% aupool -unit array-name -cng -poolno 0 -thres 70
Are you sure you want to change the threshold of usage rate in
the data pool?
(y/n[n]): y
The threshold of the data pool usage rate has been successfully
changed.
%

```

Setting the LU ownership

The load balancing function is not applied to the LUs specified as a TCE pair. Since the ownership of the LUs specified as a TCE pair is the same as the ownership of the LUs specified as a data pool, perform the setting so that the ownership of LUs specified as a data pool is balanced in advance.

Perform the following steps to set LU ownership by CLI:

1. From the command prompt, register the array to which you want to set the LU ownership, and then connect to the array.
2. Execute the `autuningluown` command to confirm an LU ownership.

```
% autuningluown -unit array-name -refer
  LU  CTL  Core  RAID Group  DP Pool  Cache Partition  Type
    0   0   X     0           0       N/A              0  SAS
    1   1   X     0           0       N/A              0  SAS
 1000   0   X     0           0       N/A              0  SAS
 1001   1   X     0           0       N/A              0  SAS
 2000   0   X     0           0       N/A              0  SAS
 2001   1   Y     0           0       N/A              0  SAS
%
```

The Core shows N/A: AMS2100/2300

3. Execute the `autuningluown` command to change the LU 2001 ownership.

```
% autuningluown -unit array-name -set -lu 2001 -ctl0 -coreX
Are you sure you want to set the LU ownership? (y/n [n]): y
The LU ownership has been set successfully.
%
```

4. Execute the `autuningluown` command to confirm an LU ownership.

```
% autuningluown -unit array-name -refer
  LU  CTL  Core  RAID Group  DP Pool  Cache Partition  Type
    0   0   X     0           0       N/A              0  SAS
    1   1   X     0           0       N/A              0  SAS
 1000   0   X     0           0       N/A              0  SAS
 1001   1   X     0           0       N/A              0  SAS
 2000   0   X     0           0       N/A              0  SAS
 2001   0   X     0           0       N/A              0  SAS
%
```

Setting the cycle time

Cycle time is the time between updates to the remote copy when the pair is in Paired status. The default is 300 seconds. You can set cycle time between 30 to 3600 seconds.

Copying may take a longer than the cycle time, depending on the amount of the differential data or low bandwidth.

To set the cycle time

1. From the command prompt, register the array to which you want to set the cycle time, and then connect to the array.
2. Execute the `autruecopyopt` command to confirm the existing cycle time. For example:

```
% autruecopyopt -unit array-name -refer
Cycle Time[sec.]      : 300
Cycle OVER report    : Disable
%
```

3. Execute the `autruecopyopt` command to set the cycle time. For example:

```
% autruecopyopt -unit array-name -set -cycletime 300
Are you sure you want to set the TrueCopy options? (y/n [n]): y
The TrueCopy options have been set successfully.
%
```

Setting mapping information

The following is the procedure for mapping information. *For iSCSI*, use the `autargetmap` command in place of `auhgmap`.

1. From the command prompt, register the array to which you want to set the mapping information, and then connect to the array.
2. Execute the `auhgmap` command to set the mapping information. The following example defines LU 0 in the array to be recognized as 6 by the host. The port is connected via host 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 mapping information has been set. For example:

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

Setting the remote port CHAP secret

iSCSI systems only. 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 6-8](#).

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:

```
% aurtmtpath -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 [Setting up the remote path on page 6-9](#) 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
  Array ID      : 85000026
  Distributed Mode : N/A

Path Information
  Interface Type      :
  Remote Array ID     :
  Remote Path Name    :

  Bandwidth [0.1 M] :
  iSCSI CHAP Secret  :

Path  Status      Local      Remote  Remote Port  TCP Port No. of
      0  Undefined  ---      ---      ---          Remote Port
      1  Undefined  ---      ---      ---          ---
%
```

iSCSI example:

```
% aurmtpath -unit array-name -refer
Initiator Information
Local Information
  Array ID      : 85000026
  Distributed Mode : N/A

Path Information
  Interface Type      : FC
  Remote Array ID     : 85000027
  Remote Path Name    : N/A
  Bandwidth [0.1 M] : 15
  iSCSI CHAP Secret  : N/A

Path  Status      Local      Remote  Remote Port  TCP Port No. of
      0  Undefined  ---      ---      ---          Remote Port
      1  Undefined  ---      ---      ---          ---

Target Information
  Local Array ID     :
%
```

- Execute the `aurmtpath` command to set the remote path.

Fibre channel example:v

```
% aurmtpath -unit array-name -set -remote 85000027 -band 15
      -path0 0A 0A -path1 1A 1B
Are you sure you want to set the remote path information?
(y/n[n]): y
The remote path information has been set successfully.
%
```

iSCSI example:

```
% aurmtpath -unit array-name -set -initiator -remote 85000027 -
secret disable
      -path0 0B -path0_addr 192.168.1.201 -band 100
      -path1 1B -path1_addr 192.168.1.209
Are you sure you want to set the remote path information?
(y/n[n]): y
The remote path information has been set successfully.
%
```

- 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
  Interface Type      : FC
  Remote Array ID     : 85000027
  Remote Path Name    : N/A
  Bandwidth [0.1 M]  : 15
  iSCSI CHAP Secret  : N/A

      Remote Port      TCP Port No. of
Path  Status          Local   Remote  IP Address  Remote
Port
  0   Normal          0A     0A   N/A         N/A
  1   Normal          1A     1B   N/A         N/A
%
```

iSCSI example:

```
% aurmtpath -unit array-name -refer
Initiator Information
  Local Information
    Array ID      : 85000026
    Distributed Mode : N/A

  Path Information
    Interface Type      : iSCSI
    Remote Array ID     : 85000027
    Remote Path Name    : N/A
    Bandwidth [0.1 M]  : 100
    iSCSI CHAP Secret  : Disable

  Path Status
    Path  Status      Local  Remote  IP Address  TCP Port No. of
    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     : 85000027

%
```

Deleting the remote path

When shutdown of the arrays is necessary, the remote path must be deleted first. The status of TCE 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 M]  :
  iSCSI CHAP Secret  :

Path      Status      Local      Remote      Remote Port      TCP Port No. of
0      Undefined    ---      ---      IP Address      Remote Port
1      Undefined    ---      ---      ---      ---
%

```

Pair operations

The following sections describe the CLI procedures and commands for performing TCE 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.
1. 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
TCE_LU0000_LU0000  0  P-VOL      0  Paired(100
%)
TrueCopy Extended Distance  0:
TCE_LU0001_LU0001  1  P-VOL      1  Paired(100
%)
TrueCopy Extended Distance  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.

- 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                : TCE_LU0000_LU0000
Local Information
  LUN                    : 0
  Attribute              : P-VOL
Remote Information
  Array ID               : 85000027
  Path Name              : N/A
  LUN                    : 0
Capacity                 : 50.0 GB
Status                   : Paired(100%)
Copy Type                : TrueCopy Extended Distance
Group Name               : 0:
Data Pool                : 0
Data Pool Usage Rate    : 0%
Consistency Time        : 2008/02/29 11:09:34
Difference Size          : 2.0 MB
Copy Pace                : ---
Fence Level              : N/A
Previous Cycle Time     : 504 sec.
%
```

Creating a pair

See prerequisite information under [Creating the initial copy on page 7-2](#) before proceeding.

To create a pair

- From the command prompt, register the local array in which you want to create pairs, and then connect to the array.
- 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 -tce -
pvol
Available Logical Units
  LUN Capacity RAID Group RAID Level Type Status
  2   50.0 GB          0    6( 9D+2P) SAS Normal
%
```

- 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 -tce -
pvol
Available Logical Units
  LUN Capacity RAID Group RAID Level Type Status
  2   50.0 GB          0    6( 9D+2P) SAS Normal
%
```

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

```
% aureplicationremote -unit local array-name -create -tce -pvol 2 -svol 2
                        -remote xxxxxxxx -gno 0 -remotepoolno 0
Are you sure you want to create pair "TCE_LU0002_LU0002"?
(y/n [n]): y
The pair has been created successfully.
%
```

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:

```
% aureplicationremote -unit local array-name -split -tce -pvol 2 -svol 2
                        -remote xxxxxxxx
Are you sure you want to split pair?
(y/n [n]): y
The pair has been split successfully.
%
```

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 re-synchronize the specified pair. For example:

```
% aureplicationremote -unit local array-name -resync -tce -pvol 2 -svol 2
                        -remote xxxxxxxx
Are you sure you want to re-synchronize pair?
(y/n [n]): y
The pair has been re-synchronized successfully.
%
```

Swapping a pair

Please review the Prerequisites in [Swapping pairs on page 7-7](#).

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 -tce -gno 1
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:

```
% aureplicationremote -unit local array-name -simplex -tce -locallun pvol
                        -pvol 2 -svol 2 -remote xxxxxxxx
Are you sure you want to release pair?
(y/n [n]): y
The pair has been released successfully.
%
```

Changing pair information

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

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

```
% aureplicationremote -unit local array-name -tce -chg -pace slow
                        -locallun pvol -pvol 2000 -svol 2002 -remote xxxxxxxx
Are you sure you want to change pair information?
(y/n [n]): y
The pair information has been changed successfully.
%
```

Monitoring pair status

To monitor pair status

1. From the command prompt, register the local array on which you want to monitor pair status, and then connect to the array.
2. Execute the `aureplicationmon -evwait` command. For example:

```
% aureplicationmon -unit local array-name -evwait -tce -st simplex -gno 0
                        -waitmode backup
Simplex Status Monitoring...
Status has been changed to Simplex.
%
```

Confirming Consistency Group (CTG) status

You can display information about a consistency group using the `aureplicationremote` command. The information is displayed in a list.

To display consistency group status

1. From the command prompt, register the local array on which you want to view consistency group status, and then connect to the array.
2. Execute the `aureplicationremote -unit unit_name -refer -groupinfo` command. For example:

Descriptions of the consistency group information that displays is shown in [Table A-1](#).

Table A-1: CTG Information

| Displayed Item | Contents |
|--|---|
| CTG No. | CTG number |
| Lapsed Time | The Lapsed Time after the current cycle is started is displayed (in hours, minutes, and seconds). |
| Remaining Difference Size | The size of the residual differential data to be transferred in the current cycle is displayed. The size of the differential data in the pair information shows a total size of the data that have not been transferred and thus remains in the local array, whereas the size of the remaining differential data does not include the size of the data to be transferred in the following cycle. Therefore, the size of the remaining differential data does not coincide with the total size of the differential data of the pairs included in the CTG. |
| Transfer Rate | The Transfer Rate of the current cycle is displayed (KB/s). During a period from the start of the cycle to the execution of the copy operation or a waiting period from completion of the copy operation to the start of the next cycle, "---" is displayed. While the Transfer Rate to be output is being calculated, "Calculating" is displayed. |
| Prediction Time of Transfer Completion | The predicted time when the data transfer is completed (Prediction Time of Transfer Completion) for each cycle of the CTG is displayed (in hours, minutes, and seconds). If the predicted time when the data transfer is completed (Prediction Time of Transfer Completion) cannot be calculated because it is maximized temporarily, "99:59:59" is displayed. During a waiting period from completion of the cyclic operation till the start of the next cycle, "Waiting" is displayed. While the predicted time to be output is being calculated, "Calculating" is displayed. |

Procedures for failure recovery

Displaying the event log

When a failure occurs, you can learn useful information from the event log. The contents of the event log include the time when an error occurred, an error message, and an error detail code.

To display the event log

1. Register the array to confirm the event log on the command prompt.
2. Execute the `auinformsg` command and confirm the event log. For example:

```
% auinformsg -unit array-name
Controller 0/1 Common
12/18/2007 11:32:11 C0 IB1900 Remote copy failed(CTG-00)
12/18/2007 11:32:11 C0 IB1G00 Pair status changed by the
error(CTG-00)
:
12/18/2007 16:41:03 00 I10000 Subsystem is ready

Controller 0
12/17/2007 18:31:48 00 RBE301 Flash program update end
12/17/2007 18:31:08 00 RBE300 Flash program update start

Controller 1
12/17/2007 18:32:37 10 RBE301 Flash program update end
12/17/2007 18:31:49 10 RBE300 Flash program update start
%
```

The event log was displayed. When searching the specified messages or error detail codes, store the output result in the file and use the search function of the text editor as shown below.

```
% auinformsg -unit array-name>infomsg.txt
%
```

Reconstructing the remote path

To reconstruct the remote path

1. Register the array to reconstruct the remote path on the command prompt.
2. Execute the `aurmtpath` command with the `-reconst` option and enable the remote path status. For example:

```
% aurmtpath -unit array-name -reconst -remote 85000027 -path0
Are you sure you want to reconstruct the remote path?
(y/n [n]): y
The reconstruction of remote path has been required.
Please check "Status" as -refer option.
%
```

Sample 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=TCE_LU0001_LU0002
REM Specify the directory path that is mount point of P-VOL and S-VOL
set MAINDIR=C:\main
set BACKUPDIR=C:\backup
REM Specify GUID of P-VOL and S-VOL
PVOL_GUID=xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxxxx
SVOL_GUID=yyyyyyyy-YYYY-YYYY-YYYY-YYYYYYYYYYYYY

REM Unmounting the S-VOL
pairdisplay -x umount %BACKUPDIR%
REM Re-synchronizing pair (Updating the backup data)
aureplicationremote -unit %UNITNAME% -tce -resync -pairname %P_NAME% -gno 0
aureplicationmon -unit %UNITNAME% -evwait -tce -pairname %P_NAME% -gno 0 -st
paired -pvol

REM Unmounting the P-VOL
pairdisplay -x umount %MAINDIR%
REM Splitting pair (Determine the backup data)
aureplicationremote -unit %UNITNAME% -tce -split -pairname %P_NAME% -gname
%G_NAME%
aureplicationmon -unit %UNITNAME% -evwait -tce -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 or Windows Server 2003/2008 is used, the CCI `mount` command is required when mounting or un-mounting a volume. The GUID, which is displayed by the Windows `mountvol` command, is needed as an argument when using the `mount` command. For more information, refer to the *Hitachi Adaptable Modular Storage Command Control Interface (CCI) Reference Guide*.

Operations using CCI

This appendix describes CCI procedures for setup and performing TCE operations.

- [Setup](#)
- [Pair operations](#)
- [Pair, group name differences in CCI and Navigator 2](#)

Setup

The following sections provide procedures for setting up CCI for TCE.

Setting the command device

The command device is used by CCI to conduct operations on the array.

- Logical units used as command devices must be recognized by the host.
- The command device must be 33 MB or greater.
- Assign multiple command devices to different RAID groups to avoid disabled CCI functionality in the event of drive failure.



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.

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

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

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

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, before performing this operation.
Are you sure you want to release the command devices? (y/n
[n]): y
The specified command device will be released.
Are you sure you want to execute? (y/n [n]): y
The command devices have been released successfully.
%

```

5. To change 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:

```

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

```

Defining the configuration definition file

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

A sample configuration definition file, HORCM_CONF, is included with the CCI software. It should be used as the basis for creating your configuration definition file(s). The system administrator should copy the sample file, set the necessary parameters in the copied file, and place the copied file in the proper directory. For 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 `horcmshutdown` 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 B-1](#) shows an example of the `horcm0.conf` file.

```

horcm0.conf - Notepad
File Edit Search Help

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

HORCM_CMD
#dev_name        dev_name      dev_name
\\.\PHYSICALDRIVE1

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

HORCM_INST
#dev_group       ip_address    service
UG01             XXXXXXXXX    5001

```

Figure B-1: Horcm0.conf Example

6. Save the configuration definition file and use the `horcmstart` 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 B-2](#) shows an example of the `horcm1.conf` file.

```

horcm1.conf - Notepad
File Edit Search Help

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

HORCM_CMD
#dev_name        dev_name      dev_name
\\.\PHYSICALDRIVE1

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

HORCM_INST
#dev_group       ip_address    service
UG01             XXXXXXXXX    5000

```

Figure B-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:\horm\etc>echo hd1-3 | .\inraid
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:\horm\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 CCI procedures for performing TCE pairs operations. In the examples provided, the group name defined in the configuration definition file is VG01.



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

Checking pair status

To check TCE pair status

1. Execute the `pairdisplay` command to display 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 PAIR
ASYNC ,85000175 2 -
vg01 oradb1(R) (CL1-B, 2, 2)85000175 2.S-VOL PAIR
ASYNC ,----- 1 -
```

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

CCI and Navigator 2 GUI pair statuses are described in [Table B-1](#).

Table B-1: Pair Status Descriptions

| 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 where copy is completed and update copy between pairs started. |
| PSUS/SSUS | Split | Update copy between pairs stopped by split. |
| PFUS | Pool Ful | Status that updating copy from the P-VOL to the S-VOL cannot continue due to too much use of the data pool. |
| SSWS | Takeover | Takeover |
| SSUS | Inconsistent | Status that updating copy from the P-VOL to the S-VOL cannot continue due to the S-VOL failure. |
| PSUE | Failure | Update copy between pairs stopped by failure occurrence. |

Creating a pair (paircreate)

To create a pair

1. Execute the `pairdisplay` command to verify that the status of the possible volumes to be copied is SMPL. The group name in the example is VG01.

```
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 [Changing copy pace on page 9-7](#) 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. For example:

```
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 PAIR Never ,85000175
2 -
VG01   oradb1(R)   (CL1-A , 1, 2 )85000175   2.S-VOL PAIR Never ,----
-      1 -
```

Splitting a pair (pairsplit)

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

To split a pair

1. Execute the `pairsplit` command to split the TCE pair in the PAIR status. The group name in the example is VG01.

```
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, Seq#,P-
LDEV# M
VG01   oradb1(L)  (CL1-A , 1, 1 )85000174   1.P-VOL PSUS ASYNC ,85000175
2 -
VG01   oradb1(R)  (CL1-A , 1, 2 )85000175   2.S-VOL SSUS ASYNC ,-----
1 -

```

Resynchronizing a pair (pairresync)

To resynchronize TCE 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 `pairevtwait` commands. The group name in the example is VG01.

```

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:

```

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 ASYNC ,85000175
2 -
VG01   oradb1(R)  (CL1-A , 1, 2 )85000175   2.S-VOL PAIR ASYNC ,-----
1 -

```

Suspending pairs (pairsplit -R)

To suspend pairs

1. Execute the `pairdisplay` command to verify that the pair to be suspended is in PAIR status. The group name in the example is VG01.

```

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 ASYNC ,85000175
2 -
VG01   oradb1(R)  (CL1-A , 1, 2 )85000175   2.S-VOL PAIR ASYNC ,-----
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 pair status changed to 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.P-VOL PSUE ASYNC ,85000175
2 -
VG01   oradb1(R)  (CL1-A , 1, 2 )85000175   2.S-VOL ----- ,-----
-----
```

Releasing pairs (pairsplit -S)

To release pairs and change status to SMPL

1. Execute the `pairsplit -S` command to release the pair. The group name in the example is VG01.

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

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

Splitting TCE S-VOL/SnapShot V-VOL pair (pairsplit -mscas)

The `pairsplit -mscas` command splits a SnapShot pair that is cascaded with an S-VOL of a TCE pair. The data to be split is the P-VOL data of the TCE pair at the time when the `pairsplit -mscas` command is accepted.

CCI adds a human-readable character string of ASCII 31 characters to a remote snapshot. Because a snapshot can be identified by a character string rather than an LU number, it can be used for discrimination of the SnapShot volumes of many generations.

Requirements

- Cascade configuration of TCE and SnapShot pairs is required.
- This command is issued to TCE; however, the pair to be split is the SnapShot pair cascaded with the TCE S-VOL.
- This command can only be issued for the TCE consistency group (CTG). It cannot be issued directly to a pair.
- The TCE pair must be in PAIR status; the SnapShot pair must be in either PSUS or PAIR status.
- When both TCE and SnapShot pairs are in PAIR status, any pair split command directly to the SnapShot pair, other than the `pairsplit` command with the `-mscas` option, cannot be executed.

Restrictions

- The operation cannot be issued when the TCE S-VOL is in Synchronizing or Paired status from a remote host.
- When even a single pair that is under the end operation (delete? synchronizing?) exists, the command cannot be executed.
- When even a single pair that is under the splitting operation exists, the command cannot be executed.
- When the pairsplit -mscas command is being executed for even a single SnapShot pair that is cascaded with a pair in the specified CTG, the command cannot be executed. The pairsplit -mscas processing is continued unless it becomes Failure or Pool Full. The processing is started from the continuation at the time of the next start even if the main switch of the primary array is turned off during the processing.

Also, review the -mscas restrictions in [Miscellaneous troubleshooting on page 10-7](#).

To split the TCE S-VOL/SnapShot V-VOL

In the example, the group name is ora. Group names of the cascaded SnapShot pairs are o0 and o1.

1. Execute the pairsplit -mscas command to the TCE pair. The status must be PAIR. For example:

```
c:\horcm\etc>pairsplit -g ora -macas Split-Marker 1
```

2. Verify that the status of the TCE pair is still PAIR by executing the pairdisplay command. The group in the example is ora.

```
c:\horcm\etc>pairdisplay -g ora
Group  PairVol(L/R) (Port#,TID, LU) ,Seq#,LDEV#.P/S,Status,Fence, Seq#,P-
LDEV# M
ora    oradb1(L)   (CL1-A , 1, 1 )85000174    1.PAIR -----,----- ---- -
ora    oradb1(R)   (CL1-B , 1, 2 )85000175    2.PAIR -----,----- ---- -
```

3. Confirm that the SnapShot Pair is split using the indirect or direct methods.
 - a. For the indirect method, execute the pairsyncwait command to verify that the P-VOL data has been transferred to the S-VOL. For example:

```
c:\horcm\etc>pairsyncwait -g ora -t 10000
UnitID CTGID   Q-Marker   Status   Q-Num
      0      3   00101231ef Done      2
```

The status may not display for one cycle after the command is issued.

Q-Marker counts up one by executing the pairsplit -mscas command.

- b. For the direct method, execute the `pairevtwait` command. For example:

```
c:\horcm\etc>pairevtwait -g o1 -s psus -t 300 10
pairevtwait : Wait status done.
```

Verify that the cascaded SnapShot pair is split by executing the `pairdisplay -v smk` command. The group in the example below is `o1`.

```
c:\horcm\etc>pairdisplay -g o1 -v smk
Group   PairVol(L/R)  Serial#  LDEV#  P/S  Status  UTC-TIME  ----
-Split-Maker-----
o1      URA_000(L)     85000175    2  P-VOL  PSUS   -        -
o1      URA_000(R)     85000175    3  S-VOL  SSUS   123456ef Split-
Marker
```

The TCE pair is released. For details on the `pairsplit` command, the `-mscas` option, and `pairsyncwait` command, refer to the *Hitachi Adaptable Modular Storage Command Control Interface (CCI) Reference Guide*.

Confirming data transfer when status is PAIR

When the TCE pair is in the PAIR status, data is transferred in regular cycles to the S-VOL. However, the P-VOL data that was settled as S-VOL data must be checked, as well as when the S-VOL data was settled.

When you execute the `pairsyncwait` command, any succeeding commands must wait until the P-VOL data at the time of the cycle update is reflected in S-VOL data.

For more information, please refer to the *Hitachi Adaptable Modular Storage Command Control Interface (CCI) Reference Guide*.

Pair creation/resynchronization for each CTG

In the pair creation/resynchronization performed with a specification of a certain group, renewal of the cycle is started from a pair for which the initial copy is completed first and the status of the pair above is changed to PAIR. A pair, initial copy for which is not completed before the first cycle renewal, renews the cycle taking the next occasion for the renewal. Therefore, the time when the status of the each pair is changed to PAIR may differ from the other one by the cycle time length.

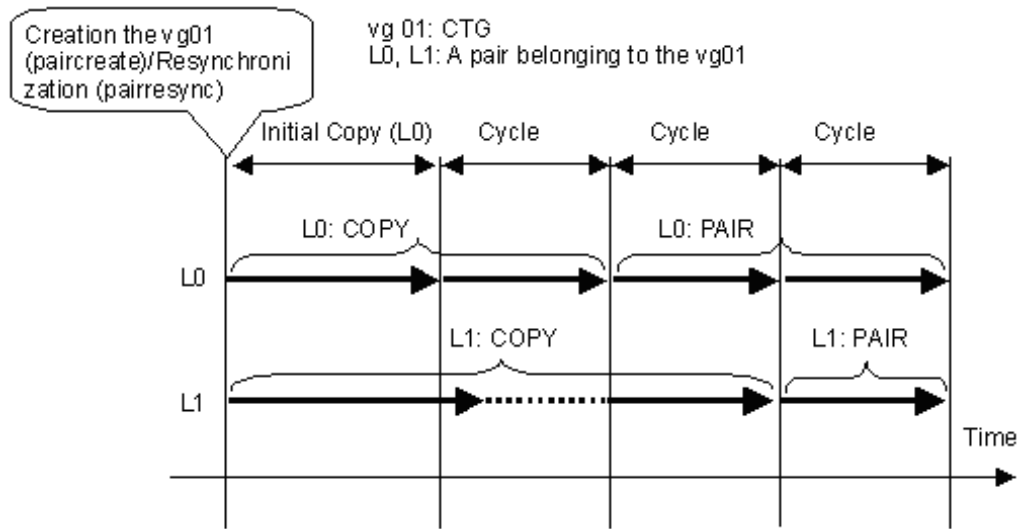


Figure B-3: Pair Creation/Resynchronization for each CTG-1

In the pair creation/resynchronization newly performed for each CTG, the time for renewing the cycle is decided by the pair for which the initial copy is completed first. The pair, the initial copy for which is completed later than the first completion of the initial copy, employs the renewed cycle starting from the cycle after next at the earliest.

When pair creation or resynchronization is performed for a group, the new cycle time begins for any pair in the group that is in PAIR status. A pair whose initial copy is not complete is not updated in the current update cycle, but will update during the next cycle. Cycle time is determined according to the first pair to complete the initial copy.

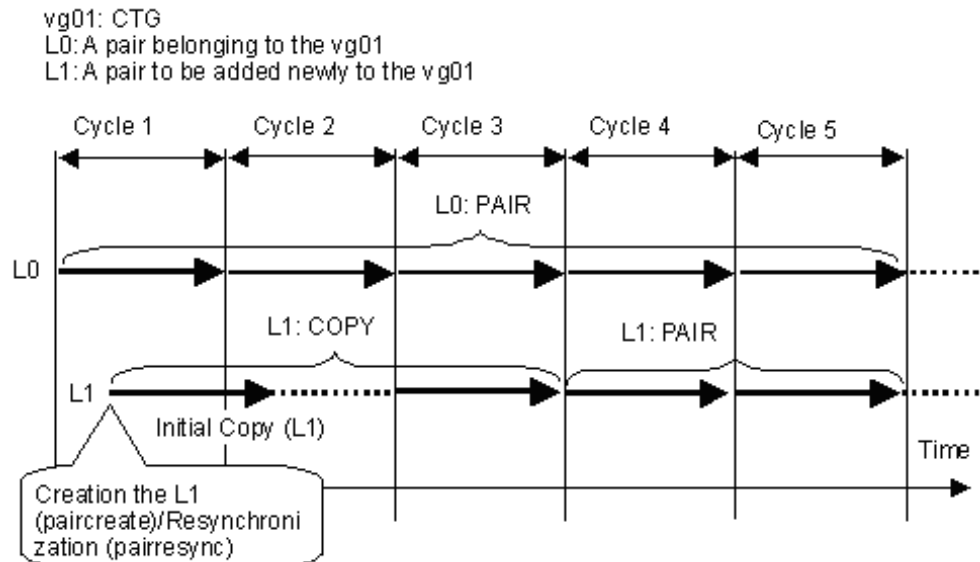


Figure B-4: Pair Creation/Resynchronization for each CTG-2

When a pair is newly added to the CTG, the pair is synchronized with the existing cycle timing. In the example, the pair is synchronized with the existing cycle from the cycle 3 and its status is changed to PAIR from the cycle 4.

- When the paircreate or pairresync command is executed, the pair undergoes the differential copy in the COPY status, undergoes the cyclic copy once, and then placed in the PAIR status. When a new pair is added to a CTG, which is already placed in the PAIR status, by the paircreate or pairresync command, the copy operation halts until the time of the existing cyclic copy after the differential copy is completed. Further, it is not placed in the PAIR status until the first cyclic copy is completed after it begins to act in time to the cycle. Therefore, the pair synchronization rate displayed by Navigator 2 or CCI may be 100% or not changed when the pair status is COPY.
- When you want to confirm the time from the stop of the copy operation to the start of the cyclic copy, check the start of the next cycle by displaying the predicted time of completing the copy using Navigator 2. For the procedure for displaying the predicted time of completing the copy, refer to section 5.2.7.

Response time of pairsplit command

A response time of a `pairsplit` command depends on a pair status and an option. [Table B-2](#) summarizes a response time for each CCI command.

In a case of splitting and deleting a pair with PAIR status, a completion of a processing takes time depending on the amount of differential data at P-VOL.

In a case of creating a remote snapshot, CCI command returns immediately but a completion of creating a snapshot depending on the amount of differential data at P-VOL. In order to check the completion, see Split-Marker of a remote snapshot is updated or a creation time of a snapshot is updated.

Table B-2: Response Time of CCI Commands

| Command | Options | Status | Response | Next Status | Remarks |
|-----------|--|-----------|-----------------------------|-----------------------------------|--|
| pairsplit | -S Delete pair | PAIR | Depend on differential data | SMPL | S-VOL data consistency guaranteed |
| | | COPY | Immediate | SMPL | No S-VOL data consistency |
| | | Others | Immediate | SMPL | No S-VOL data consistency |
| | -R Delete pair | PAIR | Immediate | SMPL (S-VOL only) | No S-VOL data consistency |
| | | COPY | Immediate | SMPL (S-VOL only) | No S-VOL data consistency Can not be executed for SSWS(R) status |
| | | Others | Immediate | SMPL (S-VOL only) | No S-VOL data consistency Can not be executed for SSWS(R) status |
| | -mscas Create remote snapshot (See note) | PAIR | Immediate | No change | A completion time depends on the amount of differential data. A completion can be check by Split-Marker and a creation time. Cycle updating process stops during creating a remote snapshot. |
| | | Others | — | — | — |
| | Others Split pair | PAIR | Depend on differential data | PSUS | S-VOL data consistency guaranteed |
| | | COPY | Immediate | PSUS | S-VOL data consistency guaranteed |
| Others | | Immediate | No change | S-VOL data consistency guaranteed | |



NOTE: Only -g option is valid. The -d option is not accepted. If there are pairs which status is not PAIR, in a CTG, a command cannot be accepted. All S-VOLs with PAIR status need to have corresponding cascading V-VOLs and MU# of these SnapShot pairs must match the MU# specified in a pairsplit -mscas command option.

Table B-3: TCE Pair Statuses and Relationship to Takeover

| Object Volume | | CCI Commands | | |
|---------------|---------|------------------|------------------|-------------|
| Attribute | Status | Paircheck Result | SVOL_Takeover | |
| | | | Data Consistency | Next Status |
| SMPL | - | To be confirmed | No | SMPL |
| P-VOL | - | To be confirmed | No | - |
| S-VOL | COPY | Inconsistent | No | COPY |
| | PAIR | To be analyzed | CTG | SSWS |
| | PSUS | Suspected | Pair | SSWS |
| | PSUS(N) | Suspected | No | PSUS(N) |
| | PFUS | Suspected | CTG | SSWS |
| | PSUE | Suspected | CTG | SSWS |
| | SSWS | Suspected | Pair | SSWS |

- Responses of paircheck.
 - To be confirmed: The object volume is not an S-VOL. Check is required.
 - Inconsistent: There is no write order guarantee of an S-VOL because an initial copy or a resync copy is on going or because of S-VOL failures. So SVOL_Takeover cannot be executed.
 - To be analyzed: Mirroring consistency cannot be determined just from a pair status of an S-VOL. However TCE does not support mirroring consistency, this result always shows that S-VOL has data consistency across a CTG not depending on a pair status of a P-VOL.
 - Suspected: There is no mirroring consistency of an S-VOL. If a pair status is PSUE or PFUS, there is data consistency across a CTG. If a pair status is PSUS or SSWS, there is data consistency for each pair in a CTG. In a case of PSUS(N), there is no data consistency.
- Data consistency after SVOL_Takeover and its response
 - CTG: Data consistency across a CTG is guaranteed.
 - Pair: Data consistency of each pair is guaranteed.
 - **No**: No data consistency of each pair.
 - Good: Response of takeover is normal.
 - NG: Response of takeover is an error. If a pair status of an S-VOL is PSUS, the pair status is changed to SSWS even if the response is an error.

See *Hitachi Adaptable Modular Storage Command Control Interface (CCI) Reference Guide* for more details about `horctakeover`.

Pair, group name differences in CCI and Navigator 2

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

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

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

Cascading with SnapShot

TCE P-VOLs and S-VOLs can be cascaded with SnapShot P-VOLs. This appendix discusses the supported configurations, operations, and statuses.

- ❑ [Cascade configurations](#)
- ❑ [Replication operations supported](#)
- ❑ [Status combinations, Read/write supported](#)
- ❑ [Guidelines and restrictions](#)
- ❑ [TCE, SnapShot behaviors compared](#)

Cascade configurations

In a cascaded system, a TCE P-VOL and/or S-VOL is shared with a SnapShot P-VOL, as shown in Figure C-1. No other configurations are supported.

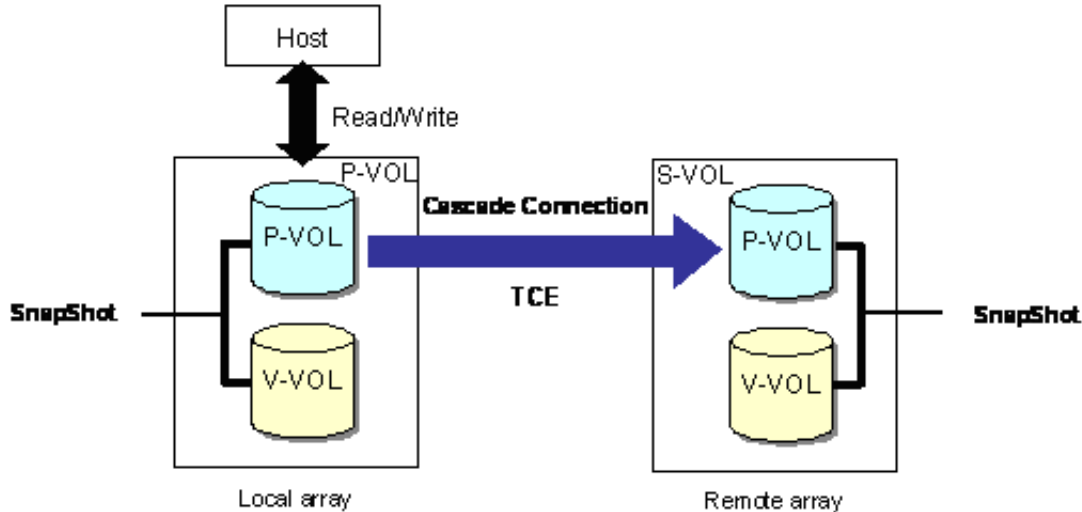


Figure C-1: Supported TCE, SnapShot Cascade Configurations

TCE cannot be cascaded with any other replication system. However, ShadowImage can be used on the same array as TCE. SnapShot can also be used outside a cascade situation on the same array as TCE.

Replication operations supported

A TCE pair operation can only be performed when the pair is in the appropriate pair status. For example, a split operation can only be performed when the TCE pair is in Paired status.

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 TCE and SnapShot operations that may be performed.

TCE operations supported

Table C-1 and Table C-2 show the TCE operations that can be performed on the shared volume when:

- Snapshot is cascaded with a TCE P-VOL
- Snapshot is cascaded with a TCE S-VOL
- The Snapshot pair has the status shown in the tables.

(SS = SnapShot)

Table C-1: Supported TCE Operations when TCE/SS P-VOL Cascaded

| TCE Pair Operation | SnapShot P-VOL Status | | | | | |
|--------------------|-----------------------|-------------------------|-------|----------------|---------|-------------------|
| | Paired | Synchronizing (Restore) | Split | Threshold over | Failure | Failure (Restore) |
| Create | Yes | No | Yes | Yes | Yes | No |
| Split | Yes | No | Yes | Yes | Yes | No |
| Re-sync | Yes | No | Yes | Yes | Yes | No |
| Restore | Yes | No | Yes | Yes | Yes | No |
| Delete | Yes | Yes | Yes | Yes | Yes | Yes |

Table C-2: Supported TCE Operations when TCE S-VOL/SS P-VOL Cascaded

| TrueCopy Pair Operation | SnapShot P-VOL Status | | | | | |
|-------------------------|-----------------------|-------------------------|-------|----------------|---------|-------------------|
| | Paired | Synchronizing (Restore) | Split | Threshold over | Failure | Failure (Restore) |
| Create | Yes | No | Yes | Yes | Yes | No |
| Split | Yes | N/A | Yes | Yes | Yes | N/A |
| Re-sync | Yes | Yes* | Yes | Yes | Yes | Yes* |
| Restore | Yes | No | Yes | Yes | Yes | No |
| Delete | Yes | Yes | Yes | Yes | Yes | Yes |
| pairsplit -mscas | Yes | N/A | No | No | No | No |

* A command will be received, but the pair status changes to Failure.

Snapshot operations supported

Table C-3 and Table C-4 show the SnapShot operations that can be performed on the shared volume when:

- Snapshot is cascaded with a TCE P-VOL
- Snapshot is cascaded with a TCE S-VOL
- The TCE pair has the status shown in the tables.

(SS = SnapShot)

Table C-3: Supported SnapShot Operations when TCE/SS P-VOL Cascaded

| SnapShot Pair Operation | TCE P-VOL Status | | | | |
|-------------------------|------------------|---------------|-------|-----------|---------|
| | Paired | Synchronizing | Split | Pool Full | Failure |
| Create | Yes | Yes | Yes | Yes | Yes |
| Split | Yes | Yes | Yes | Yes | Yes |
| Re-sync | Yes | Yes | Yes | Yes | Yes |
| Restore | No | No | Yes | Yes | No |
| Delete | Yes | Yes | Yes | Yes | Yes |

Table C-4: Supported SnapShot Operations when TCE S-VOL/SS P-VOL Cascaded

| SnapShot Pair Operation | TrueCopy S-VOL Status | | | | | | | |
|-------------------------|-----------------------|---------------|---------------------|-----|--------------|----------|------|-----------|
| | Paired | Synchronizing | Split RWR mode mode | | Inconsistent | Takeover | Busy | Pool Full |
| Create | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Split | No | No | Yes | Yes | No | Yes | Yes | No |
| Re-sync | Yes | Yes | Yes | Yes | No* | Yes | Yes | Yes |
| Restore | No | No | Yes | No | No | Yes | No | No |
| Delete | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |

*SnapShot pair status does not Split; a command cannot be received.

Status combinations, Read/write supported

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

- Table C-5 shows status and read/write allowed when the SnapShot P-VOL = a TrueCopy P-VOL.
- Table C-6 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 is allowed
- No Combined status is not allowed.
- R/W Read/Write by a host is allowed.
- R Read by a host is allowed, write is not allowed.
- W Write by a host is allowed, read is not allowed.
- Err Pair operation causes an error.

Table C-5: Read/Write, Status Allowed for TCE/SS P-VOL

| TCE P-VOL Status | SnapShot P-VOL Status | | | | | |
|------------------|-----------------------|-------------------------|---------|----------------|---------|-------------------|
| | Paired | Synchronizing (Restore) | Split | Threshold over | Failure | Failure (Restore) |
| Paired | Yes R/W | No | Yes R/W | Yes R/W | Yes R/W | No |
| Synchronizing | Yes R/W | No | Yes R/W | Yes R/W | Yes R/W | No |
| Split | Yes R/W | Yes R/W | Yes R/W | Yes R/W | Yes R/W | Err R/W |
| Pool Full | Yes R/W | Yes R/W | Yes R/W | Yes R/W | Yes R/W | Err R/W |
| Failure | Yes R | Err R/W | Yes R/W | Yes R/W | Err R/W | Err R/W |

Table C-6: Read/Write, Status Allowed for TCE S-VOL/SS P-VOL

| TCE S-VOL Status | SnapShot P-VOL Status | | | | | |
|------------------|-----------------------|-------------------------|---------|----------------|---------|-------------------|
| | Paired | Synchronizing (Restore) | Split | Threshold over | Failure | Failure (Restore) |
| Paired | Yes R | No | Yes R | Yes R | Yes R | No |
| Synchronizing | Yes R | No | Yes R | Yes R | Yes R | No |
| Split R/W mode | Yes R/W | Yes R/W | Yes R/W | Yes R/W | Yes R/W | Err R/W |
| Split R mode | Yes R | No | Yes R | Yes R | Yes R | No |
| Inconsistent | Err R/W | No | Err R/W | Err R/W | Err R/W | No |
| Takeover | Yes R/W | Yes R/W | Yes R/W | Yes R/W | Yes R/W | Err R/W |
| Busy | Yes R/W | No | Yes R/W | Yes R/W | Yes R/W | No |
| Pool Full | Yes R | No | Yes R | Yes R | Err R | No |

Guidelines and restrictions

The following provides basic guidelines for cascading SnapShot and TCE:

- SnapShot is not required for TCE.
- A cascaded SnapShot P-VOL may be paired with up to 32 SnapShot V-VOLs.
- A V-VOL cannot be cascaded with TCE.
- A SnapShot pair must be in Split status when performing a TCE pair operation.
- A TCE pair must be in Split or Failure status when performing a SnapShot pair operation.
- I/O performance on the local side is lowered when a TCE P-VOL is cascaded with a SnapShot P-VOL.
- When host I/O activity is high, performance is maximized when the TCE pair is split.
- TCE and SnapShot must use the same data pools. SnapShot data pool numbers must be the same as TCE data pool numbers.
- When SnapShot pair status for a TCE P-VOL/SnapShot P-VOL changes to Reverse Synchronizing or Failure during a restore operation, the creation or resynchronization of the TCE pair cannot be performed. The SnapShot pair must be recovered first.
- If a TCE pair is in Busy status, in which the S-VOL is in the process of being restored from the remote data pool, and this operation fails, the SnapShot pair status becomes Failure. It cannot recover unless the TCE and SnapShot pairs are deleted, and the SnapShot pair is re-created.
- When TCE pair status for a TCE P-VOL/SnapShot P-VOL changes to Failure status during restoration, the creation or resynchronization of the SnapShot pair cannot be performed.
- A horctakeover operation cannot be performed when the TCE S-VOL/SnapShot P-VOL is being restored by SnapShot.

Cascading with SnapShot on the remote side

TCE has the capability to coordinate a SnapShot backup of the remote S-VOL. This is carried out in conjunction with writing of all data from the local array's cache to the P-VOL, then updating the S-VOL before the snapshot is taken. The backup is executed as follows:

1. The host issues a remote snapshot command to the P-VOL.
2. The local array requests the creation of a remote snapshot.
3. TCE P-VOL data is updated with all data remaining in cache memory and stabilized, then updated to the S-VOL.
4. The remote SnapShot pair, if it already exists, is split.
5. A snapshot of the S-VOL is created or updated on the remote array.

The benefits of creating the remote snapshot are the following:

- TCE simplifies the snapshot backup operation by issuing only one command.
- The latest P-VOL data is backed up on the remote array.
- The timing of the required operations can be synchronized, which enables a consistent backup at the time-point that the command is issued.
- While remote snapshot processing is in progress, the TCE pair status remains Paired and the S-VOL continues to be updated.
- The remote snapshot backup frequency can be very short, from several seconds to minutes.



NOTE: When remote snapshots are made of TCE pairs in a consistency group, cycle update processing for the consistency group stops.

TCE, SnapShot behaviors compared

Table C-7 is provided to help you understand the different behaviors that can be expected due to conditions that may arise.

Table C-7: TCE, SnapShot Behaviors

| Condition | TCE Behavior | SnapShot Behavior |
|--|--|---|
| Data pool threshold over | Pair statuses do not change. Data pool status changes to Threshold Over. | Same as TCE. |
| Data pool full at local | Pair status of P-VOL with Paired status changes to Pool Full. Pair status of P-VOL with Synchronizing status changes to Failure. | Pair status changes to Failure. |
| Data pool full at remote | Pair status of P-VOL changes to Failure. Pair status of S-VOL with Paired status changes to Pool Full. Pair status of S-VOL with Synchronizing status changes to Inconsistent. | Pair status changes to Failure. |
| Data consistency when data pool full | S-VOL data stays consistent at consistency- group level. | V-VOL data is invalid. |
| How to recover from Failure | Resync the pair. | Delete then recreate the pair. |
| Failures | Failures at local: P-VOL changes to Failure. S-VOL does not change. Data consistency is ensured if pair status of S-VOL is Paired. Failures at remote: P-VOL changes to Failure. S-VOL changes to Inconsistent. No data consistency for S-VOL. | Pair status changes to Failure and V-VOL data is invalid. |
| Number of consistency groups supported | 16 | 128 |

Installing TCE when Cache Partition Manager is in use

This appendix provides important information for Cache Partition Manager when TCE is installed.

- ❑ [Initializing Cache Partition when TCE and SnapShot are installed](#)

Initializing Cache Partition when TCE and SnapShot are installed

TCE and SnapShot use part of the cache to manage internal resources, causing a reduction in the cache capacity used by Cache Partition Manager.

Cache partition information should be initialized as follows, when TCE or SnapShot are installed after Cache Partition Manager is installed:

- All the logical units should be moved to the master partitions on the side of the default owner controller.
- All the sub-partitions must be deleted and the size of each master partition should be reduced to half of the user data area after installation of TCE or SnapShot.

Figure D-1 shows an example of Cache Partition Manager usage. Figure D-2 shows an example where TCE/SnapShot is installed when Cache Partition Manager already in use.

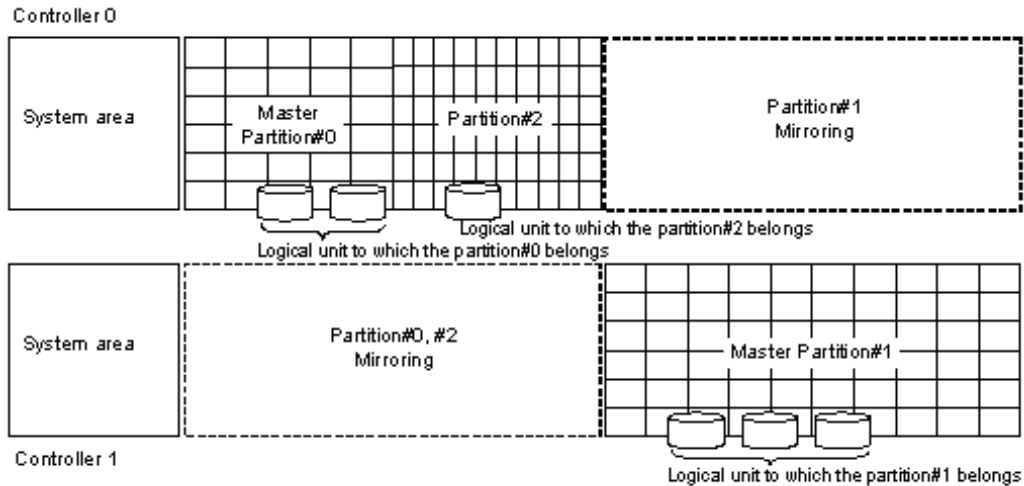


Figure D-1: Cache Partition Manager Usage

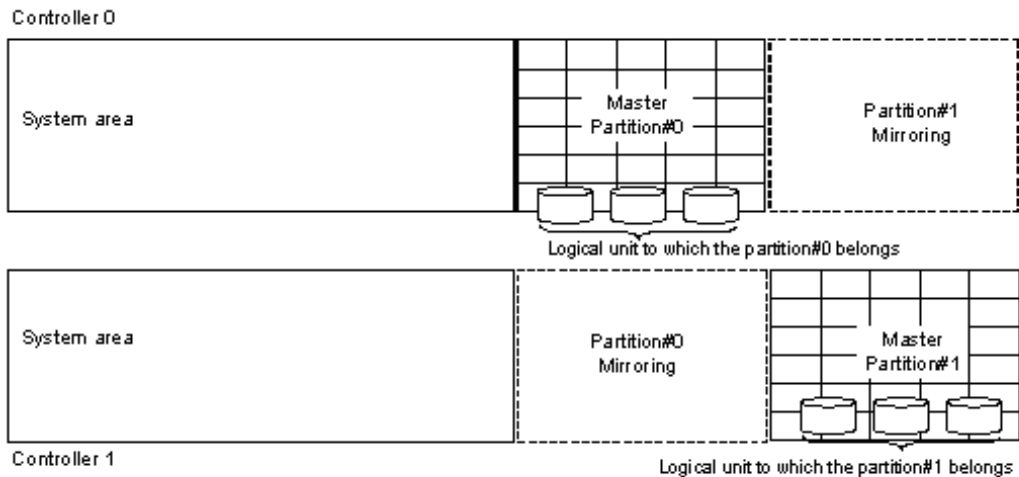


Figure D-2: TCE or SnapShot Installation with Cache Partition Manager

On the remote array, Synchronize Cache Execution mode should be turned off to avoid TCE remote path failure.

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 fibre.

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 E-1 shows an illustration of WDM.

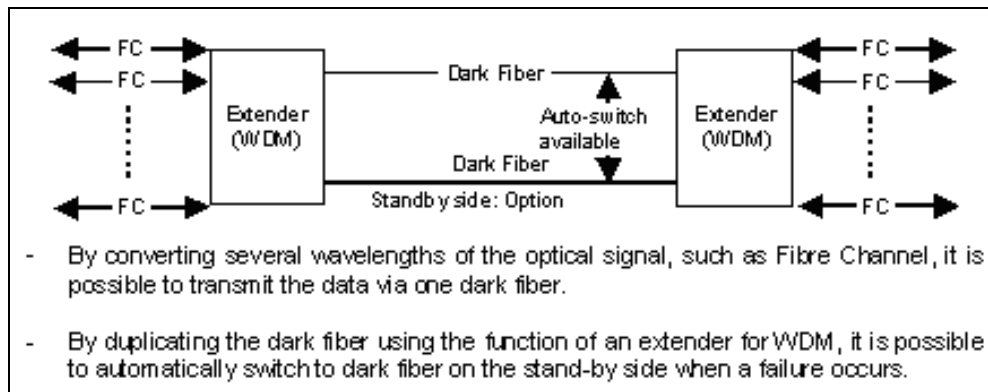


Figure E-1: Wavelength Division Multiplexing

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 a 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 E-2.

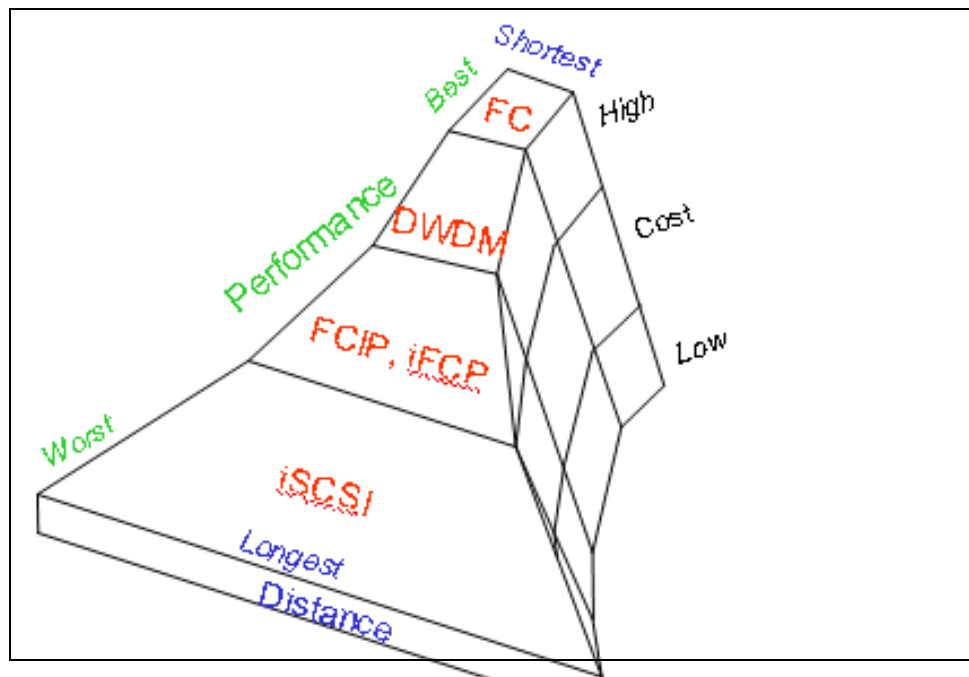


Figure E-2: Dark Fiber with WDM

- The WDM function can also be multiplexed in one dark fiber for G Ethernet.
- If switching is executed during a dark fibre failure, data transfer must be moved to another path, as shown in [Figure E-3](#).

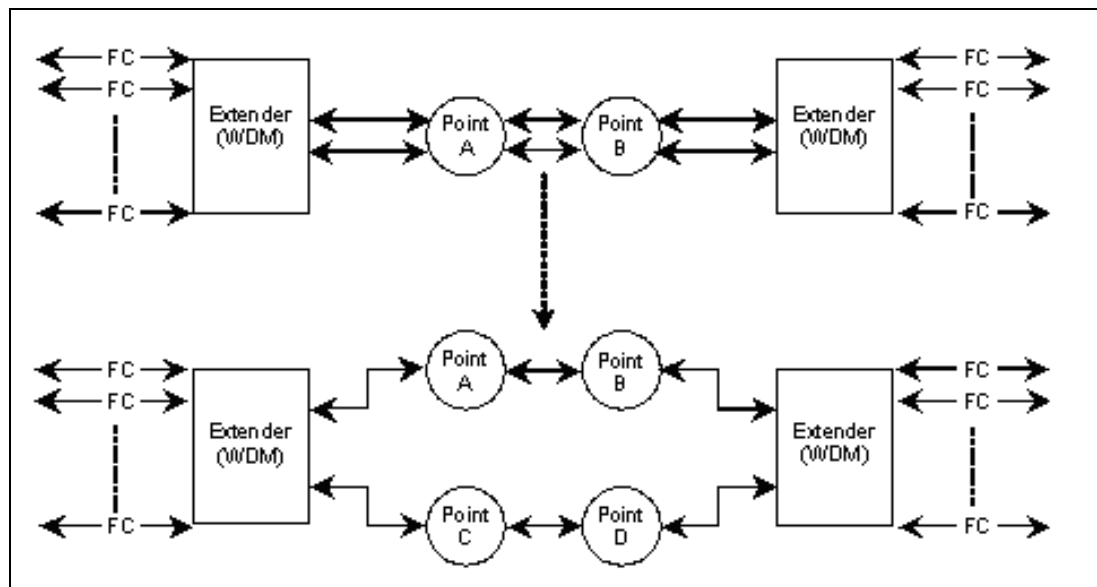


Figure E-3: Dark Fiber Failure

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

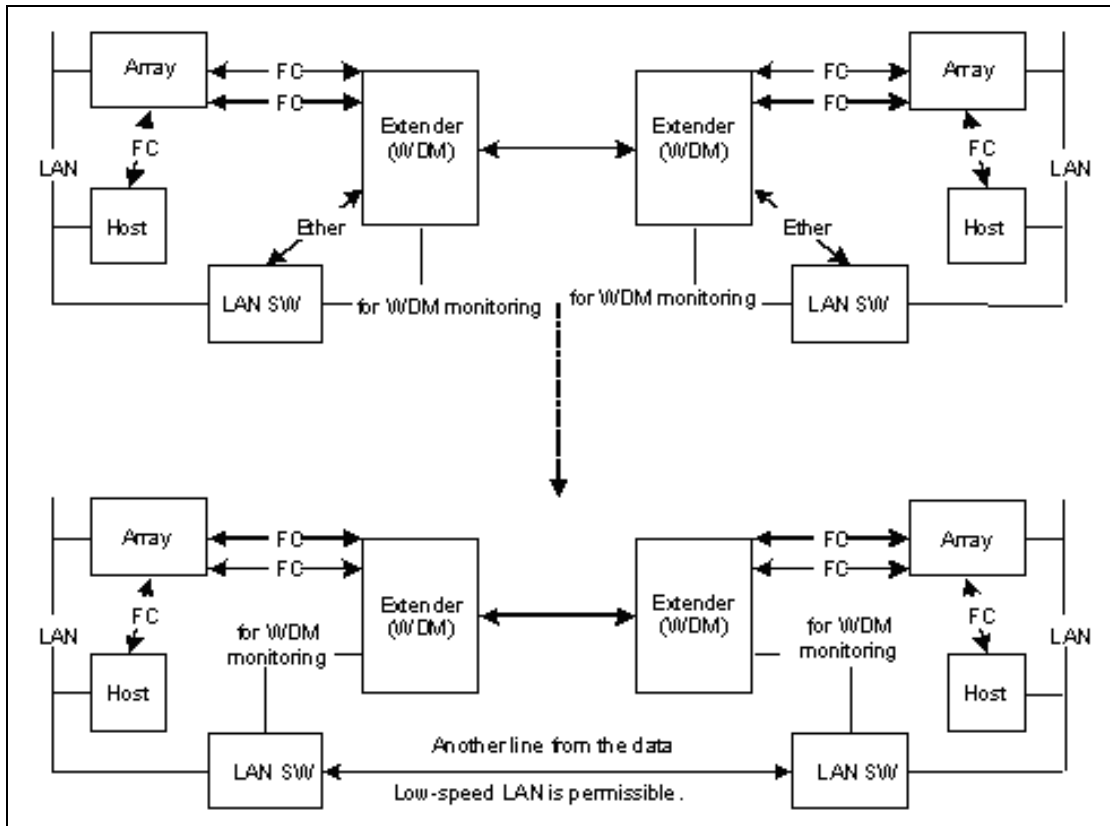


Figure E-4: Line for Monitoring



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.

A

array

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

asynchronous

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

B

background copy

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

Bps

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

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

C

cache

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

capacity

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

CCI

See command control interface.

CLI

See command line interface.

cluster

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

cluster capacity

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

command control interface (CCI)

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

command devices

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

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

command line interface (CLI)

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

concurrency of S-VOL

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

concurrent copy

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

configuration definition file

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

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

CRC

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

CTG

See Consistency Group.

cycle time

A user specified time interval used to execute recurring data updates for remote copying. Cycle time updates are set for each 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 path

See remote path.

data pool

One or more disk volumes designated to temporarily store un-transferred 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

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.

differential data control

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

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

Glossary–4

Differential Management Logical Unit (DMLU)

An exclusive volume used for storing data when the array system is powered down.

differential-data

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

disaster recovery

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

disk array

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

DMLU

See Differential Management-Logical Unit.

dual copy

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

duplex

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

E

entire copy

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

extent

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

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

F

failover

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

fallback

Refers to the process of restarting business operations at a local site using the P-VOL. It takes place after the 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(s) per second.

granularity of differential data

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

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

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

GUI

Graphical user interface.

I

I/O

Input/output.

initial copy

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

initiator ports

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

IOPS

I/O per second.

iSCSI

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

iSNS

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

L

LAN

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

load

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

logical

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

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

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.

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

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).

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

point-in-time logical copy

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

pool volume

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

primary or local site

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

primary volume (P-VOL)

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

P-VOL

See primary volume.

R**RAID**

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

Recovery Point Objective (RPO)

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

Recovery Time Objective (RTO)

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

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

remote or target site

Maintains mirrored data from the primary site.

remote path

Also called the data path, the remote path is a link that connects ports on the local storage system and the remote storage system. Two remote paths must be set up for each AMS array (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.

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

SMPL

Simplex.

snapshot

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

SNM2

See Storage Navigator Modular 2.

SSD

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

Storage Navigator Modular 2

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

suspended status

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

S-VOL

See secondary volume.

S-VOL determination

Independent of update operations, S-VOL determination replicates the S-VOL on the remote 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.

T**target copy**

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

V**virtual volume (V-VOL)**

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

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

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.

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

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

Glossary–14

Index

A

- AMS models supported for TCE pairs [4-2](#)
- array problems, recovering pairs after [10-3](#)
- arrays, swapping I/O to maintain [8-9](#)
- assigning LUs to data pools [6-6](#)
- assigning pairs to a consistency group [7-4](#)

B

- backing up the S-VOL [8-2](#)
- bandwidth
 - calculating [2-7](#)
 - changing [9-6](#)
 - measuring workload for [2-3](#)
- basic operations [7-2](#)
- behavior when data pool over [C-7](#)
- best practices for remote path [3-18](#)
- block size, checking [4-3](#)

C

- Cache Partition Manager, initializing for TCE installation [D-2](#)
- calculating data pool size [2-4](#)
- capacity, supported maximum [4-9](#)
- cascading with SnapShot [C-2](#)
- CCI
 - description [1-7](#)
 - using to
 - change command device [B-3](#)
 - create pairs [B-8](#)
 - define config def file [B-4](#)
 - monitor pair status [B-7](#)
 - release command device [B-2](#)
 - release pairs [B-10](#)
 - resync pairs [B-9](#)
 - set command device [B-2](#)
 - set environment variable [B-6](#)
 - split pairs [B-8](#)
 - suspend pairs [B-9](#)

- CCI, using to
 - set LU mapping [B-3](#)
- changing a command device using CCI [B-3](#)
- changing Threshold value [9-5](#)
- CLI
 - description [1-7](#)
 - using to
 - back up S-VOL [8-4](#)
 - create pairs [A-15](#)
 - display pair status [A-14](#)
 - enable, disable TCE [A-3](#)
 - install TCE [A-3](#)
 - resynchronize pairs [A-16](#)
 - set cycle time [A-9](#)
 - set the remote path [A-11](#)
 - set up DMLU [A-5](#)
 - set up the pool [A-6](#)
 - split pairs [A-16](#)
 - swap pairs [A-16](#)
 - uninstall TCE [A-4](#)
- collecting write-workload data [2-3](#)
- command device
 - changing [B-3](#)
 - releasing [B-2](#)
 - setup [B-2](#)
- configuration definition file, defining [B-4](#)
- configurations supported for SnapShot
 - cascade [C-2](#)
- consistency group
 - checking status with CLI [A-19](#)
 - creating, assigning pairs to [7-4](#)
 - description [1-6](#)
 - using CCI for operations [B-12](#)
- Copy Pace, changing [9-7](#)
- Copy Pace, specifying [7-5](#)
- create pair procedure [7-3](#), [7-4](#)
- CTG. See consistency group
- cycle time, monitoring, changing in GUI [9-6](#)
- cycle time, setting up with CLI [A-9](#)

D

- dark fibre [E-1](#)
- Data path, planning [3-5](#)
- data path. See remote path
- data pools
 - deleting [9-9](#)
 - description [1-4](#)
 - expanding [9-5](#)
 - measuring workload for [2-3](#)
 - monitoring usage [9-4](#)
 - setting up with CLI [A-6](#)
 - setting up with GUI [6-6](#)
 - shortage [10-2](#)
 - sizing [2-4](#)
 - Threshold field [6-6](#)
- data, measuring write-workload [2-3](#)
- deleting
 - data pool [9-9](#)
 - DMLU [9-9](#)
 - remote path [9-9](#)
 - volume pair [9-8](#)
- designing the system [2-1](#)
- Differential Management Logical Unit. See DMLU
- disaster recovery process [8-11](#)
- DMLU
 - defining [6-5](#)
 - deleting [9-9](#)
 - description [1-7](#)
 - setting up with CLI [A-5](#)
- dynamic disk with Windows Server 2000 [4-5](#)
- dynamic disk with Windows Server 2003 [4-6](#)

E

- enabling, disabling TCE [6-4](#)
- enabling, with CLI [A-3](#)
- environment variable [B-6](#)
- error codes, failure during resync [10-4](#)
- Event Log, using [10-6](#)
- expanding data pool size [9-5](#)
- extenders [E-1](#)

F

- failback procedure [8-12](#)
- Fibre Channel remote path requirements and configurations [3-5](#)
- Fibre Channel, port transfer-rate [3-10](#)

G

- Group Name, adding [7-5](#)
- GUI, description [1-7](#)
- GUI, using to
 - assign pairs to a consistency group [7-4](#)
 - create a pair [7-3](#)
 - delete a data pool [9-9](#)

- delete a DMLU [9-9](#)
- delete a pair [9-8](#)
- delete a remote path [9-9](#)
- install, enable/disable TCE [6-3](#)
- monitor data pool usage [9-4](#)
- monitor pair status [9-2](#)
- resynchronize a pair [7-6](#)
- set up data pool [6-6](#)
- set up DMLU [6-5](#)
- set up remote path [6-9](#)
- split a pair [7-5](#)
- swap a pair [7-7](#)

H

- horctakeover [8-11](#)
- host group, connecting to HP server [4-4](#)
- host recognition of P-VOL, S-VOL [4-4](#)
- host time-out recommendation [4-3](#)

I

- initial copy [7-2](#)
- installing TCE with CLI [A-3](#)
- installing TCE with GUI [6-2](#)
- interfaces for TCE [1-7](#)
- iSCSI remote path requirements and configurations [3-11](#)

L

- LAN requirements [3-3](#)
- logical units, recommendations [4-3](#)

M

- maintaining local array, swapping I/O [8-9](#)
- MC/Service Guard [4-4](#)
- measuring write-workload [2-3](#)
- migrating volumes from earlier AMS models [4-2](#)
- monitoring
 - data pool usage [9-4](#)
 - pair status [9-2](#)
 - remote path [9-6](#)
- moving data procedure [8-10](#)

O

- operating systems, restrictions with [4-3](#)
- operations [7-2](#)

P

- Pair Name field, differences on local, remote array [7-4](#)
- pair names and group names, Nav2 differences from CCI [B-17](#)
- pairs
 - assigning to a consistency group [7-4](#)
 - creating [7-3](#)
 - deleting [9-8](#)

- description 1-3
 - displaying status with CLI A-14
 - monitoring status with GUI 9-2
 - monitoring with CCI B-7
 - recommendations 4-3
 - recovering from Pool Full 10-2
 - resynchronizing 7-6
 - splitting 7-5
 - status definitions 9-2
 - swapping 7-7
 - planning
 - arrays 4-2
 - remote path 3-5
 - TCE volumes 4-3
 - Planning the remote path 3-5
 - port transfer-rate 3-10
 - prerequisites for pair creation 7-2
- R**
- RAID groups and volume pairs 4-3
 - recovering after array problems 10-3
 - recovering from failure during
 - resync 10-4
 - release a command device, using CCI B-2
 - remote array, shutdown, TCE tasks 9-10
 - Remote path
 - planning 3-5
 - remote path
 - best practices 3-18
 - deleting 9-9
 - description 3-5
 - monitoring 9-6
 - planning 3-5
 - preventing blockage 3-18
 - requirements 3-5
 - setup with CLI A-11
 - setup with GUI 6-9
 - supported configurations 3-5
 - Requirements
 - bandwidth, for WANs 2-8
 - LAN 3-3
 - requirements 5-2
 - response time for pairsplit B-15
 - resynchronization error codes 10-4
 - resynchronization errors, correcting 10-4
 - resynchronizing a pair 7-6
 - rolling averages, and cycle time 2-5
 - RPO, checking 9-8
 - RPO, update cycle 2-2
- S**
- scripts for backups (CLI) 8-2, A-21
 - setting port transfer-rate 3-10
 - SnapShot
 - behaviors vs TCE's C-7
 - cascading with C-2
 - supported operations when cascaded C-4
 - specifications 5-3
 - split pair procedure 7-5
 - statuses, pair 9-2
 - statuses, supported for cascading C-4
 - supported capacity calculation 4-9
 - supported remote path configurations 3-5
 - S-VOL, backing up 8-2
 - S-VOL, updating 7-6
 - swapping pairs 7-7
- T**
- takeover 8-11
 - TCE
 - array combinations 4-2
 - backing up the S-VOL 8-2
 - behaviors vs SnapShot's C-7
 - calculating bandwidth 2-7
 - changing bandwidth 9-6
 - create pair procedure 7-4
 - data pool
 - description 1-4
 - setup 6-6
 - sizing 2-4
 - environment 1-3
 - how it works 1-2
 - installing, enabling, disabling 6-3
 - interface 1-7
 - monitoring pair status 9-2
 - operations 7-2
 - operations before firmware
 - updating 9-10
 - pair recommendations 4-3
 - procedure for moving data 8-10
 - remote path configurations 3-5
 - requirements 5-2
 - setting up the remote path 6-9
 - setup 6-5
 - setup wizard 7-3
 - splitting a pair 7-5
 - supported operations when cascaded C-3
 - typical environment 1-3
 - Threshold field, changing 9-5
 - threshold reached, consequences 6-6
- U**
- uninstalling with CLI A-4
 - uninstalling with GUI 6-4
 - update cycle 1-2, 1-4, 2-2
 - specifying cycle time 9-6
 - updating firmware, TCE tasks 9-10
 - updating the S-VOL 7-6
- V**
- volume pair description 1-3

volume pairs, recommendations [4-3](#)

W

WAN

- bandwidth requirements [2-8](#)
- configurations supported [3-12](#)
- general requirements [3-3](#)
- types supported [3-3](#)

WDM [E-1](#)

Windows Server 2000 restrictions [4-5](#)

Windows Server 2003 restrictions [4-5](#)

wizard, TCE setup [7-3](#)

WOCs, configurations supported [3-14](#)

write order [1-4](#)

write-workload [2-3](#)

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-97DF8054-07