

| Implementing Hitachi Universal Replicator Software on Open Systems

TSI0150

Notice: This document is for informational purposes only, and does not set forth any warranty, express or implied, concerning any equipment or service offered or to be offered by Hitachi Data Systems. This document describes some capabilities that are conditioned on a maintenance contract with Hitachi Data Systems being in effect, and that may be configuration-dependent, and features that may not be currently available. Contact your local Hitachi Data Systems sales office for information on feature and product availability.

Hitachi Data Systems sells and licenses its products subject to certain terms and conditions, including limited warranties. To see a copy of these terms and conditions prior to purchase or license, please call your local sales representative to obtain a printed copy. If you purchase or license the product, you are deemed to have accepted these terms and conditions.

THE INFORMATION CONTAINED IN THIS MANUAL IS DISTRIBUTED ON AN "AS IS" BASIS WITHOUT WARRANTY OF ANY KIND, INCLUDING WITHOUT LIMITATION, ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NONINFRINGEMENT. IN NO EVENT WILL HDS BE LIABLE TO THE END USER OR ANY THIRD PARTY FOR ANY LOSS OR DAMAGE, DIRECT OR INDIRECT, FROM THE USE OF THIS MANUAL, INCLUDING, WITHOUT LIMITATION, LOST PROFITS, BUSINESS INTERRUPTION, GOODWILL OR LOST DATA, EVEN IF HDS EXPRESSLY ADVISED OF SUCH LOSS OR DAMAGE.

Hitachi Data Systems is registered with the U.S. Patent and Trademark Office as a trademark and service mark of Hitachi, Ltd. The Hitachi Data Systems logotype is a trademark and service mark of Hitachi, Ltd.

The following terms are trademarks or service marks of Hitachi Data Systems Corporation in the United States and/or other countries:

Hitachi Data Systems Registered Trademarks

Essential NAS Platform, Hi-Track, ShadowImage, TrueCopy, Universal Storage Platform

Hitachi Data Systems Trademarks

HiCard, HiPass, Hi-PER Architecture, Hi-Star, Universal Star Network

All other trademarks, trade names, and service marks used herein are the rightful property of their respective owners.

NOTICE:

Notational conventions: 1KB stands for 1,024 bytes, 1MB for 1,024 kilobytes, 1GB for 1,024 megabytes, and 1TB for 1,024 gigabytes, as is consistent with IEC (International Electrotechnical Commission) standards for prefixes for binary and metric multiples.

©2011, Hitachi Data Systems Corporation. All Rights Reserved

HDS Academy 1031

Contact Hitachi Data Systems at www.hds.com.

Product Names mentioned in courseware:

Enterprise Storage Systems

- Hitachi Virtual Storage Platform
- Hitachi Universal Storage Platform® V
- Hitachi Universal Storage Platform® VM
- Hitachi Universal Storage Platform®
- Hitachi Network Storage Controller

Legacy Products:

- Hitachi Lightning 9900™ V Series enterprise storage systems
- Hitachi Lightning 9900™ Series enterprise storage systems

Modular Storage Systems

- Hitachi Adaptable Modular Storage
- Hitachi Workgroup Modular Storage
- Hitachi Adaptable Modular Storage 2000 Family

Legacy Products:

- Hitachi Thunder 9500™ Series modular storage systems
- Hitachi Thunder 9200V™ entry-level storage

Management Tools

- Hitachi Basic Operating System
- Hitachi Basic Operating System V
- Hitachi Resource Manager™ utility package
 - ♦ Module Volume Migration Software
 - ♦ LUN Manager/LUN Expansion
 - ♦ Network Data Management Protocol (NDMP) agents
 - ♦ Logical Unit Size Expansion (LUSE)
 - ♦ Cache Partition Manager feature
 - ♦ Cache Residency Manager feature
 - ♦ Storage Navigator program
 - ♦ Storage Navigator Modular program
 - ♦ Storage Navigator Modular 2 program

Replication Software

Remote Replication:

- Hitachi Universal Replicator
- Hitachi TrueCopy® Heterogeneous Remote Replication software bundle
- Hitachi TrueCopy® Remote Replication software bundle (for modular systems)

-
- Hitachi TrueCopy® Synchronous
 - Hitachi TrueCopy® Asynchronous
 - Hitachi TrueCopy® Extended Distance

In-System Replication:

- Hitachi ShadowImage® Heterogeneous Replication (for enterprise systems)
- Hitachi ShadowImage® Replication (for modular systems)
- Hitachi Copy-on-Write Snapshot

Hitachi Command Suite

- Hitachi Chargeback
- Hitachi Device Manager
- Hitachi Dynamic Link Manager
- Hitachi Global Link Manager
- Hitachi Global Reporter
- Hitachi Path Provisioning
- Hitachi Protection Manager
- Hitachi QoS for File Servers
- Hitachi QoS for Oracle
- Hitachi Replication Monitor
- Hitachi Storage Services Manager
- Hitachi Tiered Storage Manager
- Hitachi Tuning Manager

Other Software

- Hitachi Backup and Recovery, powered by CommVault®
- Hitachi Backup Services Manager, powered by APTARE®
- Hitachi Business Continuity Manager
- Hitachi Command Control Interface (CCI)
- Hitachi Dynamic Provisioning
- Hitachi Storage Resource Management Solutions
- Hitachi Volume Migration
- Hi-Track® Monitor

Other Solutions and Terms

- Hitachi Content Platform
- Hitachi NAS Platform®
- Hitachi High-performance NAS Platform, powered by BlueArc®
- Hi-Star™ crossbar switch architecture
- Hitachi Universal Star Network™ V

Contents

INTRODUCTION	IX
Welcome and Introductions	ix
Course Description	ix
Prerequisites	x
Course Objectives	x
Course Topics	xi
Learning Paths	xii
HDS Academy Is on Twitter and LinkedIn	xiii
Terms and Acronyms	xiv
Replication Terminology	xv
 1. OVERVIEW	 1-1
Module Objectives	1-1
Purpose	1-2
Two Data Center Overview	1-3
Key Features – Journals	1-4
Key Features – MxN Consistency Group	1-6
Group Associations	1-7
Documents	1-8
Components	1-9
 2. SPECIFICATIONS	 2-1
Module Objectives	2-1
Common Storage System Specifications	2-2
Volume Specifications	2-3
Specifications: System Option Modes	2-6
Pair Volume Status	2-9
Module Review	2-13
 3. ARCHITECTURE AND INTERNAL OPERATIONS	 3-1
Module Objectives	3-1
Bitmap Areas	3-2
Bitmap Areas – Mirror Unit Numbers	3-6
Moving Data Between Storage Systems	3-7
Structure of Journal Volume	3-15
Performance Considerations	3-17
Performance Considerations and Restrictions	3-18
Best Practices: Journal Throughput	3-19
Best Practices: Sizing Journal Volumes	3-20
Best Practices: Journal Volume Physical Characteristics	3-22
Journal Cache Specifications	3-26
Configuration Planning	3-27
Review – How Universal Replicator Works	3-29
Review: Best Practice Sizing for Journal Volumes	3-31
Module Review	3-32
 4. STORAGE NAVIGATOR CONFIGURATION	 4-1
Module Objectives	4-1
Preparation: License Key Installation Procedure	4-2
Preparation: Map Logical Devices	4-3
Preparation: VSP Map Logical Devices	4-5
Preparation: Establish Paths	4-6
Preparation: Replication Links	4-8
Configuration: Overview	4-12
Configuration: Port Attributes	4-15

Configuration: Add DKC	4-17
Configuration: Journal Groups	4-28
Configuration: Journal Groups Remote Command Devices	4-31
Configuration: Journal Groups	4-34
Journal Group Configuration Details	4-35
Review: Preparation Checklist	4-53
Module Review	4-54
5. STORAGE NAVIGATOR PAIR OPERATIONS.....	5-1
Module Objectives	5-1
Preparation for Operations	5-2
Commands Overview	5-3
paircreate	5-5
pairdisplay	5-9
pairsplit	5-14
pairresync	5-16
Change Pair Option	5-18
pairsplit	5-19
Usage Monitor	5-22
History	5-27
Troubleshooting Suspended Pairs	5-28
Module Review	5-30
6. HITACHI COMMAND CONTROL INTERFACE CONFIGURATION AND OPERATIONS.....	6-1
Module Objectives	6-1
Overview	6-2
CCI Checklist	6-5
Configuration	6-6
Configuration: Command Device	6-8
Configuration: VSP Command Device	6-9
Configuration: Remote Command Device	6-10
Configuration Files	6-11
CCI Configuration	6-14
CCI Configuration: Command Device (Solaris)	6-15
Configuration Files	6-16
Configuration Files: Group and Device names	6-18
Configuration Files: Restrictions	6-19
Configuration Files: Mirror IDs	6-20
Configuration Files	6-21
Configuration: Services File	6-22
Configuration: mkconf	6-23
Configuration: Protection Facility	6-29
Configuration	6-30
Commands	6-31
Scripted Commands for Disaster Recovery	6-44
Microsoft Windows Subcommands	6-46
Configuration Setting Commands	6-49
Troubleshooting	6-50
Module Review	6-52
7. DATA PROTECTION	7-1
Module Objectives	7-1
Data Protection Bundles	7-2
Data Protection Planning Considerations	7-3
Planning Considerations: Rolling Disaster	7-5
Planning Considerations: Data Consistency	7-8

Planning Considerations: RPO and RTO	7-9
Planning Considerations: Compare Remote Copy Solutions	7-11
Planning Considerations: Resynchronization Vulnerability	7-13
Planning Considerations: Recovery	7-14
Failover: Planning Considerations	7-15
Failover: horctakeover Command	7-16
Failover Using CCI	7-21
Best Practices	7-22
Best Practices: Monitoring	7-24
Best Practices: TrueCopy Sync Consistency Group	7-25
Best Practices: Data Protection Configuration	7-26
Best Practices: ShadowImage At-Time Split	7-27
Best Practices Review	7-28
Module Review	7-29
8. THREE DATA CENTER OPERATIONS	8-1
Module Objectives	8-1
Purpose of 3DC Replication	8-2
Configurations Summary	8-3
Configurations: 2DC Cascade PassThru	8-4
Configurations: 3DC Cascade	8-8
Configurations: Cascade Mirror IDs	8-9
Configurations: Multi-target	8-10
Configurations: Multi-target Bitmaps	8-11
Configurations: 4DC	8-12
Configurations: Multi-target with Third Link and Delta Resync	8-14
Configurations: Three-Link Multi-target Bitmaps	8-15
Configurations: Mirror Unit Numbers	8-16
Configurations	8-17
Operations	8-18
Disaster Recovery with 3DC	8-19
3DC Disaster Recovery	8-25
Module Review	8-26
9. DELTA RESYNCHRONIZATION	9-1
Module Objectives	9-1
Concepts: Multi-target 3DC Review	9-2
Specifications	9-10
Configuration	9-12
Sequence of Delta Resync	9-21
CCI Support for Delta Resync	9-25
Review	9-26
Module Review	9-26
10. DATA TRANSPORT TECHNOLOGIES	10-1
Module Objectives	10-1
Supported Link Topologies	10-2
Long-Distance Transport Technology	10-7
Flow Control – Basics	10-12
Flow Control – Buffer Credits	10-13
SAN Extension Technology: Dark Fiber	10-15
SAN Extension Technology: SONET / SDH	10-16
SAN Extension Technology: CWDM	10-17
SAN Extension Technology: DWDM	10-18
SAN Extension Technology: FCIP	10-21
SAN Extension Technology: Comparison	10-25
Module Review	10-26

11. HITACHI REPLICATION MANAGER OVERVIEW	11-1
Module Objectives	11-1
Overview of Replication Manager User Interface	11-2
Features	11-6
Features: Mainframe	11-7
Features: Manage Security	11-9
Operations: Create Command Device	11-11
Operations: Configure Remote Paths	11-13
Operations: V-VOL Configuration Support	11-14
Pair Configuration Wizard: Task Management	11-15
Configuration: Task Management — Workflows	11-17
Change Pair Status Wizard	11-19
Monitoring: My Copy Groups	11-20
Monitoring: Refresh Configuration and Status	11-21
Monitoring: Refreshing Configuration Information	11-23
Alerts	11-25
Managing Alerts: Metrics for Remote Copy	11-32
Managing Events: Event Logs	11-34
Error Message Example	11-36
Application Support	11-38
Application Support: Create Replica Wizard	11-40
Module Review	11-42
12. MXN CONSISTENCY GROUPS	12-1
Module Objectives	12-1
Licensing	12-2
Concepts	12-3
Concepts: Configuration Examples	12-4
Concepts	12-6
CCI Processes for Managing	12-7
Architecture	12-8
Architecture: Order of Data Update	12-9
Managing: Sample HORCM Files	12-10
Managing: Group ID	12-11
Managing: Sample HORCM Files	12-12
Managing: Group ID	12-13
Managing: Storage Navigator Display	12-14
Managing: Sample pairedisplay Results	12-15
Managing: Sample pairedisplay Results	12-16
Managing: pairsplit Command	12-18
Restrictions	12-20
Module Review	12-21

YOUR NEXT STEPS

GLOSSARY

EVALUATING THIS COURSE

Introduction

Welcome and Introductions

- Introductions
 - Name
 - Position
 - Professional skills
 - Expectations from the course

Course Description

This five day instructor led course provides details of the Hitachi Universal Replicator and key tasks necessary to implement remote replication for Hitachi Virtual Storage Platform, Hitachi Universal Storage Platform[®] V, and Hitachi Universal Storage Platform VM systems. Learners are introduced to the Universal Replicator software specifications and operations and the corrective actions required for troubleshooting.

The course covers the use of Universal Replicator and Hitachi TrueCopy[®] Synchronous in three data center (3DC) configurations, as well as Universal Replicator Delta Resync function. Details on the use of MxN Consistency (CT) Groups, Remote Replication Data Transport, and Replication Manager are also presented.

The classroom sessions are supported by lab exercises where the participants will learn to install, configure, and using Universal Replicator, 3DC, 3DC Delta Resync, and MxN CT Groups on Hitachi storage platforms.

Prerequisites

- Prerequisite Coursework (recommended)
 - TSI1315 - Hitachi Enterprise Software Solutions
- Supplemental Coursework (recommended)
 - CSI0147 - Hitachi Enterprise Storage Replication
- Prerequisite Knowledge:
 - Basic knowledge of Hitachi enterprise storage systems
 - Familiarity with Storage Navigator and Hitachi Command Control Interface

Course Objectives

- Upon completion of this course, the learner should be able to:
 - Describe Hitachi Universal Replicator key features and benefits
 - Describe Universal Replicator components and specifications
 - Describe Universal Replicator journals, including sizing requirements and restrictions on journal volumes
 - Identify Universal Replicator performance factors
 - Use Storage Navigator to configure Universal Replicator replication links, remote storage systems, and journals
 - Issue Storage Navigator commands for Universal Replicator operations

- Upon completion of this course, the learner should be able to (*continued*):
 - Set up Command Control Interface (CCI) command devices, create CCI configuration files and operate Universal Replicator with CCI commands
 - Identify data protection best practice techniques
 - Configure and operate 3DC Multi-target
 - Configure and operate Universal Replicator Delta Resync feature
 - Discuss Remote Replication Data Transport considerations
 - Describe the features and use of Replication Manager software
 - Create and use Universal Replicator MxN Consistency Groups

Course Topics

Presentations		Activities	
1. Overview			
2. Specifications			
3. Architecture and Internal Operations			
4. Storage Navigator Configuration			
5. Storage Navigator Pair Operations		1. Hitachi Universal Replicator Software Operations	
6. Hitachi Command Control Interface Configuration and Operations			
7. Data Protection		2. Managing Universal Replicator Software with Command Control Interface (CCI)	
8. Three Data Center Operations		3. Universal Replicator Software Three Data Center	
9. Delta Resynchronization		4. Universal Replicator Delta Resync Operations	
10. Data Transport Technologies			
11. Hitachi Replication Manager Overview			
12. MxN Consistency Groups		5. Universal Replicator MxN Consistency Group Operations	

7

Learning Paths

- Are a path to professional certification
- Enable career advancement
- Are for customers, partners, and employees
 - Available on HDS.com, Partner Xchange and HDSnet
- Are available from the instructor
 - Details or copies



7

HDS.com: <http://www.hds.com/services/education/>

Partner Xchange: <https://extranet.hds.com/http://aim.hds.com/portal/dt/>

HDSnet: http://hdsnet.hds.com/hds_academy/

Please contact your local training administrator if you have any questions regarding Learning Paths or visit your applicable website.

HDS Academy Is on Twitter and LinkedIn



Follow the HDS Academy on [Twitter](#) for regular training updates.



[LinkedIn](#) is an online *community* which enables students and instructors to actively participate in online discussions related to Hitachi Data Systems products and training courses.

These are the URLs for Twitter and LinkedIn:

<http://twitter.com/HDSAcademy>

http://www.linkedin.com/groups?gid=2972602&trk=hb_side_g

Terms and Acronyms

- **LDEV** = Logical partition of an Array Group by emulation type. Size depends on Emulation.
- **LUN** = LDEV that is mapped to a port
- **Emulation** = Partitioning of Array group into OPEN-X LDEVs
 - Only one Emulation Type per Array Group is allowed.
- **LCU** = Logical Control Unit
 - Number of LCUs is dependent on cache size, number of Array Groups installed.
 - LDEVs are assigned to Logical Control Units at Install Time.
 - Max 256 LDEV per LCU
- **SSID** = Storage Subsystem ID — Unique ID assigned to each LCU

10

- **FED** = Front-End Director = Channel Adapter = CHA
- **BED** = Back-End Director = ACP Pair = Array Control Processor = DKA = Disk Adapter
- **Array Group** = RAID Group = Parity Group
 - 4 or 8 HDD of RAID 1, or RAID 5, or RAID 6
- **DKC** = Disk Controller
- **DKU** = Disk Unit = Disk Frame
- **HDD** = Hard Disk Device = disk drive

11

Replication Terminology

- **Main Control Unit (MCU, M-DKC)**
 - Contains Primary Volumes (P-VOLs) and Master Journal Volumes (M-JNL)
- **Remote Control Unit (RCU, R-DKC)**
 - Contains Secondary Volumes (S-VOLs) and Restore Journal Volumes (R-JNL)
- **P-VOL (Primary Volume)**
 - Active, online LUN (in MCU)
- **S-VOL (Secondary Volume)**
 - Remote copy of the P-VOL (in RCU)
- **Replication Links (Paths)**
 - Fibre Channel connection to carry replication control and data
- **Journal Volume**
 - Store differential (update) data if necessary
- **Journal Groups**
 - Contain both data volumes and journal volumes
 - Allow Consistency Group operation on multiple data volumes with one command

12

1. Overview

Module Objectives

- Upon completion of this module, the learner should be able to:
 - Define the purpose of Hitachi Universal Replicator software
 - Describe two data center (2DC) and three data center (3DC) configurations
 - Identify the key features of Universal Replicator
 - Discuss Universal Replicator group associations
 - Describe Universal Replicator documentation
 - Define and review Universal Replicator components

2

Purpose

- Hitachi Universal Replicator
 - Provides an asynchronous replication solution
 - Business continuity/disaster recovery
 - Data migration
 - Provides one-to-one copies at any distance between Hitachi enterprise storage systems
 - Supports three data center replication when paired with Hitachi TrueCopy® Synchronous Replication software or with another instance of Universal Replicator (3DC 2 Async configuration)
 - Supports 4DC configurations
 - Allows production volumes to stay online during normal Universal Replicator operations

3

Universal Replicator is a disaster recovery solution for large amounts of data which span multiple volumes. The Universal Replicator group-based update sequence consistency solution enables fast and accurate database recovery, even after a *rolling disaster*, without the need for time-consuming data recovery procedures.


During normal data replication operations, the primary data volumes remain online to all hosts and continue to process both read and write I/O operations. In the event of a disaster or system failure, the secondary copy of data can be rapidly invoked to allow recovery with a very high level of data integrity. Universal Replicator can also be used for data duplication and migration tasks

Once Universal Replicator operations are established, duplicate copies of data are automatically maintained asynchronously. Universal Replicator enables fast and accurate database recovery even after disasters, such as earthquakes, without the time-consuming data recovery procedures.

Supported Hitachi enterprise storage systems:

- Hitachi Virtual Storage Platform
- Hitachi Universal Storage Platform® V or VM
- Hitachi Universal Storage Platform or Hitachi Network Storage Controller

Two Data Center Overview



In the diagram and following pages,

VSP stands for Hitachi Virtual Storage Platform.

USP stands for Hitachi Universal Storage Platform.

CCI stands for Hitachi Command Control Interface (RAID Manager).

AMS stands for Hitachi Adaptable Modular Storage

WMS stands for Hitachi Workgroup Modular Storage

Key Features – Journals

- Journal Cache
 - Cache space in which Universal Replicator can build control information (metadata) and temporarily store overwritten write data blocks
- Journal Volumes
 - Offline physical LDEVs on storage system
 - OPEN-V required – Minimum two 12GB volumes
 - Used on MCU (primary) and RCU (remote) storage arrays
 - Temporarily store journal data (metadata and data) during replication
 - Provide buffer space for journal data after communication failure
 - Can be added to a group non-disruptively (*See note.*)
- Journal volumes and data volumes can be assigned to different CLPRs.
- For VSP, journal and data volumes should be managed by a dedicated Virtual Director.

5

When Universal Replicator is used, data to be copied will be temporarily stored in *journal volumes*, which are a type of physical logical devices. Universal Replicator enables you to configure and manage highly reliable data replication systems by using journal volumes to reduce chances of suspension of copy operations; copy operations can be suspended due to restrictions on data transfers from the primary site to the remote site.

The updates (sometimes called *update data*) that will be stored in journal volumes are called *journal data* or sometimes *differential data*. Because journal data will be stored in journal volumes, you can perform and manage highly reliable remote copy operations without suspension of remote copy operations. For example: Even if a communication path between the primary storage system and the remote storage system fails temporarily, remote copy operations can continue as the journal volumes provide buffering capability to maintain the replication until the communication path is recovered.

If data transfer from hosts to the primary storage system is temporarily faster than data transfer between the primary storage system and the remote storage system, remote copy operations between the primary storage system and the remote storage system can continue. Because journal volumes can contain much more *update data* than the cache memory, remote copy operations can continue if data transfer from

hosts to the primary storage system is faster for a relatively long period of time than data transfer between the primary storage system and the remote storage system.

Note: If you add a journal volume when a remote copy operation is in progress (that is, when at least one data volume pair exists for data copying), the metadata area of the journal volume will be unused and only the journal data area will be used.


To make the metadata area usable, you need to split (suspend) all the data volume pairs in the journal group and then restore (resynchronize) the pairs.

Adding journal volumes during a remote copy operation will not decrease the metadata usage rate if the metadata usage rate is high.

Adding journal volumes during a remote copy operation may not change the journal data usage rate until the journal volumes are used.

Key Features – MxN Consistency Group

- MxN Consistency Group (MxN CT Group)
 - Secondary volume consistency is preserved across multiple journal groups (may extend across multiple storage systems).
 - An MxN Consistency Group can consist of multiple storage systems at either primary or secondary sites.
 - Requires Command Control Interface software for creation and operations.



- Delivers asynchronous 2DC support for:
 - Up to four storage systems per site
 - Intermix of VSP, USP V, or USP VM at both sites
 - MxN Consistency Group over four storage systems
 - 32K pairs in a single MxN consistency group spread over four storage systems (see note)
 - Allows split of multiple journal groups as a single consistency group split based on consistency markers (CTQ-Markers) functioning as timestamps
 - Requires CCI version 01-23-03/02 or higher (01-24-03 if using VSP)


Note: 8K pairs per Journal Group X 4 Journal Groups per MxN CT Group X 2 LDKC per storage system = 64K pairs per MxN CT Group

Group Associations

■ Universal Replicator Group associations




Allowed



Mirrored Replication




Not allowed



Illegal Copy direction



Not allowed (Open Systems)




Illegal JNL Grouping

■ Allowable configurations



Allowed



MxN Consistency Group (CT Group)

Documents

- **MK-90RD7032**
 - *Hitachi Virtual Storage Platform Universal Replicator User Guide*
- **MK-96RD624**
 - *Hitachi Universal Storage Platform V, Hitachi Universal Storage Platform VM
Hitachi Universal Replicator User's Guide*
- **MK-90RD7008, MK-90RD7009, MK-90RD7010**
 - CCI documents supporting version 01-24-xx (VSP support)
- For operations with TrueCopy and ShadowImage®
 - **MK-90RD7029** *Hitachi Virtual Storage Platform TrueCopy User's Guide*
 - **MK-90RD7024** *Hitachi Virtual Storage Platform ShadowImage User's Guide*
 - **MK-96RD622** *Hitachi Universal Storage Platform V, Hitachi Universal
Storage Platform VM TrueCopy User's Guide*
 - **MK-96RD618** *Hitachi Universal Storage Platform V, Hitachi Universal
Storage Platform VM ShadowImage User's Guide*

10

If required, the equivalent documentation for the USP/NSC models:

MK-94RD223 *Hitachi TagmaStore® Universal Storage Platform and Network Storage
Controller Universal Replicator User and Reference Guide*

MK-94RD215 *Hitachi TagmaStore Universal Storage Platform and Network Storage
Controller TrueCopy User and Reference Guide*

MK-94RD204 *Hitachi TagmaStore Universal Storage Platform and Network Storage
Controller ShadowImage User's Guide*

Components

- Software Components
 - Equivalent microcode levels on both MCU and RCU
 - Storage Navigator access to both MCU and RCU
 - Universal Replicator license key is included in the Disaster Recovery Bundle which includes TrueCopy Sync.
 - 3DC and UR MxN Consistency Groups requires the Disaster Recovery Extended bundle
 - Universal Replicator license keys installed on MCU and RCU
- Optional, but recommended for disaster recovery environments
 - Hitachi Command Control Interface
 - Host failover software
 - Hitachi Replication Manager

11

2. Specifications

Module Objectives

- Upon completion of this module, the learner should be able to:
 - Identify Hitachi Universal Replicator specifications
 - Describe Universal Replicator volume status conditions

Common Storage System Specifications

Supported Storage and Emulation Types

Item			Supported emulation/spec	Remarks
Supported Storage Types: Virtual Storage Platform Universal Storage Platform V Universal Storage Platform VM Universal Storage Platform Network Storage Controller			• 3990-6E • 2105 • 2107	✓ 3990-6 (Basic) is not supported
RAID level			• RAID-5 • RAID-6 • RAID-1	✓ RAID-5 supports (3D+1P) and (7D+1P). ✓ RAID-6 supports (6D+2P). ✓ RAID-1 supports (2D+2D).
LDEV emulation type	Data volume	OPEN	• OPEN-V	
		M/F	• 3390-1,-2,-3 • 3390-3R,-9,-L, -M	
	Journal volume	OPEN	• OPEN-V	
		M/F	• OPEN-V	

3

For Open Systems environments, Controller Emulation has no effect but must be set.

Volume Specifications

- Universal Replicator usage with other features
 - TrueCopy Synchronous – P-VOLs and S-VOLs
 - ShadowImage – P-VOLs and S-VOLs
 - Copy on Write Snapshot – P-VOLs only
 - TrueCopy Asynchronous – No usage with Universal Replicator
 - Volume Migrator – If mapped to ports
 - Data Retention Utility – Certain volumes—refer to Data Retention Utility documentation
 - LUSE (Logical Unit Size Expansion) volumes – If both P-VOL and S-VOL have the same structure
 - Cache Residency Manager volumes
 - Server Priority Manager volumes
 - Externalized volumes (virtual LUNs)
 - Hitachi Dynamic Provisioning volumes

4

Notes on LUSE usage:

- Two LUSE volumes can be assigned to a Universal Replicator (UR) pair. Both of the LUSE volumes that are assigned to primary and secondary data volumes must consist of the same number of LDEVs and must have the same capacity.
- If you want to perform LUSE operation to primary or secondary data volumes in an existing UR pair, you must delete the pair first to return the volumes to SMPL status. For detailed information about LUN Expansion (LUSE), see the *LUN Expansion User's Guide*.

- **Maximum number of pairs:**
 - Virtual Storage Platform: Maximum pairs = 32,768
 - Universal Storage Platform V and VM: Maximum pairs = 32,768
 - Universal Storage Platform and NSC: Maximum pairs = 16,384

Item	Support Specifications
Emulation type	Same or different types of emulation type is available. Only OPEN-V can be used for journal volumes.
Volume capacity	Same or different amount of capacity is available.
CLPR	Journal volumes and data volumes in the same journal group can belong to different CLPRs. Journal volumes must belong to the same CLPR. A primary journal group and the corresponding secondary journal group need not belong to the same CLPR.

5

Journal volumes and data volumes in the same journal group can belong to different CLPRs. All journal volumes must belong to a single CLPR and all data volumes must belong to a single CLPR.

NSC stands for Hitachi Network Storage Controller.

■ Journal Group Specifications for Hitachi enterprise storage systems

Item	Support specifications
Number of journal groups	Up to 256 journal groups (0 to 255) per one storage system The recommended number of journal groups is up to 16.
Number of data volumes in journal group	Up to 8,192 Note that whether or not the maximum number of data volumes are available for a journal group of the local storage system depends on the number of data volumes that can be registered to a journal group of the paired storage system.
Number of journal volumes in journal group	Up to 64
Number of Mirror IDs	Up to 4 (0 to 3) If TrueCopy Sync. uses No. 0, No. 1 to 3 are available for UR.

Specifications: System Option Modes

■ Contact Tech Support for information

Note: Changes requires tech support involvement

Mode	Category	Function	Default	MCU/RCU
460	TrueCopy, TrueCopy for z/OS ShadowImage, ShadowImage for z/OS Universal Replicator, Universal Replicator for z/OS Volume Migration FlashCopy v1, v2 COW Snapshot Dynamic Provisioning	When turning off PS, the control information of the following PPs stored in shared memory will be backed up in SVP. After that, when performing volatile PS ON, the control information will be restored into shared memory from SVP. TrueCopy, TrueCopy for z/OS, ShadowImage, ShadowImage for z/OS, Volume Migration, FlashCopy v1, FlashCopy v2, Universal Replicator, Universal Replicator for z/OS, COW Snapshot, Dynamic Provisioning. Setting Mode 460 to ON is required to enable the function. Note: 1. Set mode 460 to ON when using any of the PPs listed above. 2. In case of mode 460 ON, the completion time of PS-OFF for SM backup or the completion time of volatile PS-ON after SM recovery takes up to 30 minutes. 3. When using power monitoring devices (PCI, etc.), the monitoring time for PS-OFF/ON must be set to the maximum of 30 minutes. 4. When the storage system is powered off automatically at power outage while destage mode is ON, SM is not backed up to the SVP even though mode 460 is set to ON. 5. In the case of using Dynamic Provisioning: With DKCMAIN version 60-02-04-00/00 and higher, the DP management information is stored in a dedicated area in the pool in case data are lost from the SM. However, restoring the data from the dedicated area in the pool may take more time than restoring the data from the SVP. Therefore, setting mode 460 to ON is recommended to enable the data backup and recovery function from the SVP.	OFF	MCU/RCU

7

■ System Option Modes for Universal Replicator

Mode	Category	Function	Default	MCU/RCU
448	Universal Replicator Universal Replicator for z/OS	Mode 448 = ON: (Enabled) If the SVP detects a blocked path, the SVP assumes that an error occurred and immediately splits (suspends) the mirror. Mode 448 = OFF: (Disabled) If the SVP detects a blocked path and the path does not recover within the specified period of time, the SVP assumes that an error occurred and splits (suspends) the mirror. Note: The mode 448 setting takes effect only when mode 449 is set to OFF.	OFF	-
449	Universal Replicator Universal Replicator for z/OS	Detecting and monitoring path blockade between MCU and RCU of UR/URz. Mode 449 = ON: Detecting and monitoring of path blockade will NOT be performed. Mode 449 = OFF: Detecting and monitoring of the path blockade will be performed.	OFF	-

■ System Option Modes for Universal Replicator

Mode	Category	Function	Default	MCU/RCU
466	Universal Replicator, Universal Replicator for z/OS	For UR/URz operations it is strongly recommended that the path between main and remote storage systems have a minimum data transfer speed of 100 Mbps. If the data transfer speed falls to 10 Mbps or lower, UR operations cannot be properly processed. As a result, many retries occur and UR pairs may be suspended. Mode 466 is provided to ensure proper system operation for data transfer speeds of at least 10 Mbps. Mode 466 = ON: Data transfer speeds of 10 Mbps and higher are supported. The JNL read is performed with 4-multiplexed read size of 256 KB. Mode 466 = OFF: For conventional operations. Data transfer speeds of 100 Mbps and higher are supported. The JNL read is performed with 32-multiplexed read size of 1 MB by default. Note: The data transfer speed can be changed using the Change JNL Group options.	OFF	–

9

■ System Option Modes for Universal Replicator

Mode	Category	Function	Default	MCU/RCU
506	Universal Replicator, Universal Replicator for z/OS	Enables Delta Resync with no host update I/O by copying only differential JNL instead of copying all data. The HUR Differential Resync configuration is required. - Mode 506 = ON (default): - Without update I/O: Delta Resync is enabled. - With update I/O: Delta Resync is enabled. - Mode 506 = OFF: - Without update I/O: Total data copy of Delta Resync is performed. - With update I/O: Delta Resync is enabled. Note: Even when mode 506 is set to ON, the Delta Resync may fail and only the total data copy of the Delta Resync function is allowed if the necessary journal data does not exist on the primary subsystem used for the Delta Resync operation.	ON	MCU/RCU

10

■ System option modes for Universal Replicator

Mode	Category	Function	Default	MCU/RCU
690	Universal Replicator and Universal Replicator for z/OS	This option is used to prevent Read JNL or JNL Restore when the Write Pending rate on RCU exceeds 60% as follows: - When CLPR of JNL-Volume exceeds 60%, Read JNL is prevented. - When CLPR of Data (secondary)-Volume exceeds 60%, JNL Restore is prevented. Mode 690 = ON: Read JNL or JNL Restore is prevented when the Write Pending rate on RCU exceeds 60%. Mode 690 = OFF: Read JNL or JNL Restore is not prevented when the Write Pending rate on RCU exceeds 60% (the same as before). Note: - This mode can be set online. - This mode should be set per customer's requests. - If the Write Pending status long keeps 60% or more on RCU, it takes extra time for the initial copy to be completed by making up for the prevented copy operation. - If the Write Pending status long keeps 60% or more on RCU, the pair status may become Suspend due to the JNL-Vol being full.	OFF	RCU

11

■ System option mode for Universal Replicator

Mode	Category	Function	Default	MCU/RCU
707	TrueCopy (open), Universal Replicator (open)	In a cascading configuration using TrueCopy (TC) and Universal Replicator (UR), by setting mode 707 to ON and using the registered JNL group, the copied data from the P-site can be stored in the journal volume in the I-site without writing the P-site data onto the data volume in the I-site. If P-site fails, after the R-site has taken over the production, all the data is copied back from the R-site to the P-site. Mode 707 = ON: The data on the P-site is not written onto the data volume in the I-site. Mode 707 = OFF (default): The data on the P-site is written onto the data volume in the I-site. Notes: - During the operation, this system option mode must remain ON at the I-site. - Any UR pair cannot use the storage system for usual UR operations when the system option mode is set to ON. - As for a TC pair combined with a UR pair in the JNLG that is registered after the system option is set to ON, the following operations are unavailable: - TC pair takeover, or UR pair takeover (when the P-site is recovered, all the data in the R-site is copied onto the data volume in the P-site). - Combination with a ShadowImage pair (I-site). The data is not copied. - Update by a host I/O to the data volume in the I-site. - Switching to the usual cascading configuration using TC and UR. (To switch to the usual configuration, the UR JNLG must first be deleted.) - After JNLG registration, even if the mode status is switched (ON/OFF), the status remains the same as that at the JNLG registration. - If the UR pair suspends, failure suspend occurs to the combined TC pair. - The TC/UR pair configuration is cascading, but the operation differs from the cascading operation.	OFF	MCU/RCU

12

Pair Volume Status

■ Volume status conditions

Pair Status	Description	P-VOL Volume Access	S-VOL Volume Access
SMPL	This volume is not currently assigned to a Universal Replicator data volume pair. This volume does not belong to the journal group. When this volume is added to a Universal Replicator data volume pair, its status will change to COPY.	Read / Write	Read / Write
COPY	The initial copy operation for this pair is in progress. This data volume pair is not yet synchronized. When the initial copy is complete, the status changes to PAIR.	Read / Write	NA
PAIR	This data volume pair is synchronized. Updates to the primary data volume are duplicated on the secondary data volume.	Read / Write	NA
PSUS	This data volume pair is suspended, either because the user has split the pair (pairsplit-r), or because the user has deleted the pair from the RCU (pairsplit-S). For Universal Replicator pairs, the MCU and RCU keep track of any journal data that were discarded during the pairsplit-r operation. While a pair is split, the MCU and RCU keep track of the primary data volume and secondary data volume tracks which are updated.	Read / Write	Read Only; Read and Write if write option is enabled

13

PSUS Status: Pair Suspended Synchronized

This data volume pair is not synchronized, because the user has split this pair (**pairsplit-r**), or because the user has released this pair from the remote storage system (**pairsplit-S**). For Universal Replicator pairs, the primary storage system and remote storage system keep track of any journal data that were discarded during the **pairsplit-r** operation. While a pair is split, the primary storage system and remote storage system keep track of the primary data volume and secondary data volume tracks which are updated.

When you split a pair from the primary storage system, the primary storage system changes the status of the primary data volume and secondary data volume to PSUS. When you split a pair from the remote storage system, the remote storage system changes the status of the secondary data volume to PSUS. The primary storage system detects this (if path status is normal) and changes primary data volume status to PSUS.

When you release a pair from the remote storage system, the remote storage system changes the status of the secondary data volume to SMPL. The primary storage system detects this (if path status is normal) and changes primary data volume status to PSUS. You must release the pair from the primary storage system in order to change the primary data volume status to SMPL.

■ Volume status conditions (*continued*)

Pair Status	Description	P-VOL Volume Access	S-VOL Volume Access
PSUE	This data volume pair is suspended by MCU or RCU due to an error condition. MCU and RCU keep track of any journal data that were discarded during the suspension operation. The MCU keeps track of the primary data volume tracks which are updated while the pair is suspended.	Read / Write; Read Only if Fenced	Read Only
PFUL	Universal Replicator monitors the total amount of data in the journal volume. If the amount of data exceeds the threshold (80%), the pair status changes from COPY or PAIR to PFUL. The write data that inflows is monitored during the specified time (Data Overflow Watch). The time period of monitoring can be set using the Storage Navigator software (default setting is 60 seconds). <i>Note:</i> The PFUL status is displayed by CCI (Command Control Interface). Storage Navigator displays this status as PAIR.	Read / Write	Read Only

14

PSUE: Pair Suspended Error

This data volume pair is not synchronized, because the primary storage system or remote storage system has suspended the pair due to an error condition. For Universal Replicator pairs the primary storage system and remote storage system keep track of any journal data that were discarded during the suspension operation. The primary storage system keeps track of the primary data volume tracks which are updated while the pair is suspended.

If the primary storage system detects a Universal Replicator suspension condition (see Suspension Condition), the primary storage system changes the secondary data volume status to PSUE.

If the primary storage system detects a Universal Replicator suspension condition, the remote storage system changes the secondary data volume status to PSUE. The primary storage system detects this (if path status is normal) and changes primary data volume status to PSUS.

▪ Volume status conditions (*continued*)

Pair Status	Description	P-VOL Volume Access	S-VOL Volume Access
PFUS	If Journals fill and Data Overflow Watch timer period expires (INFLOW CONTROL ON), the pair status changes from PFUL to PFUS, and the pair is suspended. <i>Note:</i> The PFUS status is displayed by CCI (Command Control Interface). Storage Navigator displays this status as PSUS.	Read / Write	Read Only; Read and Write if write option is enabled
SSWS	Result of SUSPEND SWAP (pairsplit –RS) operation in RCU which conditions S-VOLS for takeover. The data can be written into the secondary data volume that is reassigned from the primary data volume during processing of resynchronization (Takeover). <i>Note:</i> The SSWS status is displayed by CCI (Command Control Interface). Storage Navigator displays this status as PSUS or PSUE.	Read only	Read / Write

15

PFUS (Pair Full Suspended)


SSWS (Secondary Swap Suspended) – Takeover function is active.

▪ Volume status conditions (*continued*)

Pair Status	Description	P-VOL Volume Access	S-VOL Volume Access
Suspending	This pair is in transition from PAIR or COPY to PSUS/PSUE. When the split/suspend pair operation is requested, the status of all affected pairs changes to Suspending. When the split/suspend operation is complete, the status changes to PSUS/PSUE.	Read/Write	NA
Deleting	This pair is in transition from PAIR, COPY, or PSUS/PSUE to SMPL. When the pairsplit-S operation is requested, the status of all affected pairs changes to Deleting. When the pairsplit-S operation is complete, the status changes to SMPL.	Read/Write	NA

■ Transition from PAIR to PFUS

- If Journal Data exceeds 80%, group status changes to PFUL.
- Depending on Inflow Control Setting, if Journal Data attains 100%, or if it exceeds 80% for user-set period of time, group status changes to PFUS (suspended).
- PFUL = Pool (Journal) full.
- PFUS = Pool (Journal) full, suspended.



17

Universal Replicator monitors the amount of journal data. If the amount of data exceeds the threshold (80%), the pair status changes to PFUL. If the amount of journal data exceeds the threshold for a certain period of time, the volume status changes to PFUS, and the group is suspended.

Module Review

1. True or False. Any OPEN emulation type can be used for journal volumes.
2. How many journal volumes can be created in one Journal Group?
3. What does the status PSUE signify?
4. What does the status PFUS signify?

3. Architecture and Internal Operations

Module Objectives

- Upon completion of this module, the learner should be able to:
 - Describe how Universal Replicator uses bitmap areas during pair operations
 - Describe how Universal Replicator moves data between storage systems
 - Demonstrate how journal volumes are utilized to replicate data
 - Describe the structure of a journal volume
 - Identify Universal Replicator performance considerations and restrictions
 - Identify best practices for specifying journal volumes
 - Identify journal cache specifications
 - Describe Universal Replicator MxN Consistency Group functionality
 - Discuss configuration planning steps and considerations

2

Bitmap Areas

- Bitmap tracking mechanism is invoked when pairs are created.
 - In remote replication, differential bitmaps are now called *bitmap areas*.
 - Created in control memory (VSP), or shared memory (USP V)
 - The **Base Journal** process reads P-VOL cylinders, creates metadata, and sends data to RCU.
 - Changes to P-VOL data while cylinders are being copied are managed by bitmap areas.
- Bitmaps are also invoked when pairs are split or suspended.
 - Denotes **cylinders** changed by host I/O.
 - When data volumes are resynchronized, P-VOL and S-VOL bitmaps are merged and changed cylinders are sent to remote storage system.

3

Differential data (updates by write I/Os during Initial Copy or while pairs are in PSUS or PSUE status) are tracked by changes in volume's bitmap. When a split/suspended pair is resumed (pairresync), the primary storage system merges the primary data volume and secondary data volume bitmaps, and the differential data is copied to the secondary data volume.

Note: The number of bitmap areas affects the maximum possible number of pairs that can be created in the disk storage system.

- Bitmap areas during Base Journal
- **Best Practice:** Use Bitmap 1 for the UR pair, leaving the other bitmaps available for TrueCopy and Delta Resync if desired.

Primary Volume Bitmap (X denotes cylinders changed by host I/O during Base Journal)

X | X | | | X | X | | | | | | | . . .

Secondary Volume Bitmap (No host I/O access allowed to S-VOL until pairs are split)

| | | | | | | | | | | | | | | | . . .

4

Note: Differential bitmaps are initially set to all ones. As each cylinder is copied during Base Journal process, the bit corresponding to that cylinder is changed to zero.

- Bitmap usage when pairs have been split.
- If desired, write I/O on S-VOL during pairsplit is allowed.
 - When resync occurs:
 - MCU merges primary and secondary bitmaps.
 - All changes are sent to RCU under control of P-VOL bitmaps in a process similar to Base Journal operation.

Caution: Resynchronization will overwrite any changes made to the S-VOL while pairs were split.

5

For additional flexibility, Universal Replicator provides a secondary data volume write option (S-VOL Write) which enables write I/O to the secondary data volume of a split Universal Replicator pair. The secondary data volume write option can be selected by the user during the pairsplit-r operation and applies only to the selected pairs. The secondary data volume write option can be accessed only when you are connected to the primary storage system.

When you resync a split Universal Replicator pair which has the secondary data volume write option enabled, the remote storage system sends the secondary data volume bitmap to the primary storage system, and the primary storage system merges the primary data volume and secondary data volume bitmaps to determine which cylinders are out-of sync. All changed cylinders are then sent from P-VOL to S-VOL. This ensures proper resynchronization of the pair. Note that this will overwrite any changes to the S-VOL that occurred while pairs were split.

▪ **Bitmap area usage during *resynchronization***

When pairs are resynchronized, P-VOL and S-VOL bitmaps are merged in MCU and all changed cylinders are copied from P-VOL to S-VOL.

P-VOL Bitmap

X | X | | | X | X | | | | | | | . . .

S-VOL Bitmap

| X | X | | | | | | | | X | | | . . .




Merged Bitmaps

X | X | X | | | X | X | | | X | | | | . . .

Bitmap Areas – Mirror Unit Numbers

- Bitmap areas as mirror IDs (Mirror Unit Numbers)
 - Each P-VOL and S-VOL can have up to four bitmap areas.
 - Bitmaps are separate from ShadowImage differential bitmaps and do not share mirror IDs.



7

The example shows the mirror ID association of a Universal Replicator pair with multiple ShadowImage mirrors. Note that Universal Replicator has four Mirror IDs per volume while ShadowImage has three.

Moving Data Between Storage Systems

- **Journals** provide buffer space to greatly enhance UR capability to survive link outages compared to other async replications.
 - Consistency Group relationship exists between M-JNL and R-JNL.
 - Consistency Group ID is assigned when the first pair is created.
 - Can be the MxN Consistency Group ID if properly defined in CCI
 - Journal operations at the primary storage system.
 - Split all pairs in a journal group
 - Resume all suspended pairs in a journal group
 - Release all pairs in a journal group.
 - Journal operations at the remote storage system.
 - Split (pairsplit-r) all pairs in a journal group
 - Release (pairsplit-S) all pairs in a journal group regardless of their consistency status
 - If managed properly, UR guarantees multi-volume consistency on S-VOLS.

- **Base Journal operation**
 - Differential bitmaps are created.
 - Cylinder-by-cylinder copy of P-VOL data is initiated.
 - Host write operations to P-VOL are managed by bitmaps.
- **Journal Obtain operation**
 - Sends new write data to journal volume, in preparation for transport
 - Two update modes
 - Asynchronous – Normal update mode
 - Synchronous – Only invoked with Inflow Control ON (default)
- **Journal Copy operation**
 - R-DKC requests that primary journal (P-JNL) data be sent to remote storage system.
- **Journal Restore operation**
 - Updates S-VOL with data stored in secondary journal (S-JNL)
 - Maintains write order consistency of data on S-VOLs

9

Journal groups enable update sequence consistency to be maintained across a journal group of volumes. The primary data volumes and secondary data volumes of the pairs in a journal group must be located within one physical primary storage system and one physical remote storage system (1-to-1 requirement).

When more than one data volume is updated, the order that the data volumes are updated is managed within the journal group that the data volumes belong to. Consistency in data updates is maintained among paired journal groups. UR uses journal groups to maintain data consistency among data volume.

■ **Base Journal** process

- Equivalent to Initial Copy in TrueCopy Heterogeneous Remote Replication
- Initiates Universal Replicator volume pair synchronization
- Copies P-VOL data to secondary volume (S-VOL)
 - Nocopy option is available if desired (*see note*).
- **MCU** creates pointers (metadata) to P-VOL data as it is read. Metadata is stored in **Journal Cache** and destaged to the M-Journal volumes.
- The **Base Journal** process continues reading P-VOL cylinders, creating metadata, and sending data until all cylinders are read and sent to RCU.
- Bitmap mode is operational during the **Base Journal** process.
- When all cylinders have been copied (including all changes that occur while copying), pair status changes to PAIR.

10

Initial copy operations synchronize data in the primary data volume and data in the secondary data volume. They are performed independently from the host I/Os. They are performed when you create a data volume pair or when you resynchronize a suspended pair. The initial copy operation copies the Base Journal data that is obtained from the primary data volume at the Primary storage system to the Secondary storage system.

The Primary storage system reads all data of the primary data volume as the Base Journal data, in sequence. The Base Journal contains a replica of the entire data volume or a replica of updates to the data volume.

The Base Journal process denotes changes in primary volume by noting changed cylinders in the Bitmap Areas. Once all changed cylinders are copied, PAIR status is declared. When a data volume pair is suspended, the primary volume Bitmap Areas is again used to note changes to the primary volume.

Note: You can specify **None** as the copy mode for initial copy operations. If the **None** mode is selected, only volume identification information will be copied. Full Base Journal operations will not be performed. The **None** mode must be used only when you are sure that data in the primary data volume is identical to data in the secondary data volumes, or when only the volume identification information is desired.

▪ **Journal Obtain process**

- Equivalent to Update Copy in TrueCopy
- Begins after Base Journal is complete
 - Write I/O from host to P-VOL initiates **Journal Obtain** process.
 - Data is stored on P-VOL.
 - Metadata (control data) created in **Journal Cache** for each update
 - Metadata and update queued for destage to the primary journal volumes (**M-JNL Group**).
 - Metadata includes **Sequence numbers** (and **CTQ-Markers furnished by CCI** if MxN Consistency Group is specified) to preserve write order consistency across all volumes in the R-JNL groups.
 - **Journal Obtain** notification is sent to RCU.
 - Metadata and data blocks are held in cache as long as possible before destaging to M-JNL volumes.

11

▪ **Journal Obtain process (*continued*)**

- Contingent upon **Inflow Control** setting, Journal Obtain operates in two modes when Journal Cache Threshold or Journal Volume Threshold is reached.
 - **Asynchronous Mode** – Normal operation mode
 - Journal cache utilization and journal volume utilization are below threshold level.
 - **Inflow Control** is OFF.
 - Journal groups will suspend if journal cache threshold or journal volume threshold is exceeded. Host I/O will *not* be delayed.
 - **Synchronous mode** – Invoked when:
 - **Inflow Control** is ON.
 - Journal cache is full or journal volumes are full due to link failure or host write I/O rate for data to be replicated exceeding the MCU/RCU transfer rate.
 - Groups will be suspended if Overflow Watch Timer value is exceeded.
 - Will delay host I/O by the access time required to write the journal data to **Journal Volume**

12

▪ **Journal Copy** process

- Upon receipt of Journal Obtain notification, secondary storage system (RCU) responds with **Read Journal** Command to pull data from primary storage system (MCU).
- MCU responds by sending journal data (from cache if possible) as the **Journal Copy**.
 - Data is sent with **Sequence Numbers** (and CTQ-Markers if MxN Consistency Group is set).
- Sent data is received in **RCU Journal Cache**, waiting to be settled.
 - Data may be stored on secondary journal volumes (**R-JNL**) if necessary because of journal cache constraints.
- **Read Journal** commands are sent repeatedly to primary storage system until MCU sends highest sequence number.


13

▪ **Journal Restore** process

- After data is received in **RCU Journal Cache**, RCU initiates **Journal Restore**. Updates are reassembled by sequence number.
- RCU notifies MCU of highest sequence number, along with the number of updates received.
- Sequence numbers are compared to insure that all updates sent by MCU were in fact received by RCU.
- After successful comparison of sequence numbers between RCU and MCU:
 - RCU **discards** metadata and any retained R-Journal entries.
 - MCU **discards** entries in M-Journal.

14

- Sequencing guarantees write order fidelity is maintained for each JNL group.



15

- P-VOL receives write command from the host.
- MCU stores the received write data to the cache.
- For each write data, MCU obtains sequence# in the JNL Group and creates the metadata including the sequence#. MCU also allocates the area of JNL VOL and records the location the area to the metadata. (The write data and write order of them are managed with metadata for each JNL Group.)
MCU also creates JNL that can be stored in JNL VOL from write data.
- RCU issues Read Journal Commands and MCU sends the metadata and the JNL as a response to the commands. They can be performed in a different order from the write order.
- RCU stores the metadata and the JNL to the cache.
- RCU sorts the JNL in the write order by referring metadata (sequence#) for each JNL Group and reflects (writes) them to the paired S-VOL according to the write order.


■ Data transfer between MCU and RCU

- **Base Journal data** is read directly from P-VOL.
- **Update Journal data** sent to RCU comes from MCU journal cache memory or is brought in by read I/O from the journal group.
- On receipt of **Read Journal** command, MCU transfers metadata. For transferring metadata, total transfer length is adjusted to a 512 byte block boundary.
- Next, MCU transfers corresponding **JNL Data**. For transferring JNL, CRC (check code) is generated and total transfer length is adjusted to a 512 byte block boundary.
- The whole length of metadata, JNL and CRC cannot exceed **1MB in one Remote I/O operation (RIO)**.
- After completion, MCU waits for sequence number comparison. On sequence number agreement with RCU, both systems release journal data by updating pointers.

16

■ Hitachi Remote I/O operation

- Replication links use RIO.
 - Maximum of 64 updates placed per one write command
 - Up to 1MB transfer length
 - Independent of source/target devices
 - Data packed and sent from different source volumes
 - Simultaneous transfers from a single source volume
- Up to 16 concurrent writes simultaneously on a single link



17

■ Review

- One JNL Group can have up to **64** JNL VOLs.
- Each JNL VOL has one metadata extent and 32 data extents.
- When **MCU** Journal Group fills (or MCU JNL CACHE fills),
 - Synchronous mode is invoked if **Inflow Control** is ON.
 - Groups split if **Inflow Control** is OFF.
- When **RCU Journal Group fills** (or RCU JNL CACHE fills), issuance of **Read Journal** commands is postponed.
 - May cause MCU JNL VOL to fill
- For **Base Journal**, JNL data is not stored on JNL VOL. Data for the initial copy is stored in P-VOL and transferred from P-VOL without JNL VOL.
- MCU releases metadata and journal data blocks after comparing sequence numbers with RCU. (It discards journals.)
- UR Split operation **flushes** journal data to RCU.
 - Status remains **Suspending** or **Deleting** until all journal data in the M-JNL group is sent to the restore journal group.
 - New updates received from host are marked as changes in bitmap areas.

Structure of Journal Volume

■ Overview

- Journal volumes are buffers for write updates that allow Universal Replicator to survive extended periods of reduced link bandwidth or complete link failure without suspending the replication.
 - Journal volumes store journal data (metadata and data blocks created by Universal Replicator Journal Obtain process).
 - Journal volumes can be grouped into sets of 64 (VSP and USP V) or 16 (USP).
- Journal volumes are divided into extents.
 - By default there are 33 extents.
 - One extent is reserved for metadata.
 - 32 extents are reserved for journal data.
- Extents allow Universal Replicator to read data from journal volumes in parallel.


19

In MCU and RCU, journals are stored in journal volumes. One journal group can contain up to 16 journal volumes. A journal volume consists of metadata area and journal data area. The ratio of metadata area to journal data area is currently fixed at 1 to 32. A journal data area is divided into 32 extents and stores journal data. In the metadata extent, metadata is stored sequentially so that multiple metadata with neighboring sequence numbers can be read to the cache memory from the disk.

Journal data are stored in a round-robin manner so that multiple journal data with neighboring sequence numbers can be transferred to RCU in parallel. The journal data can also be read/written to the disk in parallel. In the Journal Copy function, when a MCU receives a Read Journal command from a RCU, the MCU sends the oldest journal (with the lowest sequence number) first.

Journal volumes

- All data volumes within journal group share journal volumes.
- Journal data extents are written in round-robin manner across available JNL volumes.
- Metadata is stored sequentially on each journal volume.
- Starting positions are placed differently.
- Metadata and data extents have fixed size.
- When additional journal volumes are added, extents are redistributed.
 - This allows journal volumes to be added non-disruptively.



20

Journal Volume Metadata Structure

Type	Description
Journal type	Type of journal (for example, base journal or update journal)
Original data storing position	The primary data volume slot number, and the start and end of sub-block number (data length)
LDEV Number (data)	The number of primary data volume that stores the original data
Journal data storing position	The slot number of primary journal volume, and the start sub-block number
Journal sequence number	The sequence number that is assigned when the journal is obtained
LDEV Number (journal)	The volume number of primary journal volume that stores the journal data
CTQ-Marker (for MxN Consistency Group)	Timestamp for Journal creation


21

Performance Considerations

■ Universal Replicator Link Bandwidth

- The total bandwidth between MCU and RCU **must** exceed the average throughput for write I/O (for data to be replicated) from hosts to the MCU.

Note: If there is to be no delay in transmission of data and RPO is to be minimized, then total bandwidth should exceed the peak throughput.




22

Following factors should be considered to realize the larger bandwidth between MCU and RCU than the highest throughput between the hosts and MCU:

- Type of public line (DWDM/T3/ ATM/IP)
- Number of public lines
- Number of fibre channel links
- Maximum number of links = 8 Initiators and 8 RCU Targets per MCU/RCU

Performance Considerations and Restrictions

- Data flow and performance factors of Universal Replicator
 - Throughput of M-JNL group must exceed the peak inflow data rate for data to be replicated.
 - If JNL VOL throughput is less than the peak data rate, journal cache is the only buffer available.
 - Journal group configuration planning is essential to achieve performance demands.
 - Best case: Read Journal command always finds data in MCU cache. (No read access from journal volume is required for Journal Copy.)



23

Best Practices: Journal Throughput

- Calculation of JNL throughput

- To calculate the Journal Group throughput needed, use this formula:

$$\frac{P \times N}{1024} = J$$

Where: **P** = Peak data rate* in kilobytes (for data in the Journal Group)

N = Number of write I/Os per second (for the Journal Group)

J = Journal group data throughput in MB/sec

* For data to be replicated

In general:

- The data transfer speed to the M-DKC journal volumes must exceed the **peak replicated data transfer rate** from hosts to the primary storage system.

Best Practices: Sizing Journal Volumes

- Requirement for JNL VOL Capacity
 - M-journaling will occur when the replicated data inflow rate exceeds available link bandwidth.
 - Journals provide increased resiliency by allowing Universal Replicator to maintain PAIR relationship longer than other asynchronous replications.
 - When link down occurs, all inflow data goes to M-JNL.
 - When links come up:
 - Inflow data shares link bandwidth with the M-JNL updates being sent.
 - This can result in R-JNL utilization increasing.

- JNL VOL Capacity
 - Journal volumes should have enough capacity to absorb the increase in journal data generated by the following conditions:
 - Occasional higher-than-average peak in data rate
 - Occasional link degradation or total loss of links
- Calculation for JNL VOL capacity

$$(P - L) \times T = C$$

Where:


P = Peak replicated data rate (highest aggregate inflow rate in MB/sec)

L = Link bandwidth (outflow rate in MB/sec, M-R throughput)

T = Expected duration of peak input data in seconds (customer specified)

C = Capacity of journal volumes (in MB)

- RCU Journal Group has same throughput requirement.
 - **If:** Throughput of JNL groups in **both** MCU and RCU group exceeds peak data rate for data to be replicated.
 - **Then:** MCU – RCU Link throughput does *not* have to exceed data rate for data to be replicated. Journal volumes will buffer excess data rate.




27

If user does not need a takeover environment, the performance and capacity of R-JNL group can be less than M-JNL group (not recommended).

Best Practices: Journal Volume Physical Characteristics

- JNL volume interference with other volumes in same parity group
 - Data volumes in parity groups used for JNL VOLs: **not recommended**
 - Neither JNL volume nor data volumes will achieve sufficient and stable performance if any data volumes are also in the parity group (see note).
- Recommendations:
 - **Best practice:** Assign no more than one JNL volume from any given parity group to a JNL group.
 - **Second Best Practice:** Two journal volumes per parity group per journal group.
 - While JNL volumes from any given parity group can be assigned to differing JNL groups, monitor JNL volume utilization closely to identify bottlenecks.
 - Ensure parity groups containing journal volumes are distributed across all available BEDs.



28

Note: The data transfer speed of a journal volume depends on the data transfer speed of the Parity group that the journal volume belongs. One Parity group can consist of one or more logical volumes, including journal volumes. Therefore, if frequent access is made to non-journal volumes in a Parity group, relatively fewer accesses can be made to journal volumes in the same Parity group. This can cause a drop in the data transfer speed of journal volumes. To avoid that, you will need to consider relocating the frequently accessed non-journal volumes to another parity group.

- Physical disk type
 - Access to JNL VOL is mixture of sequential write and sequential read.
 - Journal volume access is primarily random.
 - Random access requires seek time and rotational latency to position accessed data.
 - Therefore performance of physical disk (seek time, rotation speed, and transfer speed) affects JNL VOL throughput.
 - 15K RPM HDD are preferable because of faster access time than 10K RPM HDD.

29

A Journal Group can consist of physical volumes of differing number of revolutions, physical volume capacity, and physical volume RAID configurations (e.g., RAID-1 and RAID-5). Data transfer speed of Parity groups is affected by physical volumes and RAID configurations.

- **Number of JNL VOLs**

- JNL entries are distributed across all Journal volumes in round-robin manner.
- Total JNL throughput of a JNL group is the sum of each JNL VOL.
- Defining all JNL VOLs with on devices with equivalent performance is desirable.

- **Best performance of the Journal Group**

- As many volumes as possible are in the journal group.
 - Decreases extent size, therefore forcing more equable use of all volumes in the group
- All volumes are approximately the same size.
 - Makes all extents approximately the same size

30


- JNL entries are distributed across all volumes to achieve maximum performance from assigned JNL VOLs.
- Therefore total JNL throughput of a JNL group is the sum of each JNL VOL performance. More VOLs, more throughput.
- One JNL group can have up to 64 JNL VOLs.
- To obtain stable throughput, equivalent performance characteristics of all JNL Volumes are desirable. Lower performance of some JNL VOLs may strongly affect the characteristic of the JNL group throughput.

In general, best performance of the Journal Group will be realized when as many volumes as possible are incorporated into the Journal Group and all volumes are approximately the same size.

■ RAID configuration

– RAID-5 (7D+1P) is recommended.

- 7+1 parity group is more efficient – spreads I/O across multiple parity groups, BEDS.
- Better than two RAID-5 (3D+1P) groups of same capacity.
- RAID-1 sequential write performance is low in comparison.



Journal Cache Specifications

■ Overview

- Generally Universal Replicator requires a smaller amount of cache than TrueCopy Asynchronous.
 - The recommended minimum is 25% to 30% additional cache.
- JNL Cache Threshold — MCU
 - Currently it is fixed at 30% of total cache.
 - In normal use, MCU journal cache contains mainly control (metadata) information.
 - When the data rate rises above link throughput, MCU journal cache will start to fill with overwritten data blocks.
 - If JNL volume throughput is insufficient, the journal cache may fill, causing Journal group suspension.
 - If JNL volume throughput is sufficient, but data rate stays above link bandwidth for a long enough time, JNL volumes will fill, causing group suspension.

■ Overview (continued)

- **Use of Cache** parameter setting: Determines whether or not the RCU will use journal cache.
 - Requires 25% to 30% additional cache in RCU, plus 1GB per journal group
 - Provides significant RCU performance benefit
 - Default setting
- For both MCU and RCU, the only way to increase journal cache size is to increase total cache size.

Configuration Planning

- MCU and RCU must be configured properly. Consider the following configuration design steps:
 - Investigate the customer's I/O patterns (with Performance Monitor, other data collection tools).
 - Look at peak IOPS, average IOPS, read/write ratio, read/write cache hit ratio, sequential/random pattern, etc.
 - Consider the customer's limitation and issues of DKC configuration.
 - Look at number of host paths, number of DKC-DKC paths, JNL groups and RAID level, etc.
 - Design MCU and RCU configurations within the customer's limitation.
 - Estimate number of FEDs and BEDs, number of host paths, JNL groups, RAID level, etc.

1. Identify applications to be replicated.
2. Determine amount of data to be replicated.
3. Estimate inflow data rate for data to be replicated.
4. Perform necessary calculations to determine journal volume size.
5. Identify suitable storage systems.
6. Identify suitable parity groups for data volumes and journal volumes.
7. Order and install necessary hardware and software.
8. Configure MCU and RCU.
9. Test replication.
10. Begin production replication.
11. Monitor, evaluate, modify.

Thorough configuration planning is necessary to achieve performance demands.

35

Notes:

For step 3. Estimate Inflow Data Rate for data to be replicated. This determines link bandwidth needed to prevent Journal Cache overflow.

For step 5. Identify suitable storage systems. Keep in mind data may have to be migrated from other storage systems.

Review – How Universal Replicator Works

- Replication process:
 - **Initial Copy** (also called Base Journal Copy)
 - During Initial Copy process, metadata pointers to data on the P-VOL and write sequence numbers are stored in the metadata area of the journal volume.
 - The base journal data is obtained from the P-VOL.
 - The data in the S-VOL synchronizes with the data in the P-VOL using the sequence numbering scheme stored as metadata on the primary journal volume.
 - This operation is conceptually similar to Initial Copy in TrueCopy.

36

Journal obtain is the function to store the already stored data in the primary data volume as a base journal in the journal volume at the primary site. This function stores the write data as a journal data in the journal volume with every update of the primary data volume according to the write instruction from the host. The journal obtain operation is performed according to the instruction of pair create or pair resync operation from the primary site. The write sequence number from the host is assigned to the journal data. According to this information, the write sequence consistency at the remote site can be maintained. The update data from the host is kept in the cache. Therefore, the journal obtain function for the update data is performed at another time from the recipient of update data from the host or the storage of data to the data volume.

- Replication process – **Journal Obtain, Read, Copy, Restore**
 - After Base Journal completes and primary data volume is updated by write command from the host, the updated data has to be replicated on the S-VOL. This is the **Update Copy**.
 - 1. **Journal Obtain** process is invoked. JNL data is queued for destage to MCU JNL volumes
 - 2. MCU sends **Journal Obtain** notification to RCU.
 - 3. RCU issues **Read Journal**.
 - 4. **Journal Copy** pulls from MCU cache (may have to initiate Read access if JNL data has been destaged to JNL volume) and sends data to RCU cache (if **Use Cache** parameter is set). Otherwise, data goes to R-JNL.
 - 5. RCU executes **Journal Restore** and sorts JNL data by sequence number.
 - 6. When complete, RCU compares sequence numbers with MCU. If agreement is reached, data is written to S-VOLs and JNL data on both M-JNL and R-JNL volumes are discarded.

37

Update Copy starts as the **Journal Obtain** process is invoked when data is written as journal data to cache and then the journal volume. Control information (metadata) is attached.

MCU then sends **Journal Obtain** notification to RCU. This tells the RCU that pending data is now ready. Data will remain in MCU cache until it is destaged to Journal Volume.

RCU then pulls data from MCU with **Read Journal** command.

If available in cache, **Journal Copy** pulls from MCU cache and sends data to RCU cache to be stored on secondary journal volume. If not in cache, data will come from MCU Journal Volume.

RCU executes **Journal Restore** to begin assembling the Journal Data into sequence number order

After the journal data is sequenced, RCU compares sequence numbers with MCU. If both agree on number of blocks sent and the highest number sent, both discard their retained journal data.

Review: Best Practice Sizing for Journal Volumes

- JNL VOL Throughput — Requirement
 - JNL VOL throughput of M-JNL group must **exceed** the peak inflow data rate for data to be replicated.
 - If JNL VOL throughput is less than the peak data rate, Universal Replicator will suspend groups (PFUS).
 - Components of JNL VOL throughput:
 - Journal Volume performance
 - Physical disk type (15K RPM is best)
 - Number of journal volumes in the group (more is better)
 - Size of journal volumes (roughly equal is better)
 - Overall parity group utilization (low I/O load of other volumes in the parity group)
 - Transfer Bandwidth
 - Number of transfer paths (links)
 - Speed of transfer pipe

Module Review

1. True or False: Synchronous write to P-VOL is the normal mode for Universal Replicator.
2. How many journal volumes can be added to a USP V Journal Group?
3. True or False: Two large journal volumes will give better performance than many small volumes. Why?
4. How does greatly varying size of journal volumes affect Journal Group?
5. True or False: Universal Replicator replicates data synchronously.
6. True or False: Update data blocks must arrive in the RCU in the original order as sent by MCU.
7. What are the three basic operations Universal Replicator uses to move data from P-VOL to S-VOL?
8. What determines the size of a journal volume data extent?

4. Storage Navigator Configuration

Module Objectives

- Upon completion of this module, the learner should be able to:
 - Prepare the storage systems for Hitachi Universal Replicator operation
 - Map LDEVs to ports for use as P-VOLs and S-VOLs
 - Configure Universal Replicator paths and add the remote system
 - Configure journal groups

2

Preparation: License Key Installation Procedure

■ Install Universal Replicator license keys

1. Locate the license key file.
2. Open Storage Navigator in a browser.
3. Switch to Modify Mode.
4. Click the **License Key** tab.
5. Change the **Mode** to File.
6. Browse for the license key file.
7. Click **Install**.
8. Click **Apply**.
9. Confirm keys are installed.

File Go Help 29 minute(s) remaining in session. Reset Logged in as: replicationuser

License Key SNMP Information E-mail Information Partition Definition

License Key

Product Name	Key Type	Permitted Volumes	Term(days)	Status
ShadowImage	Permanent	Unlimited(2 TB)		Installed
ShadowImage for Mainframe	Permanent	Unlimited(1 TB)		Installed
Copy-on-Write Snapshot	Permanent	Unlimited(2 TB)		Installed
Compatible FlashCopy(R) V2	Permanent	Unlimited(2 TB)		Installed
Storage Navigator	Permanent	Unlimited		Installed
SNMP Agent	Permanent	Unlimited		Installed
JAVA API	Permanent	Unlimited		Installed
TrueCopy	Permanent	Unlimited(0 TB)		Installed
TrueCopy for Mainframe	Permanent	Unlimited(0 TB)		Installed

License Install

Mode ☒ KeyCode ☐ File ☐ Available

KeyCode

FileName Browse

License Status

☒ Enable ☐ Disable Set

Install Remove

Apply Cancel

Preparation: Map Logical Devices

- Map logical devices for use as P-VOLs and S-VOLs
 - Universal Replicator data volumes must be LUNs, that is, LDEVs mapped to ports.
 - The primary and remote volumes must be the same size.
 - LDEVs must be mapped to ports that have the default target attribute setting.
 - Verify that the Host Mode setting is correct for the operating system of choice.

- Using Hitachi Storage Navigator Program:
 1. Universal Replicator license key is included in the Disaster Recovery Bundle which includes TrueCopy Sync.
 2. For 3DC support, and/or MXN Extended Consistency Group, the Disaster Recovery Extended bundle is required.
 3. Create LUSE/VLL volumes as necessary using the appropriate Storage Navigator functions.
 4. Use appropriate procedure for your storage platform.
 - Map candidate P-VOL and S-VOL logical devices (LDEVs) to the desired ports.
 - Record port/LUN information.
 5. Open Universal Replicator feature in both primary and secondary systems.
 - Set port attributes for replication links.
 - Add remote DKC.
 - Create journals.
 - Perform pair operations.

Note on LDEV numbers as displayed by Storage Navigator

- An LDEV number that ends with a “#” mark indicates that the LDEV is an external volume (for example: 00:00:01#).
- An LDEV number that ends with a letter “X” indicates that the LDEV is a virtual volume used by Dynamic Provisioning (for example 00:00:01X).

- VSP LDEV mapping procedure uses the Storage Navigator browser window.

The screenshot shows the Hitachi Storage Navigator interface. On the left, the 'Logical Devices' section is highlighted. The main area displays a table of Logical Devices (LDEVs) with columns: RAID Level, Emulation Type, Capacity, Provisioning Type, Attribute, and Num. Path. The table lists several LDEVs with RAID Level 5(7D+1P) and Emulation Type OPEN-V CVS, each with a capacity of 2.00 GB and a provisioning type of Basic. Below this, the 'Add LUN Paths' wizard is shown. The wizard has four steps: 1. Select LDEVs, 2. Select Host Groups, 3. View/Change LUN Paths, and 4. Confirm. In the 'Select LDEVs' step, a table of 'Available LDEVs' is shown with columns: LDEV ID, LDEV Name, Parity Group ID, Pool Name (ID), Emulation Type, and Capacity. The table lists LDEVs with IDs from 00:00:00 to 00:00:08, all with Emulation Type OPEN-V CVS and Capacity 2.00 GB. The 'Add' button is highlighted. The 'Selected LDEVs' table is empty.

6

- VSP LDEV mapping procedure (*continued*)

The screenshot shows the 'Add LUN Paths' wizard in the 'Select Host Groups' step. The wizard has four steps: 1. Select LDEVs, 2. Select Host Groups, 3. View/Change LUN Paths, and 4. Confirm. In the 'Select Host Groups' step, a table of 'Available Host Groups' is shown with columns: Port ID, Host Group Name, Host Mode, Port Attribute, Port Security, Number of Hosts, and Host Group Name. The table lists host groups with IDs from CL1-E to CL1-D, all with Host Mode 00[Standard] and Port Attribute Target. The 'Add' button is highlighted. The 'Selected Host Groups' table is empty. Below the tables, the 'Back', 'Next', 'Finish', 'Cancel', and 'Help' buttons are visible.

7

Preparation: VSP Map Logical Devices

■ VSP LDEV mapping procedure (*continued*)

The left screenshot shows the 'Selected Host Groups' dialog. It has a table with columns: Port ID, Host Group Name, Host Mode, and Port. The first row is circled in red and contains: CL1-C, 1C-G00 (00), 00[Standard], and Te. Below the table are 'Add' and 'Remove' buttons. At the bottom, it says 'Selected: 0 of 1' and has 'Back', 'Next', 'Finish', 'Cancel', and 'Help' buttons. The 'Next' button is circled in red.


The right screenshot shows the 'Change LDEV settings' dialog. It has a table with columns: Capacity, Provisioning Type, and LUN ID(1 Set of Paths). The table contains 9 rows, all with '2.00 GB' capacity and 'Basic' provisioning type. The LUN IDs are 0000 through 0008. A checkbox labeled 'CL1-C/1C-G00' is checked. Below the table are 'Change LDEV settings' and 'Change LUN IDs' buttons. At the bottom, it has 'Back', 'Next', 'Finish', 'Cancel', and 'Help' buttons. The 'Next' button is circled in red.

Finish, then Apply

8


Preparation: Establish Paths

- VSP: Verify your connections
 - From Storage Navigator browser window, select **Actions > Component > View Port Location**.
 - You can also access from the GUI with **Go > System Information > Port Location**.
 - Connected ports have yellow highlights.



9

- USP V: After mapping LDEVs to ports, verify connectivity:
 1. Open Storage Navigator on both MCU and RCU.
 2. Select Port Status.



- USP V: Verify your connections
 - Connected ports have yellow highlights.

Information	Port Status	Status
Port Status		
Cluster-1	16FS Fibre 8ch-4mp	8S Serial 4ch-4mp
	1A 3A 1E 3E	1J 3J 1N 3N 9A 9A
	5A 7A 5E 7E	5J 7J 5N 7N DA FA
	1B 3B 1F 3F 1K 3K 1P 3P 9B BB	
	5B 7B 5F 7F 5K 7K 5P 7P DB FB	
	1C 3C 1G 3G 1L 3L 1Q 3Q 9C BC	
	5C 7C 5G 7G 5L 7L 5Q 7Q DC FC	
	1D 3D 1H 3H 1M 3M 1R 3R 9D BD	
	5D 7D 5H 7H 5M 7M 5R 7R DD FD	
Cluster-2	16FS Fibre 8ch-4mp	8S Serial 4ch-4mp
	2A 4A 2E 4E	2J 4J 2N 4N AA CA
	6A 8A 6E 8E	6J 8J 6N 8N EA GA
	2B 4B 2F 4F 2K 4K 2P 4P AB CB	
	6B 8B 6F 8F 6K 8K 6P 8P EB GB	
	2C 4C 2G 4G 2L 4L 2Q 4Q AC CC	
	6C 8C 6G 8G 6L 8L 6Q 8Q EC GC	
	2D 4D 2H 4H 2M 4M 2R 4R AD CD	
	6D 8D 6H 8H 6M 8M 6R 8R ED GD	
1A Available and connected to the other device 1A Available, but not connected to the		


Preparation: Replication Links

- Universal Replicator requires multiple Fibre Channel paths between MCU and RCU.
 - *Forward Paths* are the control paths.
 - *Reverse Paths* carry data.
 - An operational minimum configuration would comprise at least three paths:
 - One control paths (two recommended for redundancy)
 - Two data paths
 - Each path requires a pair of ports (one port on each storage system).
 - Ports have the **target** attribute by default.
- For control (forward) paths:
 - **MCU** ports must be set to **Initiator** (Transmitting Port).
 - **RCU** ports are set to **RCU Target** (Receiving Ports).
- For data (reverse) paths:
 - **MCU** ports must be set to **RCU Target**.
 - **RCU** ports are set to **Initiator**.
- Port Attribute setting *must* be done with Storage Navigator or SVP.


- VSP, USP V, and USP with one-port FED cards
 - Note the port layout below before deciding which ports to use for Initiators and RCU targets.
 - If a host-connected port is accidentally set as Initiator or RCU target, I/O errors will occur.

Note: Zone the FC switches so that no other device on the SAN can access the ports used for Universal Replicator

Cluster 1 CHA



Cluster 2 CHA




13

MP = Channel Host Interface Microprocessor

■ USP V and USP with 2-port FED card


- Note the port layout on the CHA card before deciding which ports to use for Universal Replicator paths.
 - The most common port layout is shown below.
- All path hardware (cables, switches, extenders, and converters) must be in place before beginning Universal Replicator Path configuration.

Note: If switches are used, zone the switches so that no other device on the SAN can access the ports used for Universal Replicator



14

Note: Not all CHA cards use the configuration shown. Other CHA cards assign one port per MP and present no difficulties in configuration.




15

The major components of a Universal Replicator operation using Fibre Channel interface connections are:

- Universal Replicator software running on the storage system at the Primary (production) site and on the storage system at the remote (recovery) site
- Universal Replicator Volume Pairs:
 - ♦ P-VOLs (production volumes) at primary site. The M-DKCs contain the P-VOLs, the online production data volumes.
 - ♦ S-VOLs (secondary volumes) at remote site. The R-DKCs contain the S-VOLs, which are the asynchronous copies of the P-VOLs. S-VOLs can be online to hosts only when pairs are *Split* or *Deleted*.

■ Replication Links: MCU to RCU

- “Forward” links are used for control functions between MCU and RCU.
- “Reverse” links carry journal data.




16

Configuration: Overview

- Note that Storage Navigator is required for all steps.
 - Configure **port attributes** on *both* MCU and RCU storage systems.
 - **Add DKC** definition on *both* MCU and RCU storage systems.
 - This will create and test the paths.
 - Note that the **Add DKC** operation **must** be done on both storage systems so that paths are established in the correct direction.
 - Create **journal groups** on *both* MCU and RCU storage systems.
 - Set appropriate Journal Group Options.

- Virtual Storage Platform
 - Launching Universal Replicator from Storage Navigator browser window and from inside Storage navigator GUI
- Note:* All setup must be done on *both* MCU and RCU storage systems.


Launching from browser window



Launching from SN GUI



- For Universal Storage Platform V, select Universal Replicator from Storage Navigator on *both* MCU and RCU storage systems.




19

- For Universal Storage Platform or Network Storage Controller, select Universal Replicator from Storage Navigator on *both* MCU and RCU storage systems.

Select this option

Not this...
(Mainframe)




20

Inside Storage Navigator, select the Universal Replicator software. To do so, click on the Universal Replicator menu option.

Note: There are two forms of Universal Replicator: Open Systems and Mainframe. The Open Systems Universal Replicator instance is the top Universal Replicator option.

Configuration: Port Attributes

- Configure port attributes on both MCU and RCU
 - If initiator ports and RCU remote ports are not configured on both MCU and RCU, no pairs can be created and remote copy operations will not be performed.
1. Select **Subsystem**, CHA card slot, or attribute in the tree view on the left.
 2. In the table on the right, right-click on the selected port.
 3. Set desired initiator and RCU target ports.
 4. Click **Apply**.
 5. Verify the initiator and RCU target creation.



21

Do any of the following. Ports will be displayed in the upper-right list.

- Select **Subsystem**.
- Select a channel adapter from the tree view. The port for that channel adapter will be listed on the upper right of the panel.
- Select a port attribute (RCU Target or Initiator) from the tree view.

Now Select a port in the upper-right list in the panel. Select a Fibre Channel port only.

- Right-click the selected port and select the desired port attribute (RCU remote or initiator) from the pop-up menu. The rightmost column of the list displays Modified, and the Preset list displays the changes that you have made.
- See the Preset list to check the settings that you have made. If you want to change the attribute of a port, select and right-click the port from the list and then select the new attribute. If you want to cancel a change in the attribute of a port, select and right-click the port in the Preset list and then select Cancel. If you want to cancel changes in attributes of all ports, right-click on the Preset list and select the Cancel All command.

- Click the Apply button on the lower right in the Universal Replicator panel to enable all the settings. To disable all the operations and restore the list, click the Cancel button.

Review port setting for replication


- If using USP V or USP 2-port cards, changing port attributes will change two ports. Be sure the **associated** port is not being used by a **host or other storage system**
- The associated port *can* be used for a different replication path (either Universal Replicator or TrueCopy Remote Replication).
 - ♦ Not recommended for performance reasons
- Make sure to change port attributes on both MCU and RCU storage systems.
- Make sure all cables, switches, extenders, and converters are in place.
- Only ports assigned to SLPR0 may be used for replication.

Configuration: Add DKC

- Add DKC – Required for MCU and RCU

1. On the **DKC Operation** tab page,
select **DKC**.

2. Right-click,
select **DKC
Operation >
Add DKC**.



- For Virtual Storage Platform
 1. Do on both MCU and RCU.
 2. Enter the serial number of the remote DKC.
 3. Select **Controller ID = 6**
 4. Confirm Default Path Group.
 5. Set the ports you specified as Initiator and RCU Target.
 6. Click **Option**.

Note:

Port column — **Initiator** ports

Pair-Port column — **RCU Target** ports

M-R Path	
Port	Pair-Port
CL3-C	CL4-C

23

When you assign Logical Paths, use the port allocations you set as Initiator and RCU Target. Make sure an Initiator and RCU Target are assigned together. If two Initiators are grouped together, this will cause an error.

Display Features

DKC S/N: Allows you to enter the five-digit serial number of the remote storage system.

LDKC: Enter "00"

Controller ID: Allows you to enter the controller ID (that is, storage system family ID) of the remote storage system.

Note: The controller ID for a Universal Storage Platform disk storage system is 4. The controller ID for a Universal Storage Platform V disk storage system is 5.

Path Gr. ID: Allows you to enter the path group ID. Path group IDs are used for identifying groups of logical paths. One path group can contain up to eight logical paths.

Note: In the current version, you cannot enter path group IDs. Also, you cannot clear the Default check box. The number of path groups per one remote storage system is always 1.

M-R Path: Allows you to specify logical paths from initiator ports on the primary storage system to RCU remote ports on the remote storage system.

Port: Displays a list of initiator ports on the primary storage system. Select an initiator port from this drop-down list, or type in a Port Number.

Pair-Port: Displays a list of all ports on the remote storage system. Select an RCU remote port on the remote storage system from this drop-down list, or type in a port number.

Note: When specifying a port, you can use the keyboard to enter the port number. When you enter the port number, you can abbreviate the port number into two characters. For example, you can enter 1A instead of CL1-A. You can use uppercase and lowercase letters.

Option: Opens the DKC Option panel

Cancel: Cancels the settings you made on the Add DKC panel and then closes the panel.

- For USP V or USP VM
 1. Do on both MCU and RCU.
 2. Enter the serial number of the remote DKC.
 3. Enter the **LDKC** number (default 00).
 4. Select **Controller ID** = 5
 5. Select the ports used for Universal Replicator Links.
 6. Click **Option**.

Note:

Port column — **Initiator** ports

Pair-Port column — **RCU Target** ports

M-R Path	
Port	Pair-Port

24

When you assign Logical Paths, use the port allocations you set as Initiator and RCU Target. Make sure an Initiator and RCU Target are assigned together. If two Initiators are grouped together, this will cause an error.

Display Features

DKC S/N: Allows you to enter the five-digit serial number of the remote storage system.

LDKC: Enter "00"

Controller ID: Allows you to enter the controller ID (that is, storage system family ID) of the remote storage system.

Note: The controller ID for a Universal Storage Platform disk storage system is 4. The controller ID for a Universal Storage Platform V disk storage system is 5.

Path Gr. ID: Allows you to enter the path group ID. Path group IDs are used for identifying groups of logical paths. One path group can contain up to eight logical paths.

Note: In the current version, you cannot enter path group IDs. Also, you cannot clear the Default check box. The number of path groups per one remote storage system is always 1.

M-R Path: Allows you to specify logical paths from initiator ports on the primary storage system to RCU remote ports on the remote storage system.

Port: Displays a list of initiator ports on the primary storage system. Select an initiator port from this drop-down list, or type in a Port Number.


Pair-Port: Displays a list of all ports on the remote storage system. Select an RCU remote port on the remote storage system from this drop-down list, or type in a port number.

Note: When specifying a port, you can use the keyboard to enter the port number. When you enter the port number, you can abbreviate the port number into two characters. For example, you can enter 1A instead of CL1-A. You can use uppercase and lowercase letters.

Option: Opens the DKC Option panel

Cancel: Cancels the settings you made on the Add DKC panel and then closes the panel.

- Add DKC Options: All models
 1. Keep the default settings and click **Set**.
 2. When the main screen appears, click **Apply**.



25

The DKC Option panel displays the following:

The **Minimum Paths** option allows you to specify the minimum number of paths required for each remote storage system connected to the primary storage system (default = 1).

An error occurs if you enter a larger number than the number of paths already set on the Add DKC panel, or if the number of paths falls below this number (for example, due to path failures). When an error occurs, the primary storage system suspends all affected Universal Replicator (and Universal Replicator software for z/OS) pairs to prevent remote copy operations from adversely affecting performance due to the inadequate number of paths.

If the primary storage system contains Universal Replicator pairs which contain critical data for disaster recovery, set the minimum number of paths to one, so that Universal Replicator operations continue even if there is only one path to a remote storage system.

If you need high performance at the primary storage system, set the minimum number of paths to two or more, depending on the number of pairs managed by the primary storage system.

The **RIO MIH Time** setting specifies the RIO MIH timer value, which is the wait time until data transfer from the primary storage system to the remote storage

system is complete. The RIO MIH time value must be from 10 to 100 seconds. The default setting is 15 seconds.


Note: RIO MIH is an acronym for remote I/O missing interrupt handler. Not all operating systems have Missing Interrupt Handler routines. In that case, this function still works and provides a notification of abnormally long response times for the RIO. It can be used to help identify possible intermittent link failures.

■ Connect to the RCU and repeat the Add DKC function

- At this point both fibre paths should be fully functional and able to replicate data. You can confirm that the paths are functional on the DKC Operation panel.

Note: This step is required for a normal, cross-connected configuration. It is not necessary in a Loopback Configuration.

Remote system with serial number



Check Status of DKC



All the logical paths are in normal status.



A failure occurred at some of the logical paths.

26

Display Features

Tree: Lists remote storage systems. The following information appears to the right of the icon:

- Controller ID of a remote storage system (The controller ID is a storage system family ID)
- Serial number of the remote storage system
- Path group ID

The icon indicates the status of logical paths between the primary storage system and the remote storage system

All the logical paths are in normal status. A failure occurred at some of the logical paths.

Controller ID: Displays the controller ID of a remote storage system. The controller ID is a storage system family ID of a disk storage system. The icon indicates the status of logical paths between the primary storage system and the remote storage system:

All the logical paths are in normal status.

Note: The controller ID for a TagmaStore Universal Storage Platform disk storage system is **4** and a Universal Storage Platform V and VM disk storage system is **5**.

S/N: Displays the five-digit serial number of the remote storage system.

Path Gr. ID: Displays the path group ID.

Note: In the current version, the path group ID is set to the default value.

M-R Path: Indicates the channel type of the logical paths between the primary storage system and the remote storage system. This column always displays **Fibre**.

Status:

Normal: No failure occurs to the logical paths

Failed: All the logical paths fail

Warning: Some of the logical paths fail

Num of Path: Indicates the number of logical paths

- DKC Status: All Models
 - See notes.

Check Path
Status

Hitachi Storage Navigator

DKC Status

No.	Path Status	Port	Pair-Port
1	Initialization Failed	CL3-C	CL4-C

S/N : 52322 (00)

Controller ID : 6(VSP)

Path Gr. ID: Default

M-R Path : Fibre

Minimum Paths : 1

RIO MIH Time : 15 (sec.)

DKC Registered : 02/02/2011 03:01:29

Last Update Time : 02/02/2011 03:01:29

☐ Refresh the DKC Operation tab after this panel is closed

Refresh Close

27

The DKC Status panel displays the following:

Path Status: Common conditions

- Normal -This path has been successfully established and can be used for Universal Replicator remote copy activities.
- Nothing - An operation for configuring or deleting this logical path is in progress.
- Initialization Failed - An error occurred with initialization of connection between the primary and the remote storage system. Possible causes:
 - ♦ No cable is connected to the primary storage system.
 - ♦ No cable is connected to the remote storage system.
 - ♦ No cable is connected to the network device that comes between the primary and the remote storage system.
- Communication Time Out. This status indicates one of the following:
 - ♦ A timeout error has occurred between the primary and the remote storage systems.
 - ♦ A logic error is detected between the primary and the remote storage systems.
- Resource Shortage (Port). The primary storage system rejected the function for configuring logical path connection.
 - ♦ All logical path resources in the primary storage system might be used for other connections.

- **Resource Shortage (Pair-Port).** The remote storage system rejected the function for configuring logical path connection. All logical path resources in the remote storage system might be used for other connections.
- **Serial Number Mismatch.** The serial number of the storage system connected to this logical path does not match the serial number specified by the Add DKC panel. Delete and re-create the DKC.
- **Invalid Port** The port is not an initiator port.
- **Pair-Port Number Mismatch**
 - ♦ The specified port number is incorrect.
 - ♦ The port in the remote storage system is physically disconnected from the primary storage system.
- **Pair-Port Type Mismatch.** The port on the remote storage system is not an RCU remote port.
- **Communication Failed** A communication timeout error has occurred on the path between the primary and the remote storage systems.

Port: Indicates a port on the primary storage system.

Pair-Port: Indicates a port on the remote storage system.

DKC S/N: Indicates the serial number of the remote storage system.

Path Gr. ID: Indicates a path group ID.

M-R Path: Indicates the type of channel interface between the primary and the remote storage systems. This column displays Fibre.

Minimum Paths: Indicates the minimum possible number of paths between the primary and the remote storage systems.

RIO MIH Time: Indicates the remote I/O missing interrupt handler (RIO MIH) timer value, which is the wait time until data transfer from the primary storage system to the remote storage system is complete.

DKC Registered: Indicates the date and time when the primary and the remote storage systems are associated to each other.


Last Updated: Indicates the date and time when the last operation on a logical path to the remote storage system was performed.

Refresh the DKC Operations tab after this panel is closed: If you select this checkbox, information in the DKC Operation panel will be refreshed after you close the DKC Status panel.

Configuration: Journal Groups

- Create Journal Groups in both MCU and RCU systems.
- **Best practice:** If using MxN Consistency Groups, select **unique** M-JNL and R-JNL group IDs across all participating storage systems (see note).

1. Click the **Journal Operation** tab.
2. Select a free journal group.
3. Right-click on the journal group.
4. Select **Edit JNL VOLs**.



28

Note: UR Journal Groups can be defined in CCI as MxN Consistency Groups. Journal Group IDs in MxN Consistency Groups must be unique in the group, and the MxN Consistency Group may extend across storage systems.

To avoid possible conflicts, avoid duplicate Journal Group IDs in your replication environment.

However, M-JNL and R-JNL IDs can be identical within a group, as well as within a MxN Consistency Group.

■ Create Journal Groups

1. Select the CU number or Parity Group.

2. Select LDEV.

3. Click **Add**.

4. Click **Set**.

5. Repeat this step for the RCU storage system.

Hitachi Storage Navigator

Edit Journal Volumes

Journal Volumes

Parity Group	LDKC:CU:LDEV	Capacity (GB)	Emulation	CLPR	Operation
1-2	00:01:07	8.00	OPEN-V	00:CLPR0	Add
1-2	00:01:08	8.00	OPEN-V	00:CLPR0	Add
1-2	00:01:09	8.00	OPEN-V	00:CLPR0	Add
1-2	00:01:0A	8.00	OPEN-V	00:CLPR0	Add

Free Volumes

Parity Group	LDKC:CU:LDEV	Capacity (GB)	Emulation	CLPR	Operation
1-2	00:01:07	8.00	OPEN-V	00:CLPR0	Add
1-2	00:01:08	8.00	OPEN-V	00:CLPR0	Add
1-2	00:01:09	8.00	OPEN-V	00:CLPR0	Add
1-2	00:01:0A	8.00	OPEN-V	00:CLPR0	Add
1-2	00:01:0B	8.00	OPEN-V	00:CLPR0	Add

PG/CU change : ☐ PG ☐ PG(Ext) ☒ CU 01

2DC Cascade : ☐ Enable ☒ Disable

Journal : 000

Current : 0 (0.00 GB) Total : 4 (32.01 GB)

Set Cancel

29

The Edit JNL Volumes panel displays similar information about:

JNL Volumes – Journal Volumes

Free Volumes – Not registered in journal groups

For each category of volumes:

- **Parity Group:** indicates the parity group or the external volume group where a journal volume belongs.

Note: If the letter 'E' is displayed at the beginning of a group, the group is an external volume group. In the current version, however, the panel does not display external volumes.

- **CU:LDEV:** Indicates the CU number and the LDEV number of a volume. The CU number is displayed to the left of the colon (:). The LDEV number is displayed to the right of the colon.

Note: If a pound sign (#) is displayed at the end of a volume, the volume is an external volume. In the current version, however, the panel does not display external volumes. Consider the implications carefully before selecting an external volume as a Journal Volume.

- **Capacity:** Indicates the capacity of a journal volume in gigabytes.
- **Emulation:** Indicates the emulation type of the volume.

- **Operation:** Displays one of the following:
 - ♦ **Blank:** This column usually blank
 - ♦ **Add:** Indicates a volume to be added to a journal group
 - ♦ **Delete:** Indicates a volume to be deleted from a journal group

JNL Volume buttons

- **Add:** To register volumes in a journal group, select the volumes from **Free Volumes** and click **Add**.
- **Delete:** To delete volumes from a journal group. Select the volumes from **JNL Volumes** and click **Delete**.


Parity Group/CU change

To change the display in the **Free Volumes** list:

- **Parity Group:** Select to display volumes belonging to a parity group. Specify a parity group number in the text boxes to the right, and then click the Show button.
- **CU:** Select to display volumes belonging to a CU. Then select a CU from the list to the right.


Configuration: Journal Groups Remote Command Devices

- VSP Remote Command Device
 - Useful in 3DC Delta Resync and 2DC Cascade environments



30

■ USP V Remote Command Device



31

1. Ensure that the Journal Operation window is open (see Journal Operation Window).
 2. In the tree, select a master journal group or a restore journal group. *Note:* Must be an active Journal Group (not Initial Status)
 3. Right-click on JNL Groups. Select Assign R. CMD Dev.
 4. The Assign Remote Command Device dialog box appears.
 5. Set the Mirror ID (typically 1) and the S-VOL that you want to designate as a remote command device
 6. Click **Set** and close the Assign Remote Command Device window.
 7. Review the Preview list to check the settings that you have made. If you want to modify a setting, select and right-click the setting and then select Modify.
 8. Click **Apply** to assign the remote command device to the mirror.
- **Mirror ID:** Selects mirror ID. To assign the remote command device to the mirror of TrueCopy synchronous pair in the 3DC multi-target configuration, specify 00 as the mirror ID. To assign the remote command device to only one mirror ID, select **Not set** in mirror 2. The same mirror ID cannot be specified to mirror 1 and mirror 2 in the dialog box.

- **Remote Command Device:** Selects remote command devices to be assigned to the mirrors. If a remote command device is not used, select **Not use**.
- **Set:** Applies the settings in the window to the Journal Operation window
- **Cancel:** Cancels the settings and closes the dialog box.

Configuration: Journal Groups

- Create Journal Groups (*continued*)
 6. Verify journal groups

The screenshot shows the 'Universal Replicator Journal Operation' window. On the left, there is a tree view under 'Journals' with 'Registered' and 'Free' sub-items. The main area displays a table with the following columns: JNL (LDKC), Attribute, Status, Mirror ID, SIN (LDKC), and Pair JNL. The table contains several rows of data, all with a status of 'Initial'.

JNL (LDKC)	Attribute	Status	Mirror ID	SIN (LDKC)	Pair JNL
000(00)	Initial	Initial	00		
000(00)	Initial	Initial	01		
000(00)	Initial	Initial	02		
000(00)	Initial	Initial	03		
001(00)	Initial	Initial	00		
001(00)	Initial	Initial	01		
001(00)	Initial	Initial	02		
001(00)	Initial	Initial	03		
002(00)					
003(00)					
004(00)					

Attribute: Indicates the attribute of a journal group in the local storage system.



Initial

A journal group in initial status. Journal volumes are registered in this journal group, but not primary data volumes or secondary data volumes.



Master

A master journal group. Journal volumes and primary data volumes are registered in this journal group.



Restore

A restore journal group. Journal volumes and secondary data volumes are registered in this journal group.

Blank

Neither journal volumes nor data volumes are registered in this journal group.

32

Note that *both* **Attribute** and **Status** conditions are displayed. In general, attributes provide high-level indication of the Journal Group, while status provides more detail about the current condition of the Journal Group.

Journal Group Configuration Details

■ VSP Journal Status

-  **Initial:** A mirror in initial status. Journal volumes are registered in this journal, but no data volumes (P-VOLs or S-VOLs) are registered in this journal. If a UR pair is created, data volumes are registered in a mirror and the status of the mirror changes to **Active**.
-  **Active:** Either of the following:
Initial copy is in progress. The P-VOL and the S-VOL are not synchronized.
Initial copy is finished. The P-VOL and the S-VOL are synchronized.
If a mirror is in **Active** status, some of the pairs in the mirror might be split. If this happens, the word **Warning** appears. To restore such pairs, use the **Pair Operation** window.
-  **HaltAccept:** An operation for splitting the mirror has been started. The status of the mirror immediately changes to **Halting**. **HaltAccept** can indicate status of restore journal in the mirror, but cannot indicate status of master journals in the mirror.
-  **Halting:** An operation for splitting or deleting the mirror is in progress. The P-VOL and S-VOL are not synchronized. Splitting a mirror changes the status in the following order: **Halting**, **Halt**, **Stopping**, and finally **Stopped**. Deleting a mirror changes the status in the following order: **Halting**, **Halt**, **Stopping**, **Stopped**, and finally **Initial**.
-  **Halt:** An operation for splitting or deleting the mirror is in progress. The P-VOL and the S-VOL are not synchronized.

-  **Stopping:** An operation for splitting or deleting the mirror is in progress. The P-VOL and the S-VOL are not synchronized.
-  **Stopped:** An operation for splitting or deleting the mirror is finished. The P-VOL and the S-VOL are not synchronized.
-  **Hold:** A UR pair for delta resync operation is created.
-  **Holding:** A UR pair for delta resync operation is being created.
-  **Hold(Failure):** An error occurred with the UR pair for delta resync operation.
- Blank: No data volumes are registered in this journal

■ USP V Journal Group Status

Status: Indicates the status of a journal group in the local storage system.


-  **Initial**
A journal group in initial status. Journal volumes are registered in this journal group, but not primary data volumes or secondary data volumes. When you create a UR volume pair, data volumes will be registered in a journal group. The status of the journal group will change to Active.
-  **Active**
Either of the following:
 - Initial copy is in progress. The primary data volume and the secondary data volume are not synchronized.
 - Initial copy is finished. The primary data volume and the secondary data volume are synchronized.

Note: When a journal group is in Active status, some of the data volume pairs in the journal group might be split. If this happens, the word **Warning** is displayed. To restore such data volume pairs, use the Pair Operation panel. For details on how to restore them, see section 5.6
-  **HaltAccept**
An operation for splitting the mirror has been started. The status of the journal group will immediately change to Halting.
Note: HaltAccept can indicate status of restore journal groups, but cannot indicate status of master journal groups.
-  **Halting**
An operation for splitting or deleting the mirror is in progress. The primary data volume and the secondary data volume are not synchronized.
When you split a mirror, the status will change in the following order: Halting, Halt, Stopping, and finally Stop.
When you delete a mirror, the status will change in the following order: Halting, Halt, Stopping, Stop, and finally Initial.
-  **Halt**
An operation for splitting or deleting the mirror is in progress. The primary data volume and the secondary data volume are not synchronized.
-  **Stopping**
An operation for splitting or deleting the mirror is in progress. The primary data volume and the secondary data volume are not synchronized.
-  **Stop**
Either of the following:
 - An operation for splitting the mirror is finished.
 - The operation for deleting the mirror is in progress.

The primary data volume and the secondary data volume are not synchronized.

■ VSP Journal Options

1. In the Journal Operation panel, select the Mirror ID you want to change.
2. Select **Journals > Change JNL Option**.



35

■ VSP Journal Options (*continued*)

- Specify Yes or No for **Inflow Control** setting.
- If Inflow Control is **Yes**, **Data Overflow Watch** is active. Specify the number of seconds for the system to monitor write data to the journal volume when the journal volume threshold (80%) is reached. Range is 0–600 seconds. Default is 60 seconds.
- If **Inflow Control** is **No**, **Data Overflow Watch** is disabled.
- **Use of Cache:**
 - **Use:** Journal data will be stored in the RCU cache. See note.
 - **Not Use:** Journal data is not stored in the cache.


The 'Change Journal Option' dialog box has the following fields and controls:

- Inflow Control:** Radio buttons for 'Yes' (selected) and 'No'.
- Data Overflow Watch (0 to 600 sec):** Text input field with '60' entered.
- Use of Cache:** Dropdown menu with 'Use' selected.
- Unit of Path Watch Time:** Radio buttons for 'min.', 'hour', and 'day'.
- Path Watch Time (1 to 59 min., 1 to 23 hour, 1 to 30 day):** Text input field.
- Forward Path Watch Time:** Dropdown menu.
- Copy Pace:** Dropdown menu.
- Transfer Speed (Mbps):** Dropdown menu.
- Delta resync Failure:** Dropdown menu.
- Buttons:** 'Set' and 'Cancel' at the bottom right.

Note: When there is insufficient space in the cache, journal data will also be stored into the journal volume. This setting only takes effect on the journal volumes of RAID-5 or RAID-6.


■ VSP — Changing Mirror Options

1. In the Journal Operation panel, select the Mirror ID you wish to change.
2. Select **Mirrors > Change Mirror Option**.



■ VSP — Changing Mirror Options
(continued)

- **Unit of Path Watch Time** – See note.
- **Path Watch Time** – Specify how long to defer group suspend after path failure
- **Forward Path Watch Time** – Set Path Watch Time on RCU
- **Copy Pace** – Low, Medium, High
- **Transfer Speed** – Set line speed of data transfer in MB/sec (for takeover)
- **Delta Resync Failure** – Action taken if Delta Resync operation fails



38

Unit of Path Watch Time - Specify the unit, **minute**, **hour**, or **day**.

Path Watch Time - Specify the interval from the time a path becomes blocked to when the mirror is split (suspended). The interval must be the same for master and restore journals in the same mirror (see next item).

Note: If you want a mirror to split immediately after a path becomes blocked, ask Hitachi Data Systems Support Center to set system option mode 448 to ON and set system option mode 449 to OFF.

Forward Path Watch Time

- **Yes:** The Path Watch Time value will be forwarded to the restore journal.
- **No:** The Path Watch Time value will not be forwarded to the restore journal. **No** is the default.
- **Blank:** The current setting of Forward Path Watch Time will remain unchanged.

Copy Pace - Specify the pace for initial copy activity per volume. This field cannot be specified on the remote system. **Low** is the default.

When specifying **Medium**, ensure that write I/O is 10 MB/s or less per parity group. If it exceeds 10 MB/s, pairs may be suspended.

When specifying **High**, ensure that I/O will not occur. If update I/O occurs, pairs may be suspended.


Transfer Speed - Specify the line speed (in Mbps [megabits per second]) of data transfer. Specify one of the following: **256**, **100**, or **10**.

Recommended values are as follows:

- 10 is recommended if the transfer speed is from 10 Mbps to 99 Mbps.
- 100 is recommended if the transfer speed is from 100 Mbps to 255 Mbps.
- 256 is recommended if the transfer speed is 256 Mbps and more

■ USP V – Changing Options for a Journal Group

1. In the Journal Operation panel, do either of the following:
 - In the tree, right-click a journal group and then select **JNL Groups** and **Change JNL Option** from the pop-up menu.
 - In the upper-right list, right-click the desired journal group and then select **JNL Groups** and **Change JNL Option** from the pop-up menu.
2. In the **Change JNL Option** panel, change journal group options and then select **Set**.
3. See the **Preset** list in the Journal Operation panel to check the settings that you have made.



■ USP V – Changing Options for a Journal Group (continued)

- **Inflow Control** – Slows response to host by invoking Synchronous Mode
- **Data Overflow** – Interval between journal volume full and suspension of groups (0–600 sec). Active only when Inflow Control is set
- **Copy Pace** – Low, Medium, High
- **Path Watch Time** – Specify how long to defer group suspend after path failure. To enable, set SVP Mode 449 off.
- **Forward Path Watch Time** – Set Path Watch Time on RCU
- **Use of Cache** – Hold Journal Data in RCU Journal Cache
- **Speed of Line** – Set line speed of data transfer in MB/sec (for takeover)
- **Delta Resync Failure** – Action taken if Delta Resync operation fails

Change JNL Option	
Inflow Control :	☒ Yes ☐ No
Data Overflow Watch (0 to 600 sec) :	60
Copy Pace :	Low
Unit of Path Watch Time :	☒ min. ☐ hour ☐ day
Path Watch Time (1 to 59 min., 1 to 23 hour, 1 to 30 day) :	5
Forward Path Watch Time :	No
Use of Cache :	Use
Speed of Line (Mbps) :	256
Delta resync Failure :	Entire
Set Cancel	

40

Note: Several items are locked out if pairs exist in the Journal Group. The following items cannot be changed:

- Inflow Control
- Copy Pace
- Unit of Path Watch Timer
- Forward Path Watch
- Use of Cache
- Delta Resync Failure

Inflow Control: Allows you to specify whether to restrict inflow of update I/Os to the journal volume (in other words, whether to delay response to the hosts).

- **Yes** indicates inflow will be restricted. **No** indicates inflow will not be restricted.

Note: If **Yes** is selected and the metadata or the journal data is full, the update I/Os may stop. (Journal Groups suspended)

Data Overflow Watch: Allows you to specify the time (in seconds) for monitoring whether metadata and journal data are full. This value must be within the range of 0 to 600 seconds.

Note: If **Inflow Control** is **No**, **Data Overflow Watch** does not take effect and does not display anything.

Copy Pace: Allows you to specify the pace (speed) for an initial copy activity for one volume. The default is **Low**.

- **Low:** The speed of the initial copy activity is slower than **Medium** and **High**.
- **Medium:** The speed of the initial copy activity is faster than **Low** and slower than **High**.

If you specify **Medium**, make sure that the amount of update I/Os is 10 MB/sec or less per one parity group. If it exceeds 10 MB/sec, data volume pairs may become split (suspended).

- **High:** The speed of the initial copy activity is faster than **Low** and **Medium**. If you specify High, make sure that update I/Os will not occur. If update I/Os occur, data volume pairs may become split (suspended).

Note: The remote storage system administrator cannot specify this option.

Path Watch Time: Allows you to specify the interval from when a path gets blocked to when a mirror gets split (suspended). This value must be within the range of 1 to 60 minutes.

Note: Make sure that the same interval is set to both the master and restore journal groups in the same mirror, unless otherwise required. If the interval differs between the master and restore journal groups, these journal groups will not be suspended simultaneously. For example, if the interval for the master journal group is 5 minutes and the interval for the restore journal group is 60 minutes, the master journal group will be suspended in 5 minutes after a path gets blocked, and the restore journal group will be suspended in 60 minutes after a path gets blocked.

Caution: By default, the factory enables (turns ON) SVP mode 449, disabling the path watch time option. If you'd like to enable the path watch time option, please disable mode 449 (turn it OFF).

Note: If you want to split a mirror (suspend) immediately after a path becomes blocked, please disable SVP modes 448 and 449 (turn OFF).

Forward Path Watch Time: Allows you to specify whether to forward the Path Watch Time value of the master journal group to the restore journal group. If the Path Watch Time value is forwarded, the two journal groups will have the same Path Watch Time value.

- **Yes:** The Path Watch Time value will be forwarded to the restore journal group.
- **No:** The Path Watch Time value will not be forwarded to the restore journal group. No is the default.
- **Blank:** The current setting of **Forward Path Watch Time** will remain unchanged.

Caution: This option cannot be specified in the remote site.

Use of Cache: Allows you to specify whether to store journal data in the restore journal group into the cache.

- **Use:** Journal data will be stored into the cache.

Note: When there is insufficient space in the cache, journal data will also be stored into the journal volume.

- **Not Use:** Journal data will not be stored into the cache.
- **Blank:** The current setting of **Use of Cache** will remain unchanged.

Caution: This setting does not take effect on master journal groups. However, if the CCI horctakeover command is used to change a master journal group into a restore journal group, this setting will take effect on the journal group.

Speed of Line: Allows you to specify the line speed of data transfer. The unit is Mbps (megabits per second).

- You can specify one of the following: **256, 100, or 10.**

Caution: This setting does not take effect on master journal groups. However, if the CCI **horctakeover** command is used to change a master journal group into a restore journal group, this setting will take effect on the journal group.

Delta resync Failure: Allows you to specify the processing that would take place when delta resync operation cannot be performed.

- **Entire:** All the data in primary data volume will be copied to remote data volume when delta resync operation cannot be performed. The default is Entire.
- **None:** No processing will take place when delta resync operation cannot be performed. Therefore, the remote data volume will not be updated. If Delta Resync pairs are desired, they will have to be created manually.

Caution: – This option cannot be specified in the remote site.

■ **Both VSP and USP V — System Option Modes 448 and 449**


- If you want a mirror to suspend immediately after path failure, ask the Hitachi Data Systems Support Center to set the system option mode 448 to ON and set the system option mode 449 to OFF.

Mode	Description
448	ON: If the SVP detects a blocked path, the SVP assumes that an error occurred, and then immediately splits (suspends) the mirror. OFF: If the SVP detects a blocked path and the path does not recover within the specified period of time, the SVP assumes that an error occurred, and then splits (suspends) the mirror. The mode 448 setting takes effect only when mode 449 is set to OFF.
449	ON: The SVP does not detect blocked paths. OFF: The SVP detects blocked paths and monitors the time until the mirrors get split (suspended).

41

■ **VSP – Journals after pair creation**

1. Master Journal – primary data volumes
2. Restore – Secondary Data Volumes



42

■ VSP – Journal Detail Information

1. In the Journal Operation panel navigation tree, select a journal group.
2. Right-click the selected journal group and then select **JNL Groups** and **JNL Detail** from the pop-up menu.
3. View detailed information about the journal group. (notes below)

Hitachi Storage Navigator

Journal Detail

Journal (LDKC) :	000 (00)	Attribute :	Master
Number of Journal Volumes :	2	Journal Capacity (GB) :	40.00
Data Volumes :	1	Data Capacity (GB) :	2.00
Inflow Control :	Yes	Data Overflow Watch (sec) :	60
Use of Cache :	Used		

Journal Volumes

Parity Group	LDKC:CU/LDEV	Capacity (GB)	Emulation	CLPR
5-3	00:10:18	20.00	OPEN-V	00:CLPR0
5-3	00:10:19	20.00	OPEN-V	00:CLPR0

Mirrors

Mirror ID	Attribute	Status	CTG	S/N (LDKC)	Pair JNL	CTRLID (Model Name)
00		Initial				
01	Master	Active	000	53867 (00)	001	6(VSP)
02		Initial				
03		Initial				

Previous Next Close

43

The **Journal Detail** panel displays the following:

Journal (LDKC): Indicates the number of a journal group.

Attribute: Indicates the attribute of the journal group.

- **Initial:** A journal group in initial status. Journal volumes are registered in this journal group, but no primary data volumes or remote data volumes.
- **Master:** A master journal group. Journal groups and primary data volumes are registered in this journal group.
- **Restore:** A restore journal group. Journal groups and remote data volumes are registered in this journal group.
- **Blank:** Neither journal volumes nor data volumes are registered in this journal group.

JNL Volumes: Indicates the number of journal volumes registered in the journal group.

JNL Capacity: Indicates the total capacity of all the registered journal volumes.

Data Volumes: Indicates the number of data volumes associated with the journal group.

Data Capacity: Indicates the total capacity of all the data volumes.

Inflow Control: Indicates whether to restrict inflow of update I/Os to the journal volume (whether to delay a response to the hosts).

Data Overflow Watch: Indicates the time (in seconds) for monitoring whether metadata and journal data are full.

Use of Cache: Indicates whether to store journal data in the restore journal group into the cache.

JNL Volumes: Displays a list of registered journal volumes:

Parity Group: indicates the parity group where a journal volume belongs.

CU:LDEV: Indicates the CU number and the LDEV number of a journal volume.

Capacity: Indicates the capacity of a journal volume in gigabytes.

Emulation: Indicates the emulation type of a journal volume.

CLPR: Indicates the number and the name of the CLPR where the journal volume belongs.

Mirrors: Displays a list of mirrors:

Mirror ID: Indicates a mirror ID. This column is blank if the attribute of the journal group is neither Master nor Restore.

Attribute: Indicates whether the Mirror is a Master or Restore Mirror

Status: Indicates the status of a journal group (or a mirror) in the primary storage system.

- **Initial** A journal group in initial status. Journal volumes are registered in this journal group, but not primary data volumes or remote data volumes. When you create a Universal Replicator volume pair, data volumes will be registered in a journal group. The status of the journal group will change to **Active**.
- **Active** Either of the following:

- ♦ Initial copy is in progress. The primary data volume and the remote data volume are not synchronized.
- ♦ Initial copy is finished. The primary data volume and the remote data volume are synchronized.

Note: If a journal group is in **Active** status, some of the data volume pairs in the journal group might be split. If this happens, the word **Warning** is displayed. To restore such data volume pairs, use the Pair Operation panel.

- ♦ **Halt Accept** An operation for splitting the mirror has been started. The status of the mirror will immediately change to **Halting**.

Note: **Halt Accept** can indicate status of restore journal groups, but cannot indicate status of master journal groups.

- **Halting** An operation for splitting or deleting the mirror is in progress. The primary data volume and the remote data volume are not synchronized. When

you split a mirror, the status will change in the following order: **Halting**, **Halt**, **Stopping**, and finally **Stop**. When you delete a mirror, the status will change in the following order: **Halting**, **Halt**, **Stopping**, **Stop**, and finally **Initial**.

- **Halt** An operation for splitting or deleting the mirror is in progress. The primary data volume and the remote data volume are not synchronized.
- **Stopping** An operation for splitting or deleting the mirror is in progress. The primary data volume and the remote data volume are not synchronized.
- **Stop**: Either of the following:
 - ♦ An operation for splitting the mirror is finished.
 - ♦ The operation for deleting the mirror is in progress. The primary data volume and the remote data volume are not synchronized.
- **Blank** Neither journal volumes nor data volumes are registered in this journal group.

CTG: Indicates the number of a consistency group to which the mirror belongs. This column is blank if there is no consistency group.

S/N: Indicates the serial number of the remote storage system. This column is blank if the attribute of the journal group is neither **Master** nor **Restore**.

Pair JNLG: Indicates the number of a journal group in the remote storage system. This column is blank if the attribute of the journal group is neither **Master** nor **Restore**.

Controller ID: Indicates the controller ID (that is, storage system family ID) of the remote storage system. This column is blank if the attribute of the journal group is neither **Master** nor **Restore**.

Path Watch Time: Indicates the time for monitoring blockade of paths to the remote storage system.

Pairs: Indicates number of data pairs in the Mirror

Capacity: of the data volumes

Copy Pace setting specified when the pair was created

Transfer speed setting in paircreate parameters

Delta Resync Failure setting

Remote Command Device location if defined

■ USP V – Journal Group Detail

1. In the Journal Operation panel navigation tree, select a journal group.
2. Right-click the selected journal group and then select **JNL Group** and **JNL Status** from the pop-up menu.
3. View detailed information about the journal group. (notes below)

Storage Navigator

JNL Group Detail

JNL Group (LDKC) :	00 (00)		
Attribute :	Initial		
JNL Volumes :	2	JNL Capacity (GB) :	8.00
Data Volumes :	0	Data Capacity (GB) :	0.00
Inflow Control :	Yes	Data Overflow Watch (sec) :	60
Copy Pace :	Low		
Use of Cache :	Use	Speed of Line (Mbps) :	256
Delta resync Failure :	Entire		

JNL Volumes

Parity Group	LDKC:CU:LDEV	Capacity (GB)	Emulation	CLPR
1-2	00:01:00	4.00	OPEN-V	00:CLPR0
1-2	00:01:01	4.00	OPEN-V	00:CLPR0

Mirrors

Mirror ID	Attribute	Status	CTG	S/N (LDKC)	Pair JNLG	CTRLID (Mode)
	Initial					

Previous Next Close

44

The JNL Group Detail panel displays the following:

JNL Group: Indicates the number of a journal group.

Attribute: Indicates the attribute of the journal group.

- **Initial:** A journal group in initial status. Journal volumes are registered in this journal group, but no primary data volumes or remote data volumes.
- **Master:** A master journal group. Journal groups and primary data volumes are registered in this journal group.
- **Restore:** A restore journal group. Journal groups and remote data volumes are registered in this journal group.
- **Blank:** Neither journal volumes nor data volumes are registered in this journal group.

JNL Volumes: Indicates the number of journal volumes registered in the journal group.

JNL Capacity: Indicates the total capacity of all the registered journal volumes.

Data Volumes: Indicates the number of data volumes associated with the journal group.

Data Capacity: Indicates the total capacity of all the data volumes.

Meta/Data Ratio: Indicates the ratio of metadata area to journal data area. Currently fixed at 32.

Extent: Indicates the number of extents. Currently 32.

Inflow Control: Indicates whether to restrict inflow of update I/Os to the journal volume (whether to delay a response to the hosts).

- **Yes** indicates inflow will be restricted. **No** indicates inflow will not be restricted.

Note: **Inflow Control** displays nothing if the journal group is a restore journal group.

Data Overflow Watch: Indicates the time (in seconds) for monitoring whether metadata and journal data are full.

Note: **Data Overflow Watch** displays nothing when one of the following conditions is satisfied:

- Inflow Control is No.
- The journal group is a restore journal group.

Use of Cache: Indicates whether to store journal data in the restore journal group into the cache.

- **Use:** Journal data will be stored in the cache. Requires additional cache in RCU but has significant performance benefits.

Note: When there is insufficient space in the cache, journal data will also be stored into the journal volume.

- **Not Use:** Journal data will not be stored into the cache.

Caution: This setting does not take effect on master journal groups. However, if the CCI **horctakeover** command is used to change a master journal group into a restore journal group, this setting will take effect on the journal group.

JNL Volumes: Displays a list of registered journal volumes:

Parity Group: indicates the parity group where a journal volume belongs.

CU:LDEV: Indicates the CU number and the LDEV number of a journal volume.

Capacity: Indicates the capacity of a journal volume in gigabytes.

Emulation: Indicates the emulation type of a journal volume.

CLPR: Indicates the number and the name of the CLPR where the journal volume belongs.

Mirrors: Displays a list of mirrors: – **Mirror ID:** Indicates a mirror ID. This column is blank if the attribute of the journal group is neither Master nor Restore.

Status: Indicates the status of a journal group (or a mirror) in the primary storage system.

- **Initial** A journal group in initial status. Journal volumes are registered in this journal group, but not primary data volumes or remote data volumes. When you create a Universal Replicator volume pair, data volumes will be registered in a journal group. The status of the journal group will change to **Active**.
- **Active** Either of the following:
 - ♦ Initial copy is in progress. The primary data volume and the remote data volume are not synchronized.
 - ♦ Initial copy is finished. The primary data volume and the remote data volume are synchronized.

Note: If a journal group is in **Active** status, some of the data volume pairs in the journal group might be split. If this happens, the word **Warning** is displayed. To restore such data volume pairs, use the Pair Operation panel.

- ♦ **Halt Accept** An operation for splitting the mirror has been started. The status of the mirror will immediately change to **Halting**.

Note: **Halt Accept** can indicate status of restore journal groups, but cannot indicate status of master journal groups.
- **Halting** An operation for splitting or deleting the mirror is in progress. The primary data volume and the remote data volume are not synchronized. When you split a mirror, the status will change in the following order: **Halting**, **Halt**, **Stopping**, and finally **Stop**. When you delete a mirror, the status will change in the following order: **Halting**, **Halt**, **Stopping**, **Stop**, and finally **Initial**.
- **Halt** An operation for splitting or deleting the mirror is in progress. The primary data volume and the remote data volume are not synchronized.
- **Stopping** An operation for splitting or deleting the mirror is in progress. The primary data volume and the remote data volume are not synchronized.
- **Stop:** Either of the following:
 - ♦ An operation for splitting the mirror is finished.
 - ♦ The operation for deleting the mirror is in progress. The primary data volume and the remote data volume are not synchronized.
- **Blank** Neither journal volumes nor data volumes are registered in this journal group.

CTG: Indicates the number of a consistency group to which the mirror belongs. This column is blank if there is no consistency group.

S/N: Indicates the serial number of the remote storage system. This column is blank if the attribute of the journal group is neither **Master** nor **Restore**.


Pair JNLG: Indicates the number of a journal group in the remote storage system. This column is blank if the attribute of the journal group is neither **Master** nor **Restore**.

Controller ID: Indicates the controller ID (that is, storage system family ID) of the remote storage system. This column is blank if the attribute of the journal group is neither **Master** nor **Restore**. *Note:* The controller ID for a TagmaStore™ USP and NSC disk storage system is **4**.

Path Watch Time: Indicates the time for monitoring blockade of paths to the remote storage system.

■ Deleting Journal Volumes

1. In the Journal Operation panel navigation tree, select a journal group.
2. Select one or more Journal Volumes.
 - **Group Attribute** must be *Initial* (no pairs) or **Group Status** must be *STOP* (all pairs suspended) or *HOLD* (Delta Resync pairs suspended).



Note that the last Journal Volume in the group cannot be deleted if there are any pairs in the group, without regard to the pair status.

Review: Preparation Checklist

- Install necessary **license keys** on both storage systems.
- Map **logical devices** to ports if necessary.
- Establish **physical paths** between storage systems.
 - Make sure all cables, switches, extenders, and converters are in place.
- Configure **replication port attributes** on both MCU and RCU storage systems.
- **ADD DKC** on both MCU and RCU storage systems.
 - This will create and test the Universal Replicator links between MCU and RCU.
 - Note that the ADD DKC operation *must* be done on both storage systems so that paths are established in the correct direction.
- **Create journal groups** on both MCU and RCU storage systems.
 - Set appropriate Journal Options after creating journal groups.

46

Module Review

1. If the minimum path setting is higher than the Universal Replicator paths available, what results will occur?
2. When Universal Replicator is configured through Storage Navigator, what is the first step?
3. What is the minimum number of replication links used for Universal Replicator links?
4. When a Journal Group is created, what Journal Group Attribute is shown?
5. What LDEV emulation types can be used for Journal Volumes?

5. Storage Navigator Pair Operations

Module Objectives

- Upon completion of this module, the learner should be able to:
 - Prepare storage systems and use Universal Replicator commands to perform operations in Storage Navigator
 - Describe the features and operation of Usage Monitor
 - Invoke Universal Replicator history function
 - Describe troubleshooting procedures for path failures and suspended pairs

2

Preparation for Operations

- Preparation for Universal Replicator operations
 - Universal Replicator license keys installed on candidate primary (MCU) and remote (RCU) Hitachi enterprise storage systems
 - At least two logical fibre paths configured and activated between the storage systems
 - At least one Journal Group present in each MCU and RCU
 - A list of candidate LDEVs for P-VOLs and associated S-VOLs showing current port mapping details:
 - Port ID
 - Host Group ID
 - LUN number

3

Hitachi enterprise storage systems include:

- Virtual Storage Platform
- Universal Storage Platform V
- Universal Storage Platform VM
- Universal Storage Platform
- Network Storage Controller

Commands Overview


Command	Function/Description
Pairedisplay	To view detailed information about a pair of data volumes Status transition: N/A
Paircreate	Creates a Universal Replicator volume pair Status transition: SMPL→COPY→PAIR
Pairsplit -S	Deletes a Universal Replicator volume pair Status transition: Any status/SMPL → SMPL
Pairsplit -r	Splits a pair Status transition: Any status/SMPL and PSUE →PSUS
Pairresync	Resynchronizes a pair Status transition: PSUS/PSUE→PAIR

Status: Indicates status of volume pairs

 SMPL	The volume is not currently assigned to a pair. When the initial copy is started by a paircreate operation, the volume status changes to COPY..
 COPY	The initial copy operation is in progress. Data in the primary data volume is not synchronized with data in the secondary data volume. When the initial copy is complete, the status will change to PAIR.
 PAIR	The volume is paired with another volume. The two volumes are fully synchronized. All updates from the host to the primary data volume are duplicated at the secondary data volume.
 PSUS	The pair has been split by the user (pairsplit-r), or the pair has been released by the user from the secondary storage system (pairsplit-S). The primary data volume and the secondary data volume are not synchronized.
 PSUE	The pair has been split by the primary storage system or the secondary storage system due to an error or a failure. The primary data volume and the secondary data volume are not synchronized.
 Suspending	The primary data volume and the secondary data volume are not synchronized. This pair is in transition from the PAIR or COPY status to the PSUS/PSUE status.
 Deleting	The primary data volume and the secondary data volume are not synchronized. This pair is in transition from the PAIR, COPY, or PSUS/PSUE status to the SMPL status.

■ Pair Operations

1. Open Storage Navigator on primary storage system (MCU).
2. Click the **Universal Replicator** icon.
3. Click on the **Pair Operation** tab.
4. Select port from the tree view.
5. Select candidate P-VOL from the list on the right.
6. Right-click and select the operation.



6

To begin data replication:

1. Open Storage Navigator in **Modify** mode and navigate to the Universal Replicator interface.
2. Click on the **Pair Operation** tab. All mapped LUNS are displayed on the left side of the window.
3. Select the Universal Replicator production volumes from the list.

Note: Universal Replicator only supports OPEN-V type volumes.


paircreate

■ VSP Paircreate panel

– Volume information →

– Journal Information →

– Detail Information →



The screenshot shows the 'Hitachi Storage Navigator' window with the 'Paircreate' tab selected. The panel is divided into several sections:

- Port-GID-LUN**:
 - P-VOL: CL1-C - 00 - 000 (00:00:00)
 - S-VOL: CL2-C (dropdown), 00 (dropdown), 000 (dropdown)
 - Select Other S-VOL(s): Increment (dropdown)
- M-JNL, Mirror ID, R-JNL**:
 - Mirror: 000 (dropdown), 01 (dropdown), 001 (dropdown)
 - CT Group: 000 (dropdown)
 - S/N, CTRLID, Path: 53867(00), 6(VSP), Fibre (dropdown)
- Initial Copy, Priority, Error Level**:
 - Initial Copy: Entire (dropdown)
 - Priority: Entire (dropdown)
 - Error Level: Mirror (dropdown)
- M-JNL Information**:
 - Current Mirror(s): 0
 - Total Mirror(s): 0


At the bottom right, there are 'Set' and 'Cancel' buttons.

7

Note: You can select and right-click more than one volume if you want to create more than one pair at one time. Choose all the remote volumes from the same remote storage system.

▪ **paircreate:** Volume information

- P-VOL is selected already.
- Select S-VOL by entering:
 - Port ID
 - Host Group ID
 - LUN number
- If S-VOL information is not known, open Storage Navigator on remote storage system and look at LUN manager.



8

When the dialog box appears, select the appropriate S-VOL.

S-VOL Values:

- Port – S-VOL Port

You can specify the port number with two characters. For instance, you can abbreviate CL1-A to 1A (not case-sensitive).

- GID – Host Group number
- LUN – LUN number

If you need a reference, please look at the LUN map listing in the Pair Operation tab to find the S-VOL you are looking for.

If a logical volume is an external volume, the symbol "#" appears after the LDEV number. For detailed information about external volumes, please refer to the *Universal Volume Manager User's Guide*.

If you selected more than one primary data volume, select the remote data volume for the primary data volume being displayed. The remote data volumes for the rest of the primary data volumes are automatically assigned according to the LUN.

For example, if you select three primary data volumes and select LUN01 as the S-VOL for the first primary data volume, the remote data volumes for the two other primary data volumes will be LUN02 and LUN03.

▪ **paircreate:** Journal Group information

- **M-JNL:** JNL Group in Master DKC
- **Mirror ID:** Recommended: Use Mirror ID 1 for UR pairs (see note)
- **R-JNL:** JNL Group in Remote DKC
- **CT Group:** Consistency Group Number for this M-JNL/R-JNL association (see note)
- **DKC:** Select appropriate remote DKC

The screenshot shows a command-line interface for the 'paircreate' command. It features three rows of options, each with a label in a dark blue box and a value in a light blue box with a dropdown arrow. The first row is for 'Mirror', with columns for 'M-JNL' (value 00), 'Mirror ID' (value 01), and 'R-JNL' (value 02). The second row is for 'CT Group', with a value of '00*'. The third row is for 'DKC', with a value of '10156(00), 5(USP V/USP VM), Fibre'. Above the 'DKC' row, the text 'S/N, CTRLID, Path' is displayed.

	M-JNL	Mirror ID	R-JNL
Mirror :	00	01	02
CT Group :	00*		
DKC :	10156(00), 5(USP V/USP VM), Fibre		

9

Mirror:

- M-JNL: *Master* Journal Group.
- Mirror ID: Set to 1 even if defining a 2DC configuration. This will allow 3DC configuration at a later date. (TC requires Mirror ID 0.)
- R-JNL: *Restore* Journal Group

CT Group:

Note: This is *Not* the MxN Consistency Group number. That can only be defined with Command Control Interface software.

Assign a Consistency Group number for this particular M-JNL/R-JNL association. Ensure that the CT Group selected is not in use by ShadowImage or TrueCopy Async.

If a Universal Replicator volume pair already exists in the Journal Group, the CT Group setting will have already been made and does not have to be set again. There will be an asterisk (*) next to the C/T group number. Also, the corresponding pairs of journal volumes will appear automatically.

■ paircreate: Detail information

– Initial Copy:

- Entire – all cylinders
- None – no cylinders
- Delta – create Delta Resync pairs

– Select data copy scheduling **Priority**

- 1–256 where 1 is highest

– Error Level:

- Group – All volume pairs to suspend on error
- LU – Only affected pair will suspend on error

Initial Copy : Entire

Priority : 32

Error Level : Group

M-JNL Information :
Current Mirror(s) : 0 Total Mirror(s) : 0

Set Cancel

10

■ paircreate: Set/Apply

- Volume status changes from SMPL to Copy.
- When copy completes, refresh screen to see status change to Pair.


Export Previous Next				
VOL	Status	JNL-MirrorID	S/N (LDKC)	CTRLID
CL1-C - 00 - 000 (00:00:00)	PAIR	000-01	53867(00)	€
CL1-C - 00 - 001 (00:00:01)	PAIR	000-01	53867(00)	€
CL1-C - 00 - 002 (00:00:02)	PAIR	000-01	53867(00)	€
CL1-C - 00 - 003 (00:00:03)	PAIR	000-01	53867(00)	€
CL1-C - 00 - 004 (00:00:04)	PAIR	000-01	53867(00)	€
CL1-C - 00 - 005 (00:00:05)	SMPL			
CL1-C - 00 - 006 (00:00:06)	SMPL			
CL1-C - 00 - 007 (00:00:07)	SMPL			
CL1-C - 00 - 008 (00:00:08)	SMPL			
Used Capacity: Unlimited Total Pairs : 10				
Preview				

11

pairedisplay

■ pairedisplay Command

- Displays detailed pair information
 - Sync Rate = Base Journal Copy Progress
 - Journal groups
 - Mirror ID
- Can also be displayed through Storage Navigator on remote storage system
 - **pairedisplay** is reversed.



12

Both the primary storage system administrator and the remote storage system administrator can perform this operation.

Pairedisplay Panel:

Status: Indicates the status of the pair. If the pair is split (or suspended), the suspend type is displayed. If the pair is waiting for initial copy, the word *Queuing* is displayed.

Sync.: If the volume in the primary storage system is a primary data volume, progress of an initial copy operation is displayed. If the volume in the primary storage system is a remote data volume, **Sync.** displays information in the following ways:

- If the volume pair is not split, nothing is displayed.
- If the volume pair is split and therefore is in PSUS or PSUE status, Sync. usually displays synchronization rate (that is, concordance rate) between the remote data volume before it became split and the remote data volume after it became split. For example, the synchronization rate is 100 percent if the contents of the remote data volume are the same before and after the volume pair became split.

Note: If a failure in the initial copy operation causes the volume pair to be split, Sync. displays nothing. If a failure occurs in the initial copy operation, the Detailed Information dialog box displays the phrase "Initial copy failed."

P-VOL: Indicates the primary data volume. The first line displays the port number, the GID, the LUN and LDKC:CU:LDEV (the number of LDKC, the number of CU, and the number of LDEV) of the primary data volume; the GID is a group number for a host group.

- If the primary data volume exists in the primary storage system, the first line also displays the CLPR number and the CLPR name.
- If the primary data volume is an LUSE volume, the LUN is the LDEV number of the top LDEV (that is, the smallest LDEV number in the group of LDEVs that are combined as an LUSE volume).

The second line displays the device emulation type. The third line displays the volume capacity.

Notes:

- If a volume is an external volume, the pound or gate symbol (#) appears after the LDEV number. For detailed information about external volumes, see the *Universal Volume Manager User's Guide*.
- If the primary data volume is a volume of Universal Storage Platform or Network Storage Controller, "00" is displayed as for the LDKC number.
- If there is no primary data volume in the primary storage system, the port ID, GID, and LUN of the volume that you specified in the remote storage system when creating the pair will be displayed. If you change or delete the port ID, GID, or LUN of the volume in the remote storage system, incorrect information will be displayed. So, unless you have any special circumstances, do not change or delete the port ID, GID, or LUN that you specified when creating the pair.

S-VOL: Indicates the remote data volume.

- The first line displays the port number, the GID, the LUN and LDKC:CU:LDEV (the number of LDKC, the number of CU, and the number of LDEV) of the remote data volume; the GID is a group number for a host group.
- If the remote data volume exists in the primary storage system, the first line also displays the CLPR number and the CLPR name.
- If the remote data volume is an LUSE volume, the LUN is the LDEV number of the top LDEV (that is, the smallest LDEV number in the group of LDEVs that are combined as an LUSE volume).
- The second line displays the device emulation type. The third line displays the volume capacity.

Notes:

- ♦ If a volume is an external volume, the pound (#) appears after the LDEV number. For detailed information about external volumes, see the *Universal Volume Manager User's Guide*.
- ♦ If the remote data volume is a volume of Hitachi Universal Storage Platform™ and Hitachi Network Storage Controller, “00” is displayed as the LDKC number.
- ♦ If there is no remote data volume in the primary storage system, the port ID, GID, and LUN of the volume that you specified in the remote storage system when creating the pair will be displayed. If you change or delete the port ID, GID, or LUN of the volume in the remote storage system, incorrect information will be displayed. So, unless you have any special circumstances, do not change or delete the port ID, GID, or LUN that you specified when creating the pair.
- **CLPR:** Indicates the CLPR number and the CLPR name of the volume in the primary storage system.
- **M-JNL Group:** Indicates the master journal group.
- **R-JNL Group:** Indicates the restore journal group.
- **Mirror ID:** Indicates the mirror ID.
- **CT Group:** Indicates the consistency group number.
- **S/N (CTRL ID):** displays the five-digit serial number and the control ID of the remote storage system. The control ID is enclosed by parentheses.
- **Path Type:** Indicates the channel type of the path interface between the storage systems (Fiber).

Note: In the current version, the channel type is always displayed as **Fiber**.

Initial Copy Priority: Indicates priority (scheduling order) of the initial copy operations. The value can be within the range of 1 to 256 (disabled when the status becomes PAIR).

Error Level: Indicates the range used for splitting a pair when a failure occurs. The default is **Group**.

- **Group:** If a failure occurs with a pair, all pairs in the consistency group where the pair belongs will be split.
- **LU:** If a failure occurs with a pair, only the pair will be split.

S-VOL Write: Indicates whether write I/O to the remote data volume is enabled or disabled (enabled only when the pair is split).

Other Information: Displays the following:

- **Established Time:** Indicates the date and time when the volume pair was created.
- **Updated Time:** Indicates the date and time when the volume pair status was last updated.

Refresh the Pair Operation window after this window is closed: If this check box is selected, the Pair Operation window will be updated when the Detailed Information dialog box closes.


Previous: Displays the pair status information for the previous pair in the list (the pair in the row above).

Next: Displays the pair status information for the next pair in the list (the pair in the row below)

Notes:

- The Display Filter settings (see Filtering Information in the List in the Pair Operation Window) can affect how **Previous** or **Next** is recognized.
- The list displays a maximum of 1,024 rows at once. The **Previous** and **Next** buttons on the Detailed Information dialog box can only be used for the currently displayed 1,024 rows.

- JNL Status after **paircreate**
 - Note that the Master and Restore Journals correspond to the Mirror ID specified during paircreate.



pairsplit

■ Pairsplit -r: Normal Pairsplit

- **S-VOL Write:**
 - Disabled by default
 - Enabled: Allows R/W of S-VOL after split
- **Range:** Suspend Mirror (group) or volume level
- **Suspend Mode:**
 - Flush – Send update to S-VOL
 - Purge – Convert Journal Update data to changed cylinders and suspend immediately
- **Result:** PSUS for all pairs

VOL	Status	JNL-MirrorID	S/N (LDKC)	CTRLID
CL1-C-00-000 (00:00:00)	PAIR	000-01	53867(00)	
CL1-C-00-001 (00:00:01)	PAIR	000-01	53867(00)	
CL1-C-00-002 (00:00:02)	PAIR	000-01	53867(00)	
CL1-C-00-003 (00:00:03)	PAIR	000-01	53867(00)	
CL1-C-00-004 (00:00:04)	PAIR	000-01	53867(00)	
CL1-C-00-005 (00:00:05)	PAIR	000-01	53867(00)	
CL1-C-00-006 (00:00:06)	PAIR	000-01	53867(00)	
CL1-C-00-007 (00:00:07)	PAIR	000-01	53867(00)	
CL1-C-00-008 (00:00:08)	PAIR	000-01	53867(00)	

Used Capacity: Unlimited

Set Cancel

VOL	Status	JNL-MirrorID	S/N (LDKC)	CTRLID
CL1-C-00-000 (00:00:00)	PSUS	000-01	53867(00)	
CL1-C-00-001 (00:00:01)	PSUS	000-01	53867(00)	
CL1-C-00-002 (00:00:02)	PSUS	000-01	53867(00)	
CL1-C-00-003 (00:00:03)	PSUS	000-01	53867(00)	
CL1-C-00-004 (00:00:04)	PSUS	000-01	53867(00)	
CL1-C-00-005 (00:00:05)	SMPL			
CL1-C-00-006 (00:00:06)	SMPL			
CL1-C-00-007 (00:00:07)	SMPL			
CL1-C-00-008 (00:00:08)	SMPL			

Used Capacity: Unlimited

Total Pairs: 10

Set Cancel

14

S-VOL Write: Allows you to specify whether to permit hosts to write data to the remote volume. The default is **Disable** (that is, do not permit):

- **Disable:** Hosts cannot write data to the remote volume while the pair is split.
- **Enable:** Hosts can write data to the remote volume while the pair is split. This option is available only when the selected volume is a primary volume.

Range: Allows you to specify the split range. The default is **Mirror** (all volumes in the consistency group)

- ♦ **LU:** Only the specified pairs will be split.

Note: If you select pairs with PAIR status and other than PAIR status in the same consistency group, an unexpected suspension may occur during the pair operations (pairsplit -r, pairsplit -S, and pairresync) under heavy I/O load conditions. You can estimate whether the I/O load is heavy or not from the rate of journal cache (around 30%), or if you cannot see the journal cache rate, from the frequency of host I/O. The suspend pair operations should be performed under light I/O load conditions.

Group: All pairs in the same consistency groups as the selected pairs will be split.

Note: If the following two conditions are satisfied and you select **Apply**, a warning message will be displayed and processing cannot be continued:

- The Preset list contains two or more pairs belonging to the same consistency group.
- The Range column displays Group for at least one of the above pairs.
- To be able to continue processing, do either of the following:
- Ensure that the Range column **displays LU** for all pairs in the same consistency group.
- In the Preset list, select all but one pair in the same consistency group, right-click the selected pairs, and then select **Delete**.

Suspend Mode: Allows you to specify how to deal with update data that has not been copied to the remote volume. The default is **Flush**:


- **Flush:** When you split the pair, update data will be copied to the remote volume.
- **Purge:** When you split the pair, update data will not be copied to the remote volume. Instead it will convert to changed cylinders in P-VOL differential bitmap. If you resync the pair later, the changed P-VOL data will be copied to the remote volume.

Set: Applies the settings to **the Preset** list in the Pair Operation panel.

pairresync

■ pairresync

1. In the list, select and right-click the pair that you want to restore.
 - The pair status must be PSUS or PSUE.
2. From the pop-up menu, select **Pairresync**.
3. Range: Specify **Mirror** or **LU**
4. Resync Mode: Default is **Normal**. May also be Delta or Return to Standby. See Student notes.
5. Result: **PAIR** status



Export Previous Next				
VOL	Status	JNL-MirrorID	S/N (LDKC)	CTRLID
CL1-C-00-000 (00:00:00)	PSUS	000-01	53867(00)	
CL1-C-00-001 (00:00:01)	PSUS	000-01	53867(00)	
CL1-C-00-002 (00:00:02)	PAIR	000-01	53867(00)	
CL1-C-00-003 (00:00:03)	PAIR	000-01	53867(00)	
CL1-C-00-004 (00:00:04)	PAIR	000-01	53867(00)	
CL1-C-00-005 (00:00:05)	SMPL			
CL1-C-00-006 (00:00:06)	SMPL			
CL1-C-00-007 (00:00:07)	SMPL			
CL1-C-00-008 (00:00:08)	SMPL			

Used Capacity: Unlimited Total Pairs: 10

Preview

15

Note: If the primary or remote Universal Storage Platform and Network Storage Controller storage system is powered off and its backup batteries are fully discharged while pairs are suspended, the P-VOL and S-VOL bitmaps will not be retained. In this unlikely case, the affected storage system will mark all cylinders of all suspended volumes as modified, so that the primary storage system will perform the equivalent of an entire initial copy operation when the pairs are resumed.

The Hitachi Universal Storage Platform™ V and Hitachi Universal Storage Platform™ VM storage systems have the capability of storing Shared Memory (containing the P-VOL and S-VOL bitmaps) to the SVP hard drive. Shared Memory can thus be recovered in the event that battery backup has not retained power to the Shared Memory.

If any pair was suspended due to an error condition (use the Pairedisplay panel to view the suspend type), make sure that the error condition has been removed. The primary storage system will not resume the pairs until the error condition has been removed.

The Pairresync panel displays the following:

Range: Allows you to specify the restore range. The default setting is **Mirror**. The default setting cannot be changed when in **HOLD** or **HLDE** status.

- **LU:** Only the specified pairs will be restored.
- **Group:** All pairs in the same consistency groups as the selected pairs will be restored.

Note: If the following two conditions are satisfied and you select **Apply**, a warning message will be displayed and processing cannot be continued:

- The **Preset** list contains two or more pairs belonging to the same consistency group.
- The **Range** column displays Group for at least one of the above pairs.

To be able to continue processing, do either of the following:

- Ensure that the **Range** column displays **LU** for all pairs in the same consistency group.
- In the **Preset** list, select all but one pair in the same consistency group, right click the selected pairs, and then select **Delete**.

Priority: Allows you to specify the desired priority (1-256) (scheduling order) for the pair-restoring operations.

Note: If the **Range** is **Group**, you cannot change the **Priority** option.

DKC: Indicates the storage system.

Resync Mode: Indicates the processing after recovery of the pairs.

- **Normal:** Split pair whose status is **PSUS** or **PSUE** will be recovered.
- **Delta:** Delta resync operation will be performed.
- **Return to standby:** The status of pairs will be recovered from **HLDE** to **HOLD**.


Error Level: Allows you to specify the range used for splitting a pair when a failure occurs:

- **Group:** If a failure occurs with a pair, all pairs in the consistency group where the pair belongs will be split.
- **LU:** If a failure occurs with a pair, only the pair will be split.

Note: If the **Range** is **Group**, you cannot change the **Error Level** option.

Change Pair Option

- Allows change of **Error Level** from Mirror to LU (VSP only)




pairsplit

■ pairsplit -S Options

– Deleting a Pair

- **Range:** Mirror or LU
If LU is selected, only that pair is deleted.



Export Previous Next				
VOL	Status	JNL-MirrorID	S/N (LDK)	CTRLID
CL1-C-00-000 (00:00:00)	SMPL			
CL1-C-00-001 (00:00:01)	PAIR	000-01	53867(00)	
CL1-C-00-002 (00:00:02)	PAIR	000-01	53867(00)	
CL1-C-00-003 (00:00:03)	PAIR	000-01	53867(00)	
CL1-C-00-004 (00:00:04)	PAIR	000-01	53867(00)	
CL1-C-00-005 (00:00:05)	SMPL			
CL1-C-00-006 (00:00:06)	SMPL			
CL1-C-00-007 (00:00:07)	SMPL			
CL1-C-00-008 (00:00:08)	SMPL			
Used Capacity : Unlimited Total Pairs : 8				

17

The Pairsplit-S panel displays the following:

Range: Allows you to specify the delete range. The default is Mirror.

- **LU:** Only the specified pairs will be deleted.

If you select pairs with PAIR status and other than PAIR status in the same consistency group, an unexpected suspension may occur during the pair operations (Pairsplit-r, Pairsplit-S, and Pairresync) under heavy I/O load conditions. You can estimate whether the I/O load is heavy or not from the rate of journal cache (around 30%), or if you cannot see the journal cache rate, from the frequency of host I/O. The pair operations should be performed under light I/O load conditions.

- **Group:** All pairs in the same consistency groups as the selected pairs will be deleted.

Note: Do not use this option when deleting pairs at the remote storage system during disaster recovery.

Note: If the following two conditions are satisfied and you click **Set**, a warning message is displayed and processing cannot continue:

- ♦ The Preset list contains two or more pairs belonging to the same consistency group
- ♦ The Range column displays Group for at least one of the above pairs

To be able to continue processing, do either of the following:

- Ensure that the Range column displays LU for all pairs in the same consistency group.
- In the Preset list, select all but one pair in the same consistency group, right-click the selected pairs, and then select **Delete**.

Delete Mode: Allows you to specify whether to delete the pairs forcibly. When the status of the pairs to be deleted is SMPL or Deleting, the default setting is Force. Otherwise, the default setting is Normal.

- **Force:** The pairs will forcibly be deleted even if the primary storage system is unable to communicate with the remote storage system.



Note: If issued on S-VOL, FORCE will delete S-VOLS without regard for (or changing) P-VOL status. To recover pair, delete P-VOL and issue paircreate again.

- **Normal:** The pairs will be deleted only if the primary storage system is able to change the pair status of the primary and remote volumes to SMPL
- **Set:** Applies the settings to the Preset list in the Pair Operation panel

■ pairsplit -S Options

– Deleting a Pair

- **Range:** Mirror
- **Delete Mode:**
 - Normal
 - **Force** – When issued at remote site on S-VOLs, forces S-VOLs to simplex without regard for P-VOL status




Usage Monitor

- Monitoring operations are performed by Usage Monitor
 - Collects I/O Stats for all LDEVs on Hitachi enterprise storage systems
- Functions
 - Start and stop monitoring
 - Display Usage Graph
 - Export Usage Data Monitor File

- Start Usage Monitoring on VSP


1. On Storage Navigator Browser screen, navigate to **Performance Monitor**.
2. Scroll right and select **Edit Monitoring Switch**.
3. Set Monitoring Switch to **Enable**.



20

If you set 1 minute for **Gathering Interval**, the sampling data will be held one day. If you set 15 minutes for **Gathering Interval**, the sampling data will be held 15 days. When **Gathering Interval** is changed, the data obtained before changing is deleted.


- Start Usage Monitoring on USP V
 1. Navigate to **Performance Manager > Monitoring Options**.
 2. Select **Enable** in the **Monitoring Switch Current Status** box.
 3. Set desired Gathering Interval.
 4. **APPLY** the change.
- USP – Enable and apply Usage Monitor directly in **URz Monitor** page.



21

If you set 1 minute for **Gathering Interval**, the sampling data will be held one day. If you set 15 minutes for **Gathering Interval**, the sampling data will be held 15 days. When **Gathering Interval** is changed, the data obtained before changing is deleted.

1. Select the **Usage Monitor** tab.
2. Right-click on one of the quadrants.
3. Select **Display Item** from the dialog.




22

If you set 1 minute for **Gathering Interval**, the sampling data will be held one day. If you set 15 minutes for **Gathering Interval**, the sampling data will be held 15 days. When **Gathering Interval** is changed, the data obtained before changing is deleted.

■ Display Item panel


1. Select appropriate item in **Select Volume** section.
2. In the **Monitor Data** box, select the appropriate items. You must select at least one.
3. Click **Set** to close the Display Item panel.

The Usage Operations panel now displays a graph showing the selected I/O statistics data for the selected LUs.



23

■ Sample graphs



24

History

- LIFO file – **Export** function saves operation history into a .tgz file
 - Same function on earlier systems is called **Snapshot**.

Journal Operation Pair Operation DKC Operation Usage Monitor History Optional Operation				
Universal Replicator History				
Status :	Complete			
Last Update Time :	2010/08/05 09:00:02			
Page :	1 / 1			
		Export	Previous	Next
Operation Time	Operation	Journal	Mirror ID	VC
2010/07/31 15:27:38	Status Change by MCU(PAIR to SMPL)	002	00	00.4
2010/07/31 15:27:38	Pairsplit-S Complete	001	00	00.4
2010/07/31 15:27:15	Status Change by MCU(PSUS/PSUE to SMPL)	002	00	00.4
2010/07/31 15:27:15	Pairsplit-S Complete	001	00	00.4
2010/07/31 15:27:15	Pairsplit-S Start	001	00	00.4
2010/07/31 15:27:15	Pairsplit-S Start	001	00	00.4
2010/07/31 15:06:52	Pairresync Complete	001	00	00.4
2010/07/31 15:06:51	Status Change by MCU(COPY to PAIR)	002	00	00.4
2010/07/31 15:06:32	Status Change by MCU(PSUS/PSUE to COPY)	002	00	00.4
2010/07/31 15:06:33	Pairresync Start	001	00	00.4
2010/07/31 15:05:01	Status Change by MCU(PAIR to PSUS/PSUE)	002	00	00.4
2010/07/31 15:05:01	Status Change by MCU(PAIR to PSUS/PSUE)	002	00	00.4
2010/07/31 15:05:01	Pairsplit-r Complete	001	00	00.4
2010/07/31 15:05:01	Pairsplit-r Complete	001	00	00.4
2010/07/31 15:04:42	Pairsplit-r Start	001	00	00.4
2010/07/31 15:04:42	Pairsplit-r Start	001	00	00.4
2010/07/31 15:04:39	Paircreate Complete	001	00	00.4
2010/07/31 15:04:39	Paircreate Complete	001	00	00.4
2010/07/31 15:04:15	Pair definition	001	00	00.4
2010/07/31 15:04:15	Pair definition	001	00	00.4

25

Troubleshooting Suspended Pairs

Suspend Type	Applies To	Description	Corrective Action
PSUE, by RCU	Primary data volume	The primary storage system suspended a pair because the primary storage system detected an error condition at the remote storage system. The suspend type for the remote data volume is by MCU.	Clear the error condition at the remote storage system or remote data volume. If you need to access the remote data volume, delete the pair from the primary storage system. If data in the remote data volume has been changed, delete the pair from the primary storage system and then recreate the pair by using the Paircreate panel. If data in the remote data volume has not been changed, restore the pair from the primary storage system.
PSUE, S-VOL Failure	Primary data volume	The primary storage system detected an error during communication with the remote storage system or detected an I/O error during update copy. In this case, the suspend type for the remote data volume is usually by MCU.	Check the path status on the DKC Status panel (see Table 10-2). Clear any error conditions at the remote storage system and the remote data volume. If you need to access the remote data volume, delete the pair from the remote storage system. If data in the remote data volume has been changed, delete the pair from the primary storage system and then recreate the pair by using the Paircreate panel. If data in the remote data volume has not been changed, restore the pair from the primary storage system.

26

Suspend Type	Applies To	Description	Corrective Action
PSUE, MCU IMPL	Primary data volume, remote data volume	The primary storage system could not find valid control information in its nonvolatile memory during the IMPL procedure. This error occurs only if the primary storage system is without power for more than 48 hours (power failure and fully discharged batteries).	Restore the pair (pairresync) from the primary storage system. The primary storage system will perform an initial copy operation in response to the pairresync request, so that the entire primary data volume will be copied to the remote data volume.
PSUE, Initial Copy Failed	Primary data volume, remote data volume	The primary storage system suspended this pair during the initial copy operation. The data on the remote data volume is not identical to the data on the primary data volume.	Delete the pair from the primary storage system. Clear all error conditions at the primary storage system, primary data volume, remote storage system, and remote data volume. Restart the initial copy operation by using the Paircreate panel.
PSUE, MCU P/S-OFF	Secondary data volume	The primary storage system suspended all Universal Replicator pairs because the primary storage system was powered off.	None. The primary storage system will automatically restore these Universal Replicator pairs when the primary storage system is powered on.

■ SIMs

Classification	Causes of Suspension	SIM	Recovery Procedure
Primary storage system hardware or remote storage system hardware	Hardware redundancy has been lost due to some blockade condition. As a result, one of the following could not complete: primary-remote storage system communication, journal creation, copy operation, restore operation, staging process, or destaging process. Journals cannot be retained because some portion of the cache memory or the shared memory has been blocked due to hardware failure. The primary storage system failed to create and transfer journals due to unrecoverable hardware failure. The remote storage system failed to receive and restore journals due to unrecoverable hardware failure. The drive parity group was in correction-access status while the Universal Replicator pair was in COPY status.	DC0x DC1x DC2x	According to SIM, remove the hardware blockade or failure. Restore the failed volume pairs (pairresync).

28

■ SIMs

Classification	Causes of Suspension	SIM	Recovery Procedure
Communication between the primary and remote storage systems	Communication between the storage systems failed because the remote storage system or network relay devices were not running. Journal volumes remained full even after the timeout period elapsed.	DC0x, DC1x	Remove the failure from the primary and remote storage systems or the network relay devices. If necessary, increase resources as needed (e.g., the amount of cache, the number of paths between primary and remote storage systems, the parity groups for journal volumes, etc.). Restore the failed pairs (pairresync).
RIO overload or RIO failure	An unrecoverable RIO (remote I/O) timeout occurred because the storage system or networks relay devices were overloaded. Or, RIO could not be finished due to a failure in the storage system.	DC2x	Delete failed pairs (pairsplit-S). If necessary, increase resources as needed (e.g., the amount of cache, the number of paths between primary and remote storage systems, the parity groups for journal volumes, etc.). Re-establish failed pairs (paircreate).
Planned power outage to the primary storage system	The Universal Replicator pairs were temporarily suspended due to a planned power outage to the primary storage system.	DC8x	No recovery procedure is required. The primary storage system will automatically remove the suspension condition when the storage system is powered on.

29

Module Review

1. List the Universal Replicator commands provided by Storage Navigator.
2. With a pair split, what is the default Read/Write attribute set on an S-VOL?
3. True or False: A normal pair split is performed using the **pairsplit -S** option.
4. Describe the steps to start the Usage Monitor.
5. True or False. A volume pair can be resynchronized if the storage system has reported an error and the error has not been fixed. Why or why not?

6. Hitachi Command Control Interface Configuration and Operations

Module Objectives

- Upon completion of this module, the learner should be able to:
 - Describe the history of Hitachi Command Control Interface (CCI, sometimes called RAID Manager)
 - Identify CCI software components
 - Describe the checklist used to configure
 - Create a physical configuration
 - Create and revise the configuration files
 - Describe the command set
 - Identify CCI Microsoft Windows subcommands
 - Discuss configuration setting commands
 - Perform troubleshooting and corrective actions for operations

Overview

- Origins:

- Hitachi Open Remote Copy (HORC) is the original name for Hitachi TrueCopy Remote Replication software.
- Hitachi Open Remote Copy Manager (HORCM) is the original name for the management software now called Command Control Interface (CCI).
- Hitachi Multi-RAID Coupling Facility (HMRCF) is the original name for the Hitachi ShadowImage Heterogeneous Replication software.
- Both products were managed with HORCM, now CCI.
- The acronyms HORCM and MRCF are still used internally by CCI.

- CCI hardware components

- Command Device
 - Communication path for CCI
 - Accept commands from CCI for execution by replications
- Remote Command Device
 - Externally attached command device to allow CCI to remotely manage a replication
- Volumes to be used as replication P-VOLs and S-VOLs
- Replication Licenses must be installed on all participating storage systems.

- All Hitachi enterprise class storage systems support CCI software operations.

- CCI software versions relate to storage microcode level.
- Check documentation for current versions.


- Latest CCI documentation for all enterprise platforms
 - **MK-90RD7008** *Hitachi Command Control Interface Installation and Configuration Guide*
 - **MK-90RD7009** *Hitachi Command Control Interface Command Reference*
 - **MK-90RD7010** *Hitachi Command Control Interface User and Reference Guide*

5

- CCI software
 - Installs on SAN-attached servers
 - Communicates with the storage systems using the command devices
 - No communication with the devices containing data to be replicated
- CCI Software Internals
 - HORCM Instances – two required (minimum)
 - Run as Services
 - Primary instance manages P-VOLs
 - Secondary instance manages S-VOLs
 - HORCM Configuration Files – two required (minimum)
 - Define servers and communication paths
 - Defines location of Command Device
 - Define devices used by the replication
 - Define location of server running other Instance
 - HORCM Commands
 - Monitor and control copy operations

6

Two Servers – Two HORCM Instances



7

Shown here are the four components mentioned on the previous slide. The relationship between these components include:

- There are always at least two instances, a sending instance and a receiving instance.
 - ♦ Instance 0 is the sending instance and Instance 1 is the receiving instance.
- Each instance relies on a configuration file in order to communicate with the other instance, as well as to communicate with the system.
- The configuration file defines the volumes that will be *paired up*.

If you have two instances, you will have two corresponding configuration files.

When a command is issued, usually via a script, the instance sends the command to the CMD Device. The system then actuates the command.

CCI Checklist

- CCI Checklist Steps
 - Install Command Control Interface software
 - Map P-VOLs and S-VOLs as LUNs to desired host ports
 - Note that CCI only needs access to Command Device
 - CCI does not need access to P-VOLs and S-VOLs
 - Create and map CMD Device to port accessible to CCI
 - Edit and save (or create new) HORCM configuration files
 - Minimum two files will be required
 - Edit the Services File
 - Set Environment Variables (optional)
 - Start Instances
 - Issue Commands

Configuration

- Install CCI
 - Microsoft Windows Server
 - Download latest version
 - Copy to appropriate folder
 - Run **Setup.exe**
 - UNIX
 - Download latest version to appropriate directory
 - Run **RMinstsh** script

- All Hitachi replications required LUNs (mapped LDEVs)
- If not available, create candidate P-VOLS and S-VOLs
 - Use procedures appropriate to the model of storage system, map desired LDEVs to the appropriate ports
 - Record Port, Host Group, LUN numbers
 - Release any ShadowImage **RESERVED** volumes
 - Reserved Volumes cannot be used by CCI.

10

CCI depends on S-VOLs being defined ahead of time using LUN Manager. You cannot define LUNs through CCI.

Reserved Volumes cannot be used by CCI. Remember that Storage Navigator allows setting of ShadowImage S-VOL reserves. However, reserving S-VOLs blocks access to the LDEV by the host. Since CCI resides on the host, the S-VOLs that were reserved through the GUI will be blocked.

Make sure S-VOL LUNs are mapped to the appropriate ports. An example backup scenario might be a P-VOL is mapped to production server ports and S-VOL is mapped to backup server ports.

Configuration: Command Device

- Create a Command Device
 - Use appropriate procedure for your storage platform
 - Create minimum size volume (46MB)
 - Map the device to the correct ports
 - Mark it as the Command Device
 - Right-click and select CMD DEV

Note: Command Device is a raw device. It is either labeled (UNIX) or has a signature (Microsoft Windows). No file systems or partitions are created in this process.

11

In order to run ShadowImage via the command prompt, a Command Device is required. 46MB is the smallest OPEN-V LDEV that can be created. The volume designated as the command device is used only by the primary storage system and is a raw device. Multiple command devices in one storage system are allowed.

Procedure: Use Storage Navigator VLL function to create the smallest possible OPEN-X volume. Set Command Device attribute ON with LUN Manager.

Make sure that the volume to be selected as the command device does not contain any user data. Otherwise it will be inaccessible to the UNIX/PC server host. A LUSE volume cannot be used as a command device.

Configuration: VSP Command Device

- VSP Command Device using Storage Navigator browser window

Step 1

Step 3

Step 2

RAID Level	Emulation Type	Capacity	Provisioning Type	Attribute	Num Path
5(7D+1P)	OPEN-V CVS	4.00 GB	Basic	-	
5(7D+1P)	OPEN-V CVS	4.00 GB	Basic	-	
5(7D+1P)	OPEN-V CVS	4.00 GB	Basic	-	
5(7D+1P)	OPEN-V CVS	1678.80...	Basic	-	
5(7D+1P)	OPEN-V CVS	0.09 GB	Basic	-	
5(7D+1P)	OPEN-V CVS	0.09 GB	Basic	-	
5(7D+1P)	OPEN-V CVS	0.09 GB	Basic	-	
5(7D+1P)	OPEN-V CVS	0.09 GB	Basic	-	
5(7D+1P)	OPEN-V CVS	0.09 GB	Basic	-	
5(7D+1P)	OPEN-V CVS	0.09 GB	Basic	-	
5(7D+1P)	OPEN-V CVS	0.09 GB	Basic	-	
5(7D+1P)	OPEN-V CVS	0.09 GB	Basic	-	

- Set VSP Command Device using Storage Navigator browser window

Step 1

Step 2

Step 3

Command Device : ☒ Enable ☐ Disable

Command Device Attributes :

Command Device Security : ☐ Enable ☒ Disable

User Authentication : ☐ Enable ☒ Disable


Device Group Definition : ☐ Enable ☒ Disable

Task Name : 110209-EditCommandDevices (Max. 32 Characters)

LDEV ID	LDEV Name	Parity Group ID	Emulation Type	Capacity	Provisioning Type	Command Device
00:07:00		1-5	OPEN-V CVS	0.09 GB	Basic	Enabled

Configuration: Remote Command Device

- CCI Instance can operate replication in remote storage system through the remote command device mapped in the local storage system.
- Useful in 3DC Delta Resync and 2DC Cascade environments.
- Usage will be discussed in Delta Resync Module.



■ VSP Remote Command Device

File Go Help 30 minute(s) remaining in session Reset Logged in as: replicationuser

Journal Operation Pair Operation DKC Operation Usage Monitor History Optional Operation

Universal Replicator Journal Operation

Journals

JNL (LDKC)	Attribute	Status	Mirror ID	S/N (LDKC)	Pair JNL
000(00)		Initial	00		
000(00)		Initial	01		
000(00)		Initial	02		
000(00)		Initial	03		
001(00)		Initial	00		
001(00)		Initial	01		
001(00)		Initial	02		
001(00)		Initial	03		
002(00)					
003(00)					
004(00)					
005(00)					
006(00)					

Preview

Hitachi Storage Navigator

Assign Remote Command Device

Mirror ID	Remote Command Device
Mirror 1: 00	Not Use
Mirror 2: Not Set	Not Use

Set Cancel

Configuration Files

- Edit HORCM.CONF sample configuration file
 - Delete comments.
 - Leave these parameters:
 - HORCM_MON
 - HORCM_CMD
 - HORCM_LDEV (or HORCM_DEV if desired)
 - HORCM_INST
 - HORCM_CTQM (for MxN Consistency Group) see note
 - Minimum of two HORCM files are required.
 - Numbered instances configured by **horcmn.conf** (for example, horcm0.conf).
 - Best Practice: Always configure P-VOLs in even instance numbers and S-VOLs in odd instance numbers regardless of whether or not both instances are on same server.

16

MxN Consistency Group will be covered in a later module.

■ Sample HORCM configuration file (CCI V01-17-03 and later)

HORCM_MON Identifies the primary host	HORCM_MON #IP 172.16.0.44	service horcm0	poll 10003000	timeout
HORCM_CMD Identifies the command device	HORCM_CMD /dev/rdisk/c4t0d255s2 \\.\CMD-30095			
HORCM_LDEV ** Identifies the volumes	HORCM_LDEV #dev_group dev_name Serial# CU:LDEV(LDEV#) MU# URGRP1 dev1 30095 00:00 1 URGRP1 dev2 30095 00:01 1			
HORCM_INST Identifies the remote host	HORCM_INST #dev_group IP URGRP1 172.16.0.44	service horcm1		

New parameters have been added to support Open UR MxN Consistency Group. See MxN Consistency Group Module for details.s

17

**CCI Version 01-17-03 introduced the HORCM_LDEV parameter. Absolute LUN numbers are no longer required, which eases the work required to produce valid HORCM configuration files.

■ HORCM.CONF File Overview

- Provides a definition of hosts, command devices, groups, and volumes to the Command Control Interface instance
- After CCI installation, a sample text file is located at /HORCM/etc/horcm.conf.
- Edit the sample file to remove all commented lines except for the lines associated with these parameters.
- Create two files named horcm0.conf and horcm1.conf and copy to:
 - /etc directory (UNIX)
 - /HORCM/etc and /WINNT (NT) or /WINDOWS (XP)
- Assigned numbers are arbitrary. Class examples are:
 - horcm0.conf controls P-VOLs
 - horcm1.conf controls S-VOLs

Ensure that there are no hidden characters or spaces at the end of lines.

When saving horcm.conf files in Microsoft Windows, do **not** save as .txt files.

CCI Configuration

- Command Device
 - Accepts CCI commands that are executed by the storage system.
 - The Command Device also returns status requests to the UNIX/PC host.
 - Used only by the CCI and is a raw device.

- Identifying the Command Device in Windows
 - Newer CCI versions can use this format to identify the Command Device.

 - `\\.\CMD-30095`
 - Use the first Command Device found on Subsystem Serial Number 30095

 - `\\.\CMD-30095-250`
 - Use LDEV 250 (decimal)

 - `\\.\CMD-30095-250-CL1-A-1`
 - Port-Host Group

CCI Configuration: Command Device (Solaris)

- Identifying the Command Device in Solaris
 - Older method: Issue **format** command. Find disk with the **CM** string
For example: `c2t0d1 <HITACHI - Open-V - CVS - CM>`
- Defining the Command Device in the horcm.conf files
 - Old Method: Raw disk pathname - `/dev/rdisk/c2t0d1s2`
 - New methods:
 - `\\.\CMD-30095:/dev/rdisk/`
 - Use first available command device
 - `\\.\CMD-30095-250:/dev/rdisk/`
 - Use LDEV#250
 - `\\.\CMD-30095-250-CL1-A-1:/dev/rdisk/`
 - Full specification for S/N 30095, LDEV 250 (decimal) connected to Port CL1-A, host group 1

Configuration Files

- For **CCI version 01-16-00** and **earlier**, find absolute LUN numbers
 1. Enter correct information in **HORCM_MON** and **HORCM_CMD** parameters.
 2. Leave **HORCM_DEV** and **HORCM_INST** parameters commented out.
 3. Start the instance (example: # horcmstart 0).
 4. Execute **raidscan -p portID -fx**.

```
HORCM_MON
#IP          service poll    timeout
172.16.0.44  horcm0      1000   3000


HORCM_CMD
/dev/rdisk/c4t0d255s2    /dev/rdisk/c4t0d254s2

#HORCM_DEV
#dev_group dev_name port  TID      LUN  MU#

#HORCM_INST
#dev_group  IP          service
```


- Execute **raidscan -p port -fx** to find CCI absolute LUN numbers:

- CCI scans ports and assigns absolute LUN numbers as they are detected
- CCI assigns its own TIDs and they may be different from the Fibre Channel assigned TIDs



Example: raidscan -p CL1-A -fx

LUN Manager		CCI Absolute LUN #s
LUN	CU:LDEV	
0000	00:02	0
0001	00:40	1
0002	00:41	2
0000	00:00	3
0001	10:00	4
0002	3E:00	5
0003	3F:00	6
0004		
0005		
0006		
0007	00:01	7
0008	10:01	8

22

LUNs on a Group are mapped for the host system so that the remote ID and LUN which is used by CCI will be different from the remote ID and LUN shown by the host system. Use the remote ID and LUN indicated by the **raidscan** command.


Example of **raidscan** Command

raidscan -p CL1-A -fx

Port# / ALPA	TID#	LU#	NUM(LDEV#)
CL1-A / DE/	1	5-0	1(05) LUN 5, Mirror unit 0
CL1-A / DE/	1	5-1	1(05) LUN 5, Mirror unit 1
CL1-A / DE/	1	5-2	1(05) LUN 5, Mirror unit 2
CL1-A / DE/	1	7-0	3(25, 26, 27) 3 Volume LUSE LUN

Configuration Files: Group and Device names


- Naming the groups and devices
 - Each DEV_GROUP is made up of one or more DEV_NAMES.
 - Try to be simple yet descriptive.
 - Names can be up to 31 characters.
- Match names to the S-VOLs. Ensure that:
 - Device names match for both the P-VOLs and the S-VOLs.
 - Device group names must match for both P-VOLs and S-VOLs.



Configuration Files: Restrictions

- Configuration File Considerations
 - Anytime you make a change or update to the configuration file, stop the HORCM instance, then restart the it.
 - Watch out for extra spaces and tabs in the file.
 - Group names and device names are:
 - Case sensitive
 - Must be unique within the file
 - Must match in both files
 - Make sure the HORCM configuration files have been saved in the correct directory.
 - UNIX: /etc
 - Windows: Highest level Windows folder

Configuration Files: Mirror IDs



25

Configuration Files

- For CCI V01-17-03 and later, use **HORCM_LDEV** parameter to create HORCM files for both P-VOLs and S-VOLs.

Note: Universal Replicator software MU# must be **1** in both configuration files

HORCM_MON				
#ip_address	service	poll(10ms)	timeout(10ms)	
10.17.105.4	horcm0	1000	3000	

HORCM_CMD				
#dev_name				
\\CMD-10145				

HORCM_LDEV				
#dev_group	dev_name	Serial#	CU:LDEV(LDEV#)	MU#
urgrp	pair1	10145	00:00	1
urgrp	pair2	10145	00:01	1

HORCM_INST		
#dev_group	ip_address	service
urgrp	10.17.105.5	horcm1

HORCM_MON				
#ip_address	service	poll(10ms)	timeout(10ms)	
10.17.105.5	horcm1	1000	3000	

HORCM_CMD				
#dev_name				
\\CMD-10156				

HORCM_LDEV				
#dev_group	dev_name	Serial#	CU:LDEV(LDEV#)	MU#
urgrp	pair1	10156	02:00	1
urgrp	pair2	10156	02:01	1

HORCM_INST		
#dev_group	ip_address	service
urgrp	10.17.105.4	horcm0

26

When naming these files, they should be called *horcm + decimal number + .conf*. This is necessary for CCI to detect which configuration files to use during the process. CCI needs the specified number to start the instance and recognize which configuration file to use with that instance.

Mirror Unit Number (MU#) for 3DC TrueCopy/Universal Replicator environment:

Defines the mirror unit number (0 - 3) of one of four possible TrueCopy/Universal Replicator bitmap associations for an LDEV in a Cascaded or Multi-remote 3DC environment. If this number is omitted, it is assumed to be zero (0).

- The MU# for TrueCopy must specify either *blank* or *0*.
- The MU# for Universal Replicator must specify *h1* in both horcm files.

Best practice: Always use Mirror ID h1 for UR. This will allow a 2DC UR replication to be converted easily to 3DC at a later date.

Configuration: Services File

- Edit the Services File
 - Solaris: `/etc/services`
 - Microsoft Windows: `C:\Windows\system32\drivers\etc\services`
 - This step can be omitted if the default service number of 52323 is coded in the conf files.

nfs	2049/udp		# sun nfs
knetd	2053/tcp		# Kerberos de-multiplexor
eklogin	2105/tcp		# Kerberos encrypted rlogin
rmt	5555/tcp	rmt	
mtb	5556/tcp	mtbd	# mtb backup
man	9535/tcp		# remote man server
w	9536/tcp		
mantst	9537/tcp		# remote man server, testing
bnews	10000/tcp		
rscs0	10000/udp		
queue	10001/tcp		
rscs1	10001/udp		
poker	10002/tcp		
rscs2	10002/udp		
gateway	10003/tcp		
rscs3	10003/udp		
remp	10004/tcp		
rscs4	10004/udp		
rscs5	10005/udp		
rscs6	10006/udp		
rscs7	10007/udp		
rscs8	10008/udp		
rscs9	10009/udp		
rscsa	10010/udp		
rscsb	10011/udp		
qmaster	10012/tcp		
qmaster	10012/udp		
horcm0	11000/udp		# horcm instance 0
horcm1	11001/udp		# horcm instance 1

**For Microsoft Windows:
Make sure to include a CR/LF
at the end of the last line.**

27

Edit Services file to add the horcm instances.

Configuration: mkconf

- CCI Configuration Files – **mkconf** Command Tool Functionality
- **UNIX Procedure:**
 1. Make a configuration file for only HORCM_CMD by executing **inqraid -sort -CM -CLI**.
 2. Start a HORCM instance with descriptions for HORCM_DEV and HORCM_INST commented out.
 3. Make a configuration file that include HORCM_DEV and HORCM_INST by executing **raidscan -find conf** from a special file (raw device file) provided via STDIN.
 4. Uncomment HORCM_DEV and HORCM_INST parameters. Start the HORCM instance for verification of the configuration file.
 5. Execute **raidscan -find verify** to verify the correspondence of the device file and the configuration file.
 6. Execute **mkconf** command tool.

28

Notes:

A unit ID is added to the Serial# order. If two or more command devices exist in the storage system, then this option selects the multiple device files.

- If the remote device is the command device, then the remote device is suppressed as comment as shown below:

```
# ERROR [CMDDEV] /dev/rdisk/c23t0d7 SER = 61456 LDEV = 259 [ OPEN-3-CM ]
```

- If the remote device has shared an LDEV among multiple device files and an LDEV is displayed by another remote device already, then its remote device is suppressed as comment as shown below:

```
# ERROR [LDEV LINK] /dev/rdisk/c24t0d3 SER = 61456 LDEV = 195 [FIBRE FCTBL = 4]
```

- If the remote device does not have a valid MU#, then its remote device is suppressed as comment as shown below:

```
# ERROR [INVALID MUN (2 < 1)] /dev/rdisk/c24t0d3 SER = 61456 LDEV = 195 [ OPEN-3 ]
```

- If the remote device has been mixing with difference between RAID TYPE, then its remote device is suppressed as comment as shown below:

```
# ERROR [MIXING RAID TYPE] /dev/rdisk/c24t0d3 SER = 61456 LDEV = 195  
[ OPEN-3 ]
```


■ CCI Configuration Files: **mkconf** usage

- Example: **mkconf.sh -g ORA -i 0 -m 1** will use as input the results of the **raidscan** command and create a horcm0.conf file for the dev_group ORA with Mirror Unit numbers set to 0.

```
HORCM_MON
#ip_address      service      poll(10ms)      timeout(10ms)
127.0.0.1        52323              1000             3000

HORCM_CMD
#dev_name
#UnitID 0 (Serial# 30095)
/dev/rdsk/c23t3d0

HORCM_DEV
#dev_group      dev_name      port#      TargetID      LU#      MU#
# /dev/rdsk/c23t0d0 SER = 30095 LDEV = 192 [ FIBRE FCTBL = 4 ]
ORA             ORA_000      CL2-B      0             0             1

HORCM_INST
#dev_group      ip_address      service
ORA             127.0.0.1      52323
```

29

Actual output of **mkconf** usage on Microsoft Windows server is below. Note that the procedure differs from the UNIX procedure shown in a previous slide. Some lines in the Device File were deleted for clarity.

C:\HORCM\etc>

C:\HORCM\etc>raidscan -p cl1-a-0 -CLI -fx

PORT# /ALPA/C TID# LU# Seq# Num LDEV# P/S Status Fence P-Seq# P-LDEV#

```
CL1-A-0 ef 5 1 0 20068 1 40 SMPL - - - -
CL1-A-0 ef 5 1 1 20068 1 41 SMPL - - - -
CL1-A-0 ef 5 1 2 20068 1 42 SMPL - - - -
CL1-A-0 ef 5 1 3 20068 1 43 SMPL - - - -
CL1-A-0 ef 5 1 4 20068 1 44 SMPL - - - -
CL1-A-0 ef 5 1 5 20068 1 45 SMPL - - - -
CL1-A-0 ef 5 1 6 20068 1 46 SMPL - - - -
CL1-A-0 ef 5 1 7 20068 1 47 SMPL - - - -
CL1-A-0 ef 5 1 8 20068 1 48 SMPL - - - -
CL1-A-0 ef 5 1 9 20068 1 49 SMPL - - - -
```

```
CL1-A-0 ef 5 1 10 20068 1 4a SMPL - - - -
```

```
CL1-A-0 ef 5 1 11 20068 1 4b SMPL - - - -
```

```
CL1-A-0 ef 5 1 12 20068 1 700 SMPL - - - -
```

```
C:\HORCM\etc>raidscan -p cl1-a-0 -CLI -fx >> temp
```

Note that **raidscan** output is sent to a temp file. Next we pipe temp to the **mkconf** command as shown below:

```
C:\HORCM\etc>type temp | c:\horcm\tool\mkconf.exe -g test -i 5 -m 0
```

starting HORCM inst 5

HORCM inst 5 starts successfully.

HORCM Shutdown inst 5 !!!

A CONFIG file was successfully completed.

starting HORCM inst 5

HORCM inst 5 starts successfully.

DEVICE_FILE	Group	PairVol	PORT	TARG	LUN	M	SERIAL	LDEV
Harddisk1	test	test_000	CL1-A	1	0	0	20068	64
Harddisk5	test	test_001	CL1-A	1	4	0	20068	68
Harddisk1	test	test_000	CL1-A	1	0	0	20068	64
Harddisk1	test	test_000	CL1-A	1	0	0	20068	64
Harddisk1	test	test_000	CL1-A	1	0	0	20068	64
Harddisk5	test	test_001	CL1-A	1	4	0	20068	68

.

.

.

```
Harddisk12 test test_011 CL1-A 1 11 0 20068 75
```

```
Harddisk1 test test_000 CL1-A 1 0 0 20068 64
```

HORCM Shutdown inst 5 !!!

Please check

'C:\HORCM\etc\horcm5.conf', 'C:\HORCM\etc\log5\curlog\horcm_*_log.txt',
and modify 'ip_address and service'.

```
C:\HORCM\etc>
```

The **mkconf** command creates the following file. Note that the loopback IP address is used, along with the default service number instead of a configuration file name.

Using the service number eliminates the need to edit the etc/services file. Some commented entries were removed for clarity:

Created by mkconf on Thu Oct 25 13:02:44

HORCM_MON

#ip_address	service	poll(10ms)	timeout(10ms)
127.0.0.1	52323	1000	3000

HORCM_CMD

#dev_name

#UnitID 0 (Serial# 20068)

\\.\PhysicalDrive13

HORCM_DEV

#dev_group	dev_name	port#	TargetID	LU#	MU#
# Harddisk1	SER = 20068	LDEV = 64	[FIBRE FCTBL = 2]		
test	test_000	CL1-A	1	0	0
# Harddisk5	SER = 20068	LDEV = 68	[FIBRE FCTBL = 2]		
test	test_001	CL1-A	1	4	0
# Harddisk2	SER = 20068	LDEV = 65	[FIBRE FCTBL = 2]		
test	test_002	CL1-A	1	1	0
# Harddisk3	SER = 20068	LDEV = 66	[FIBRE FCTBL = 2]		
test	test_003	CL1-A	1	2	0
# Harddisk4	SER = 20068	LDEV = 67	[FIBRE FCTBL = 2]		
test	test_004	CL1-A	1	3	0
# Harddisk6	SER = 20068	LDEV = 69	[FIBRE FCTBL = 2]		
test	test_005	CL1-A	1	5	0
# Harddisk7	SER = 20068	LDEV = 70	[FIBRE FCTBL = 2]		
test	test_006	CL1-A	1	6	0
# Harddisk8	SER = 20068	LDEV = 71	[FIBRE FCTBL = 2]		
test	test_007	CL1-A	1	7	0
# Harddisk9	SER = 20068	LDEV = 72	[FIBRE FCTBL = 2]		
test	test_008	CL1-A	1	8	0
# Harddisk10	SER = 20068	LDEV = 73	[FIBRE FCTBL = 2]		
test	test_009	CL1-A	1	9	0

```
# Harddisk11      SER = 20068 LDEV = 74 [ FIBRE FCTBL = 2 ]
test      test_010  CL1-A      1   10   0
# Harddisk12      SER = 20068 LDEV = 75 [ FIBRE FCTBL = 2 ]
test      test_011  CL1-A      1   11   0
# ERROR [LDEV LINK] Harddisk1      SER = 20068 LDEV = 64 [ FIBRE
FCTBL = 2 ]
HORCM_INST
#dev_group  ip_address  service
test      127.0.0.1    52323
```

Configuration: Protection Facility

- CCI recognizes only volumes permitted by the Protection Facility
- The Protection Facility ON/OFF is controlled by Storage Navigator or the SVP, or SNMP.
 - The Protection Facility uses an enhanced command device defined using the LUN Manager software (or SNMP).
 - When you define the command device, the Protection Facility is turned ON or OFF to each command device.
 - CCI recognizes the Protection attribute on startup.

Configuration

- Start Instances
 - Microsoft Windows
 - `horcmstart <instance number>`
 - UNIX
 - `horcmstart.sh <instance number>`
 - Commands
 - `pairstdisplay` – **Important! Confirm copy direction**
 - `paircreate`
 - `pairsplit`
 - `pairresync`
-
- Environment parameters can now be set in command line.
 - `-I[instance#]` – specifies Instance number
 - Instance 4 example: `pairstdisplay -g urgrp -I0`
 - `-IH[instance#]` or `-ITC[instance#]` – specifies the command as a Universal Replicator software or TrueCopy Remote Replication software operation,
 - Instance 4 example: `pairstdisplay -g urgrp -fxce -IH0`
 - `-IM[instance#]` or `-ISI[instance#]` – specifies the command as ShadowImage software command
 - Instance 0 example: `pairstdisplay -g siggrp -IM10`

Commands

- Universal Replicator enhanced CCI Commands
 - **pairedisplay** displays pair information* Enhanced for MxN Consistency Group
 - **paircreate** creates Universal Replicator pairs* Enhanced for MxN Consistency Group
 - **pairsplit** splits a Universal Replicator pair*
 - **pairresync** resynchronizes a Universal Replicator pair
 - **pairevtwait** waits for a event transition
 - **pairmon** monitors and reports on pair events
 - **raidscan** displays port configuration*
 - **raidqry** displays CCI host configuration
 - **raidar** displays port I/O activity
 - **raidvchkscan** checks volume information*
 - **pairsyncwait** waits for split pairs to synchronize*
 - **inqraid** displays relation between special files and PDEVs*
 - **paircurchk** checks the currency of remote volumes by evaluating the data consistency based on pair status and fence level*
 - **horctakeoff** converts 3DC Multi-remote to 3DC Cascade*
 - **horctakeover** reverses MCU/RCU relationship (failover)*

* UR Enhanced command shown in red. Other commands are the same as for TrueCopy software.

33

▪ **pairedisplay** – specific options for HUR

- fxce** – Displays LDEVs in hex, % synchronized, journal and external information
 - v jnl** – Displays journal information such as Q-MARKER (sequence number) and Q-COUNT (remaining sequence numbers)
 - v jnlt** – Displays journal information plus timeout value settings
- Equivalent to same parameters in **raidvchkscan** command, discussed later.

```
#pairedisplay -g urgrp -fxce -IH0
```

Group Seq#, LDEV P/S, Status, Fence, %, P-LDEV# M CTG JID AP EM E-Seq# E-LDEV#

CTG CT group ID for UR pair. (MxN Consistency Group if defined)

JID UR journal group ID

AP Number of active paths for UR links

EM External connection mode

E-Seq# Displays the production (serial) number of the external LUN, “Unknown” is shown as ‘-’.

E-LDEV# Displays the LDEV# of the external LUN.

34

CCI supports the **-fdxe** option with the **pairedisplay** command so that you can discover the external LUNs on the pair volume. This will show you additional information on the pair volumes.

▪ **pairedisplay** (within a shell script)

```
pairevtwait -g $datagroup -nowait | grep " PSUS" > /dev/null
if [ $? -ne 0 ]
then
    pairedisplay -g $datagroup -fxc -IH0
    exit 1
else
    echo "$W Pair '$datagroup' already Split"
fi
```


- **paircreate** — options specific to UR
 - jp <id> — M-JNL group ID
 - js <id> — R-JNL group ID
 - f **async** [CTGID] — Specify Fence Level **async** for **Universal Replicator**. CTGID Defaults to next sequential ID if not specified. See note for usage in MxN Consistency Group environments
- HUR Paircreate example


```
#paircreate -g urgrp -vl -f async 20 -jp 20 -js 20 -IH0
```

-g	CCI grp_name
-vl	Establishes copy direction Local to Remote (normal) direction
-f async 20	UR pair, CT ID 20
-jp 20	UR M-JNL ID
-js 20	UR R-JNL ID
-IH0	Instance 0 UR environment parameter

Note: If CCI configuration files are set up to use **MxN Consistency Group**, the CTGID specified in the **-fence** parameter will be the **MxN Consistency Group ID**. Otherwise, the CTGID will only apply to the mirrors in the specified journal.

- **pairsplit**
 - P Purge remaining Journal data to bitmaps without updating S-VOL.
 - R Force S-VOL to simplex mode forcibly. Issued at secondary site.
 - RS Set S-VOLs to SSWS status. Allows full host access.
 - RB Change S-VOL from SSWS to SSUS (normal split status).
 - I Split local volume only. Issued at primary site.

Best practice: Use HUR and ShadowImage Consistency Group AT-TIME Split together to provide a gold copy of remote data

To create Open ShadowImage consistency group, use **-m grp NN** option in **paircreate** command (NN = Consistency Group ID 00-7F)

Operational sequence:

Issue pairsplit to HUR group. Allow S-VOLs to go to SSUS status

Issue AT-Time split to ShadowImage Consistency group.

Note: ShadowImage Consistency Group can be split while HUR Journal is in PAIR status but does NOT ensure multi-volume consistency of ShadowImage S-VOLs. For best results, always split UR groups first.

■ **pairsplit** — examples

- # **pairsplit -g urgrp -IH6**
- # **pairsplit -g urgrp -d vol1 -IH6**
- # **pairsplit -g urgrp -rw -IH6**
- # **pairsplit -g urgrp -S -IH6**
- # **pairsplit -g urgrp -RS -IH7** (Executed at secondary site to force S-VOLs to SSWS status)

Example:

```
pairsplit -g urgrp -d vol1 -IH0  
pairedisplay -g urgrp -d vol1 -fxc -IH0
```

Group	PairVol(L/R)	(Port#,TID,LU),	Seq#,	LDEV#.	P/S,	Status,	Fence,	% ,	P-LDEV#
urgrp	vol1(L)	(CL1-A , 1, 0)	10145	0..	P-VOL	PSUS, ASYNC	98	200	
urgrp	vol1(R)	(CL1-A , 0, 0)	10156	200..	S-VOL	SSUS, ASYNC	100	0	

In the event of a temporary primary site outage, **pairsplit -RS** can be executed at the Local or Remote Site to force S-VOLs into SSWS status so that host applications can use the S-VOLs.

■ **pairsplit** (within a shell script)

```
for datagroup in $VGNames
do
  pairsplit -g $datagroup -IH0
  if [ $? -ne 0 ]
  then
    echo "Error Pairsplit command failed for '$datagroup' \n"
    exit 1
  else
    echo "Split command '$datagroup' successful"
  fi
  pairevtwait -g $datagroup -s psus -t 7200
  if [ $? -ne 0 ]
  then
    echo "Error Pairsplit not complete for '$datagroup' \n"
    exit 1
  else
    echo "Split '$datagroup' successful"
  fi
done
```

■ pairresync

– Examples

```
# pairresync -g urgrp -IH6
```

```
# pairresync -g urgrp -d oradb0 -IH6
```

```
# pairresync -g urgrp -swaps -IH7 (P-VOL/S-SVOL swap)
```

Example after doing **pairresync -g urgrp -IH6**

```
C:\HORCM\etc>pairdisplay -g urgrp -d oradb0 -fxc -IH6
```

Group	PairVol(L/R)	(Port#,TID,LU), Seq#, LDEV#.	P/S,Status,Fence, % ,P-LDEV#
urgrp	vol1(L)	(CL1-A , 1, 0) 10145 0..	P-VOL RSCOPY , ASYNC 98 204
urgrp	vol1(R)	(CL1-A , 0, 0) 10156 204..	S-VOL RSCOPY , ASYNC100 0

Note: **pairresync -swaps** is the second command of a failover operation at a remote site. It attempts to change S-VOLs in SSWS status (see **pairsplit -RS**) into P-VOLs. If successful, MCU/RCU relationship is reversed.

■ pairresync

– Parameters:

-swaps — Executed from the secondary site to change replication direction. Swaps P-VOL and S-VOL

-swapp — Executed from the primary site to set the primary volume back to P-VOL

-l — Sets local volume (must be P-VOL) back to PAIR status regardless of S-VOL status

▪ **pairresync** (within a shell script)

```
for datagroup in $VGNames
do
pairresync -g $datagroup -IH6
  if [ $? -ne 0 ]
  then
    echo "Error Pairresync command failed for '$datagroup' \n"
    exit 1
  else
    echo "Start of resynchronization of '$datagroup' successful"
  fi
  echo "Wait for end of resynchronization"
  $pairevtwait -g $datagroup -s pair -t 7200
  if [ $? -ne 0 ]
  then
    echo "Error Resync not complete \n"
    exit 1
  fi
  pairstatus -g $datagroup -fcx
done
```

■ **raidvchkscan** – Enhanced for HUR; checks journal information

Parameter	Details
Command Name	raidvchkscan: Validation checking confirmation command
Format	raidvchkscan { -h -q -z -v jnl [t] [unit#] [-s <Seq#>] [-f[x]] }
Options	-h: Displays Help/Usage and version information. -q: Terminates the interactive mode and exits the command. -z or -zx: Makes the raidvchkscan command enter the interactive mode. The -zx option guards performing of the HORCM in the interactive mode. When this option detects a HORCM shut down, interactive mode terminates. -l[H][M][instance#] or -l[TC][SI][instance#] Specifies the command as [HORC]/[HOMRCF], and used for specifying instance# of HORCM. -s <Seq#>: Used to specify the Seq# (serial#) of the storage system when this option cannot specify unitID which is contained for "-v jnl" option. If this option is specified, the unitID which is contained in "-v jnl" is invalid. -fx: Displays the LDEV number in hexadecimal notation. -v jnl: Displays information for the journal volume (see Figure 4.46). -v jnl:t: Displays the DOW, DPW, and APW time-out values for controlling the journal (see Figure 4.48).

▪ **raidvchkscan** parameters

-v jnl [unit#] finds the journal volume lists by Unit Number.

-v jnl -s [Serial #] finds the journal volume lists by Serial Number

Example: # **raidvchkscan -v jnl -s 62500**

JID	MU	CTG	JNLS	AP	U(%)	Q-Marker	Q-CNT	D-SZ(BLK)	Seq#	Nm	LDEV#
001	0	1	PJNN	4	21	43216fde	30	512345	62500	2	265
002	1	2	PJNF	4	95	3459fd43	52000	512345	62500	3	270
003	0	3	PJSN	4	0	-	-	512345	62500	1	275
004	0	4	PJSF	4	45	1234f432	78	512345	62500	1	276
005	0	5	PJSE	0	0	-	-	512345	62500	1	277

Q-Marker shows the latest sequence # on the P-JNL volume

Q-CNT: the number of remaining Q-Markers within each journal volume.

D-SZ: journal volume capacity

Num: The number of journal volumes in the group.

44

Details of **raidvchkscan** output table:

JID: Displays the journal group ID

MU: Displays the mirror descriptions on TrueCopy/Universal Replicator

CTG: Displays the CT group ID

JNLS: Displays the following status in the journal group

- SMPL → this means the journal volume which does not have a pair
- P(S)JNN → this means "P(S)-VOL Journal Normal Normal"
- P(S)JSN → this means "P(S)-VOL Journal Suspend Normal"
- PJNF → this means "P(S)-VOL Journal Normal Full"
- P(S)JSF → this means "P(S)-VOL Journal Suspend Full"
- P(S)JSE → this means "P(S)-VOL Journal Suspend Error" including Link failure

AP: Displays the active paths

U(%): Displays the usage rate of the journal data

Q-Marker: Displays the sequence # of the journal group ID, called the Q-marker

For PJNL, Q-Marker shows the latest sequence # on the P-JNL volume

For S-JNL, the Q-marker shows the latest sequence # of the cache (DFW).

Q-CNT: Displays the number of remaining Q-Markers within each journal volume

D-SZ: Displays the capacity for the journal data on the journal volume

Seq#: Displays the serial number of the RAID

Num: Displays the number of LDEV configured the journal volume

LDEV#: Displays the first number of LDEV configured the journal volume

▪ **pairevtwait**

- Waits until a specific pair status is achieved before returning control
- Useful for scripts that need to wait until a specific pair status is achieved

pairevtwait -g <group> -s <status> -t <timeout>

The **pairresync** command terminates before resynchronization of the remote (or primary) volume is complete. Use pair event waiting to verify that the resync operation completed successfully (status changes from COPY to PAIR).

▪ **pairmon**

- Obtains the pair status transition of each volume pair and reports it
- Runs in background
- Issues messages on pair status changes

-allsnd

Reports all events if there is pair status transition information.

```
# pairmon -allsnd -nowait
```

```
Group Pair vol Port targ# LUN# LDEV# Oldstat code > Newstat code
```

```
oradb oradb1 CL1-A 1 5 145 SMPL 0x00 > COPY 0x01
```

```
oradb oradb2 CL1-A 1 6 146 PAIR 0x02 > PSUS 0x04
```


■ raidqry

Displays the configuration of the connected host and storage system.

raidqry -l

No	Group	Hostname	HORCM_ver	Uid	Serial#	Micro_ver	Cache (MB)
1	---	HOSTA	01-22-03/06	0	30053	50-04-00/00	8192
1	---	HOSTA	01-22-03/06	1	30054	50-04-00/00	8192

47

■ raidar

Displays port statistics

raidar -p cl1-a 15 6 -p cl1-b 14 5 -p cl1-a 12 3 -s 3

TIME[03]	PORT	T	L	VOL	STATUS	IOPS	HIT(%)	W(%)	IOCNT
----------	------	---	---	-----	--------	------	--------	------	-------

13:45:25	-	-	-	-	-	-	-	-	-
----------	---	---	---	---	---	---	---	---	---

13:45:28	CL1-A 15 6	SMPL	---			200.0	80.0	40.0	600
----------	------------	------	-----	--	--	-------	------	------	-----

	CL1-B 14 5	P-VOL PAIR				133.3	35.0	13.4	400
--	------------	------------	--	--	--	-------	------	------	-----

	CL1-A 12 3	P-VOL PSUS				200.0	35.0	40.6	600
--	------------	------------	--	--	--	-------	------	------	-----

IOPS: # of I/Os (read/write) per second (total I/O rate)

HIT(%): Hit rate for read I/Os (read hit rate).

W(%): Ratio of write I/Os to total I/Os (percent writes).

IOCNT: number of write and read I/Os

48

▪ **pairsyncwait**

- fq** This option is used to display the number of remaining Q-Markers within CT group.
- m** [Q-marker] can confirm the transfer of a specific sequence number to remote site. In other words, determines whether or not pairs are synchronized while in PAIR STATUS.

pairsyncwait -g oradb -nowait -fq -IH6

UnitID	CTGID	Q-Marker	Status	Q-Num	QM-Cnt
0	3	01003408ef	NOWAIT	2	120

pairsyncwait -g oradb -nowait -m 01003408e0 -fq -IH7

UnitID	CTGID	Q-Marker	Status	Q-Num	QM-Cnt
0	3	01003408e0	NOWAIT	2	105

pairsyncwait -g oradb -t 50 -fq -IH6

UnitID	CTGID	Q-Marker	Status	Q-Num	QM-Cnt
0	3	01003408ef	TIMEOUT	2	5

49

Details of **pairsyncwait** output table:

UnitID: The Unit ID, in case of multiple DKC connection

CTGID: The CTGID within the UnitID

Q-Marker: The latest sequence # of the MCU P-VOL when the command was received

Status: The status after the execution of the command

Q-Num: The number of processes queued to wait for synchronization within the CTGID of the Unit

QM-Cnt: The number of remaining Q-Marker within CT group of the Unit
TrueCopy/Async sends a token called "dummy record set" at regular interval time.
Therefore **QM-Cnt** always shows "2" or "3" even if Host has no writing.

- When specifying "**-nowait -fq**" the "**QM-Cnt**" will be shown as the number of remaining Q-Marker at this time within CT group.
- When specifying "**-nowait -m <marker> -fq**" the "**QM-Cnt**" will be shown as the number of remaining Q-Marker from the specified <marker> within CT group.
- When specifying "TIMEOUT" without "**-nowait**" the "**QM-Cnt**" will be shown as the number of remaining Q-Marker at this timeout within CT group. "**QM-Cnt**" will be shown as "-", if the status for Q-Marker is invalid (i.e. status is "BROKEN" or "CHANGED").

■ inqraid -fw

Displays cascading volume status (3DC)

Example:

```
- # ls /dev/rdsd/* | inqraid -CLI -fw -IH0
- DEVICE_FILE PORT SERIAL LDEV CTG C../B../ SSID R:Group PRODUCT_ID
- c1t2d10s2 CL2-D 62500 266 - Psss/P/PP-----0005 1:01-02 OPEN-3 -SUN
- c1t2d11s2 CL2-D 62500 267 - ssss/s/ss----- 0005 1:01-02 OPEN-3 -SUN
```

Note: P = P-VOL, S=S-VOL, s=simplex

This option is added to display all of the cascading volume statuses on STD Inquiry Page.

■ paircurchk (Pair Currency Check)

Object Volume			Currency	
Attribute	Status	Fence	Paircurchk	SVOL_takeover
SMPL	-	-	To be confirmed	-
P-VOL	-	-	To be confirmed	-
S-VOL	COPY	Data	Inconsistent	Inconsistent
		Status		
		Never		
	PAIR	Async	Inconsistent	Inconsistent
		Data	OK	OK
		Status	OK	OK
	PAIR	Never	To be analyzed	To be analyzed
		Async	To be analyzed	OK (assumption)
		PFUL	To be analyzed	OK (assumption)
	PSUS	Data	Suspected	Suspected
		Status	Suspected	Suspected
		Never	Suspected	Suspected
	PSUS	Async	Suspected	Suspected
		PFUS	Suspected	OK (assumption)
		PSUE	Data	OK
PDUB	Status	Status	Suspected	Suspected
		Never	Suspected	Suspected
		Async	Suspected	OK (assumption)
	SSWS	Data	Suspected	-
		Status	Suspected	
		Never	Suspected	
		Async	Suspected	

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

51

The CCI **paircurchk** command checks the currency of remote volumes by evaluating the data consistency based on pair status and fence level.

Scripted Commands for Disaster Recovery

■ horctakeover

- Useful for temporary loss of primary or remote site
- Execution:
 - Checks the specified volume or group attributes (paircurchk)
 - Decides which takeover function to implement based on the attributes
 - Takeover Switch executes the chosen takeover function
 - Temporary loss of Remote Site – P-VOL-takeover
 - Both sites operational – Swap-takeover
 - Temporary loss of primary site – S-VOL-takeover

See Data Protection Module for detailed description.

■ **horctakeoff**

- Change 3DC multi-target to 3DC cascade while host applications are running.
- Conditions: Temporary primary site failure, failover to Local site has occurred.
- Execute **horctakeoff**.
- If successful, execute **horctakeover** to switch operations to the Remote site without affecting host operations at the Local site.
- Operates on either individual volume or volume group.

Note: See Data Protection Module for detailed description.

53

After horctakeoff execution, **horctakeover** command will failover replication from the Local Site to the Remote Site without stopping the host application.

Microsoft Windows Subcommands

- Provides UNIX-like functionality for Microsoft Windows hosts
- Can be executed by appending
 -x <subcommand> to any CCI command
- Recommendation: Use a non-destructive command (such as **pairdisplay**) when using these subcommands
 - findcmddev
 - drivescan
 - portscan
 - sync
 - syncd
 - mount
 - umount
 - umountd

- **sync** and **syncd**
 - Flushes the system cache to disk.
 - **syncd** flushes the system cache to disk and waits (30 sec) for delayed (paging) IO then dismounts the drive.
 - **sync** and **syncd** do not propagate to volume mount points under a specified drive letter.

pairdisplay -x sync all

pairdisplay -x sync D:

pairdisplay -x sync D:\mountpoint

pairdisplay -x syncd hdisk3 hdisk4

pairdisplay -x syncd \Vol2

▪ Mount

- Mounts the specified volume to a drive letter or volume mount point

pairdisplay -x mount

pairdisplay -x mount D: hdisk3 p1

pairdisplay -x mount E: \Vol2

pairdisplay -x mount F:\mountpoint \Vol3

56

▪ Umount

- Unmount volume (deletes drive letter or volume mount point mapping)
- Will flush the system cache to disk prior to unmounting

pairdisplay -x umount D:

pairdisplay -x umount \Vol2

pairdisplay -x umount F:\mountpoint

57

- Special Facilities for Windows 2008/2003/2000 Systems
 - Signature changing facility (see CCI section 4.20.1)
 - Directory mount facility (section 4.20.2)
 - CCI supports saving/restoring the GUID Diskid of the GPT Basic disk to the **inqraid** command.

58

- Special Facilities for Windows 2008/2003/2000 Systems
(*continued*)
 - LDM (Logical Disk Manager) Volume Discovery
 - Volume Discovery Function
 - **Physical level** – Use **\$Physical** as *Key Word* for the discovery.
 - **LDM volume level** – Use **\$Volume** as *Key Word* for the discovery.
 - **Drive letter level** – Use **\$LETALL** as *Key Word* for the discovery.
 - System Buffer Flushing Function
 - System buffers associated with logical drives can be flushed.

59

Configuration Setting Commands

- RAIDCOM Commands — CCI configuration setting commands:
 - Enable non-replication operations to be executed by CCI
 - Supported operations:
 - Creation, modifying, and deletion of Dynamic Provisioning and Copy-on-Write Pools and virtual volumes
 - Setting remote replication link attribute, creation of RCU and R-DKC
 - Creation, modification, and deletion of UR Journals
 - Creation, modification, and deletion of external storage paths and external volumes
 - For overview of CCI Configuration Setting operations, see *Provisioning Operations with CCI* section in *Hitachi Command Control Interface User and Reference Guide*, **MK-90RD7010**
 - For details on individual commands, see *Configuration Setting Commands* section in *Hitachi Command Control Interface Command Reference*, **MK-90RD7009**

Troubleshooting

- Instance will not start:
 - Check CCI logs:
 - **curlog** for useful information
 - Check that CMD device path is correct in the configuration file.
 - Check that CMD device is labeled; if not label it.
 - Check that the configuration file is saved in the correct OS directory and that file attribute is .conf.
 - Check environment variables.
 - Check configuration files for hidden characters.

- If **paircreate** fails, check to see if:
 - Correct devices are specified in the configuration file:
 - Use **raidscan**
 - Remember that LUNs are displayed in hex in LUN Manager so LUN 10 is actually LUN 16 in the configuration file.
 - Remember that there may be absolute LUNs to consider.
 - Run **raidscan** to find the absolute LUN value if using an older version of CCI.
 - S-VOLs are not reserved through Storage Navigator ShadowImage function.
 - Links are in place and working properly.

- CCI Error Codes

- Check /HORCM/log*/curlog/horcmlog_HOST/horcm.log where:

- * is the instance number.
 - HOST is the host name.

- Example: 11:06:03-37897-10413- SSB = 0x**SSB1,SSB2**

- Error codes appear on the right of the equal symbol (=)

- Error codes are specific to each replication and each storage platform.

- Consult the *Hitachi Universal Replicator User Guide* for your particular model of Hitachi enterprise storage.

- For error codes not listed, contact the Hitachi Data Systems Support Center.

Module Review

1. What Storage Navigator operations are required before RAID Manager CCI can be used to manage Universal Replicator pairs?
2. True or False: The Command Device must be mapped to a port so that user applications can access it.
3. If using the new HORCM_LDEV parameter, do you need to code the absolute LUN numbers?

7. Data Protection

Module Objectives

- Upon completion of this module, learner should be able to:
 - Identify current Data Protection bundle specifications
 - Evaluate Data Protection planning considerations
 - Rolling Disaster
 - Data Consistency
 - RPO and RTO
 - Remote Copy Solutions
 - Resynchronization Vulnerability
 - Perform failover operations with Hitachi CCI software
 - Identify best practices for Data Protection

2

Data Protection Bundles


- Data Protection bundle includes these software products:
 - Hitachi TrueCopy Synchronous software
 - Hitachi Universal Replicator (for point to point asynchronous remote copy operation)
- Data Protection Extended bundle includes the following features:
 - 3Data Center cascade and multi-target
 - Universal Replicator MxN Consistency Groups

Data Protection Planning Considerations

Data Protection Classifications – Two categories of disaster against which a business may protect its data

- Disasters of **scale** include:
 - **Point disaster** representing a single event at a single point of time. This could be caused by human error, or isolated hardware failure.
 - **Site-wide disaster** impacting operations at an entire facility, caused by fire, earthquake or hurricane.
- Disasters of **time** include:
 - **Immediate disaster** where a single, distinct event impacts all components at the same time, such as a tsunami or explosion.
 - **Rolling disaster** where several components fail at different points in time, for example a fire takes down a server, then storage, network, eventually the entire site. Typically this is the most severe type of disaster.
- Regulatory requirements driving data recovery
 - **Sarbanes-Oxley** – Public Company Accounting Reform and Investor Protection Act of 2002, Title III Corporate Responsibility, Section 404 Assessment of Internal Control
 - Commonly interpreted to require specific information technology controls on data archiving and retention
 - **Basel II, III** – International banking accord
 - Defines operational risk guidelines for international banking
 - **Email archiving** – Litigation defense under business records laws

Weigh cost of recovery against recovery time, cost and value of the data



6

■ Data Protection Tiers

Technology Tier	RPO Range	RTO Range	Minimum # Disk Copies	Distance
Tape Backup ¹	24-168 hours	2-168 hours	N/A	Any
Virtual Tape ¹	12-48 hours	1-24 hours	Disk Pool	Any if replicated
Disk Point-in-Time ¹	2-36 hours	15 minutes- 12 hours	3 ²	Any
Synchronous Remote Copy	0-2 minutes	1-8 hours	2 ²	Limited
Synchronous Remote Copy w/Failover	0-2 minutes	5-60 minutes	2 ²	Limited
Asynchronous Remote Copy	0-5 minutes ³	30 minutes- 8 hours	2 ²	Any
Asynchronous Remote Copy w/Failover	0-5 minutes ³	30-90 minutes	2 ²	Any
Three Data Center	0-2 minutes	1-8 hours	3-7 ^{2,4}	Any


7

Planning Considerations: Rolling Disaster

■ Rolling Disaster: Worst Case

A rolling disaster occurs when an unplanned outage event takes place over a span of time — anywhere from a few minutes to several hours. For example:

- A fire starts in the network area.
- This causes a link failure while remote backups are running.
- A few moments later fire reaches the other converter cabinet and the other links fail.
- Fire spreads to the data center floor and storage systems are affected.



8


During a rolling disaster, not all systems, storage and network connections fail at precisely the same moment. In this situation, a system may still be able to process transactions and issue updates to primary storage devices, but due to earlier failures, updates may not replicate successfully to the secondary site. Rolling disasters pose a challenge because they may result in corrupted and unusable data at the remote site, requiring difficult and very lengthy recovery processes.

To protect against rolling disasters, a data replication technology must be able to freeze remote replicas at a point in time prior to or during the onset of the outage. This ability to create point-in-time images of data is what differentiates remote copy technology from simple mirroring.

Because the remote and local I/O of a synchronous replication succeed or fail together, this replication approach does not introduce data inconsistencies following a disaster. Rolling disasters are primarily a challenge for remote asynchronous replication, and one of the principle areas of concern is write order fidelity.


- Rolling Disaster: Worst Case (*continued*)

- If there is no control over dependent writes to the backup volumes, out of sequence data is committed to the remote volumes.



- Data received at the remote disk control unit (RCU) will be incomplete, with dependent writes missing.

- Mitigating effects of Rolling Disaster – **ensure write order fidelity**
 - Replication technology must employ techniques to automatically preserve **write order fidelity** at the remote site.
 - Ensures the integrity of asynchronously replicated data in a rolling disaster.



10

In the context of data replication, data consistency represents the ability to recover from a failure or disruptive event. A fundamental concept of data consistency that enables quick recovery is the *dependent write*, the pervasive logic among complex data structures comprising databases, file systems, etc., that determines the sequence in which writes are issued. A dependent write is a data update that cannot be executed until a previous write, on which it is dependent, has been executed. This logic preserves the integrity, the consistency, of the data and allows systems and applications to restart after a sudden failure.

Planning Considerations: Data Consistency

■ Data Consistency

- Data consistency represents the ability to recover from a failure or disruptive event by preserving dependent writes.
 - A *dependent write* is a data update that cannot be executed until a previous write, on which it is dependent has been executed, thus preserving write order fidelity.
- There are three levels of data consistency with different implications within the application and data architecture.

■ Data Consistency Levels

- **I/O consistency**, or crash recovery consistency, refers to data that is not necessarily transaction consistent, but is still in a restartable state.
 - Provided by Hitachi remote replication products and other vendor's replications.
- **Transaction consistency** is achieved when an application is shut down (quiesced), or when the application/database or other system component rolls back or rolls forward after a restart. A transaction is a logical unit of work that may include hundreds or thousands of updates.
 - Local snapshot mirroring can assist by providing generational rollback data
- **Application consistency** is the state in which individual components have each been recovered to a transaction consistent state. Collectively, the components need to be synchronized based on the application requirements. This can thought of as “user consistency.”
 - Provided by application and operating system tools

Planning Considerations: RPO and RTO

- Recovery Point Objectives (RPO) and Recovery Time Objectives (RTO)
 - Evaluation of risk tolerance is stated in terms of how much data must be recovered to resume operations, the RPO, and the outage duration, the RTO.
 - **RPO:** The worst case time between the interruption in operations and the last recoverable backup, where potentially lost data weighs against cost
 - A measure of the amount of data that has to be recovered
 - **RTO:** The time to resume operations after the interruption



13

When evaluating the cost of business continuity solutions, the greatest cost component is usually the bandwidth needed to support remote replication. The greatest benefit is maintaining a minimal RPO.

- To maintain a continuous replica copy, bandwidth must exceed the average write workload that occurs during any given RPO interval subject to the capacity limitations of the buffering mechanism.
 - Establish desired RPO.
 - Identify the matching interval with greatest write activity.
 - Calculate link bandwidth and buffer capacity required to keep up with this traffic.
 - In practice, data on hand can be used to revise resource requirements.

Planning Considerations: Compare Remote Copy Solutions

- Universal Replicator: Hardware Asynchronous Remote Copy
 - Benefits:
 - Journals provide buffers to increase Universal Replicator resistance to link problems
 - RPO less dependent on link bandwidth than TrueCopy Async (See note)
 - Maintains consistency of remote data with consistency grouping, sequence numbering scheme
 - No CPU involved; CU does copy
 - Consistent host I/O delays averaging 0.6 millisec (not distance dependent)
 - Restrictions:
 - RPO 1 to 2 min due to loss of Update Copy, RTO 0 to 8 hours (0 for high availability environment with takeover function)
 - Journal volume throughput becomes limiting factor (See Note)
 - Exclude temp volumes
 - Available for Universal Storage Platform family only

15

Note: When evaluating the cost of business continuity solutions, the greatest cost component is usually the bandwidth needed to support remote replication. Universal Replicator reduces the cost outlay for link bandwidth. Some of that cost is transferred to the one-time cost of providing maximum possible journal volume throughput. Link bandwidth is still important, because unsent data on the MCU's Journal Volume represents an increase in RPO.

- TrueCopy Synchronous: Hardware Synchronous Remote Copy
 - Benefits: Low RPO, low RTO (See note 1)
 - No CPU involvement in remote copy
 - Writes go to remote cache:
 - Automatically maintains single-volume data consistency
 - Restrictions:
 - No inherent grouping function is available to ease management and provide consistency of secondary volumes.
 - With optical link latency, I/O delay is 0.9 millisec at zero distance.
 - Turnaround times become unacceptable at longer distances.
 - Exclude Temp data from replication.

Note 1: While RPO and RTO are not zero due to controller latency, they are as close as is possible to attain in a remote replication.

- Synchronous versus Asynchronous Replication
 - TrueCopy Synchronous software
 - Pros:
 - Very fast recovery time (low RTO)
 - Good data currency (low RPO)
 - Cons: Performance impact to the host; varies with distance; no grouping
 - Universal Replicator software
 - Pros:
 - Less performance impact than Synchronous
 - Protection against regional disasters
 - Cons: Less favorable RPO and RTO
 - Combine Synchronous and Asynchronous replications with three data center
 - Fast recovery and excellent data currency for local site failures
 - Advanced protection from regional disasters

Planning Considerations: Resynchronization Vulnerability

- **Minimize Resync Time** – Plan your operations to minimize pair resync time.
 - Remote volume data is vulnerable during resync, so minimize split time.
- Take advantage of new **ShadowImage Consistency Group At-Time Split** function to provide a “gold copy” of replicated data at remote site.
 - This will guard against loss of data if problem occurs during resync of remote pairs.

Planning Considerations: Recovery

- File Recovery:
 - When the primary or secondary storage system suspends a remote pair due to a disaster, the secondary data volume may contain in-process data.
 - File recovery procedures will be needed.
 - These procedures should be the same as those used for recovering data volumes that becomes inaccessible due to control unit failure.
- Lost Updates:
 - No remote replication provides any function for retrieving lost updates.
 - Requires application features such as database log file.
 - Check consistency of S-VOL data, then switch operations to the secondary site.
 - Verify currency of files used for file and database recovery.

Failover: Planning Considerations

- Switching operations to the secondary site can be done manually or automated by using the CCI **horctakeover** scripted command on remote replication pairs.
 - This command performs the necessary step to prepare the secondary site to function as if it were a primary site, using secondary data volumes.
- The **horctakeover** command checks the pair status of secondary data volumes and splits consistency groups to ensure consistency of secondary data volumes and makes the secondary data volumes usable.
- If possible, the **horctakeover** command attempts to restore pairs to reverse primary data volumes and secondary data volumes (change S-VOLs to P-VOLs).
- If the **horctakeover** command executes successfully, operations can be started at the secondary site, using secondary data volumes that are now P-VOLs.
 - For detailed information about how to use CCI, please refer to the Command Control Interface (CCI) user and reference guides.

20

Failover: horctakeover Command


■ horctakeover Command

- Executes several operations
 - Checks the specified volume's or group's attributes (**paircurchk**)
 - Decides the takeover function based on the attributes
 - Executes the chosen takeover function, and returns the result
- Takeover-switch process determines which of the following functions to execute:
 - **Swap-takeover**
 - **PVOL-takeover**
 - **SVOL-takeover**

21

The TrueCopy Synchronous/Asynchronous software **takeover** command (**horctakeover**) is a scripted command for executing several TrueCopy Synchronous/Asynchronous software operations. The takeover command checks the specified volume's or group's attributes (**paircurchk**), decides the takeover function based on the attributes, executes the chosen takeover function, and returns the result. A paired volume or a group can be specified as the target of the TrueCopy Synchronous/Asynchronous software takeover command. If S-VOL-takeover is specified for a group, the data consistency check is executed for all volumes in the group, and all inconsistent volumes are found in the execution log file and displayed (same as **paircurchk** command). The takeover command allows swapping of the primary and secondary volumes, so that if the primary or secondary volume is switched due to a server error or package transfer, duplex operations can be continued using the reversed volumes. When control is handed over to the current node, swapping the volumes again eliminates the need to copy them. The takeover command also allows the secondary volume to be separated for disaster recovery operations.

■ horctakeover Flowchart



22

- **PVOL Takeover** – This function allows the takeover node to use the primary volume (i.e., reading and writing are enabled), on the assumption that the remote node (possessing the secondary volume) cannot be used.
 - This function executes the following two commands:
 - **PVOL-PSUE takeover:**
 - Changes the primary volume to the suspend (PSUE , PSUS) state which enables write I/Os to all primary volumes of the group.
 - The action of the PVOL-PSUE- takeover causes PSUE and/or PSUS to be intermingled in the group.
 - **PVOL-SMPL takeover:**
 - Changes the primary volume to the simplex (SMPL) state. First, PVOL-takeover executes PVOL-PSUE takeover further than PVOL-SMPL takeover. If the PVOL-PSUE takeover function fails, the PVOL-SMPL takeover function is executed.

23

The PVOL-takeover function releases the pair state as a group, since that maintains the consistency of the secondary volume at having accepted horctakeover command when the primary volume is fenced (*data* or *status* and *PSUE* or *PDUB* state, *PSUE* or *PDUB* volume are contained in the group). This function allows the takeover node to use the primary volume (for example, reading and writing are enabled), on the assumption that the remote node (possessing the secondary volume) cannot be used. PVOL takeover can be specified for a paired volume or a group.

For Asynchronous software: PVOL-takeover will not be executed.

▪ **swaptakeover**

- The **swap-takeover** function is used by the HA control script when a package is manually moved to an alternate data center while all hardware is operational.
- This function internally executes the following commands to swap the primary and secondary volumes:
 - **pairsplit -RS** forces the secondary volume into SSWS mode.
 - **pairresync -swaps** swaps the secondary volume with the primary volume.
- If successful:
 - S-VOL becomes P-VOL (MCU – RCU relationship is reversed).
 - New S-VOL is resynchronized.

When the P-VOL status of the remote node is PAIR and the S-VOL data is consistent, it is possible to swap the primary and secondary volumes. The swaptakeover function is used by the HA control script when a package is manually moved to an alternate data center while all hardware is operational. Swaptakeover can be specified for a paired volume or a group.

The swaptakeover function internally executes the following commands to swap the primary and secondary volumes:

1. Execute **suspend** for swapping the local volume (S-VOL). If this step fails, swaptakeover is disabled and an error is returned.
2. Execute **resync** for swapping to the S-VOL (local volume). This will swap P/S-VOL and redirect the copy direction. The result of this is:
 - The old P-VOL will be the new S-VOL
 - The old S-VOL will be the new P-VOL

And of course, this will also change the remote copy direction and it will synchronize the pair. To move back to the original state you need to repeat this procedure (where, of course, P/S-VOL are the new current ones). If this step fails, swap-takeover returns at SVOL-SSUS-takeover, and the local volume (SVOL) is maintained in SSUS(PSUS) state which allows and keeps track of write I/Os using a bitmap for the SVOL. This special state is displayed as SSWS using the **-fc** option of the **pairedisplay** command.

The swaptakeover function does not use SMPL or No Copy mode for swapping to guarantee mirror consistence, and this is included as a function of SVOL takeover.

■ S-VOL Takeover

- The **SVOL takeover** function allows the takeover node to use the secondary volume (except in COPY state) on the assumption that the remote node (possessing the primary volume) cannot be used.
 - **pairsplit -RS** forces the secondary volume into SSWS mode.
- If the primary and secondary volumes are not consistent, the **SVOL takeover** function fails.
- If primary and secondary volumes are consistent, the **SVOL takeover** function attempts to switch to the primary volume using resync for swapping.
- **SVOL takeover** can be specified for a paired volume or a group.
 - If the **SVOL takeover** is specified for a group, a data consistency check is executed for all volumes in the group, and all inconsistent volumes are displayed.

25

The data consistency of the UR S-VOL is evaluated by its pair status and fence level. If successful, the SVOL takeover function returns swap takeover as the return value of the horctakeover command. If not successful, the SVOL takeover function returns SVOL-SSUS-takeover as the return value of the **horctakeover** command. In case of a host failure, Swap takeover is returned. In case of an ESCON/FC or P-VOL site failure, SVOL-SSUS-takeover is returned.


Failover Using CCI

- Suspend Swap (***pairsplit -RS*** at remote Site)
 - Checks the pair status of secondary data volumes and splits journal groups to ensure consistency of secondary data volumes.
 - Makes the remote data volumes usable (SWAPPING status).
 - Remote volumes can now be used without restriction.
 - Allows Data Protection testing that requires updating on the secondary site in the Suspend status.

- Resync Swap (***pairresync -swaps*** at remote site)
 - Reverses primary data volumes and secondary data volumes if primary site is still usable.
 - Attempts to do a resync of the new secondary data volume (now in the former primary site).
 - If successful, reverses the MCU/RCU relationship and prepares the secondary site to function as if it were a primary site.
 - Operations can be started at the secondary site, using secondary data volumes that are now primary volumes.

Best Practices


- Situation — MCU consistency group has P-VOL = PAIR or COPY after pairsplit -RS.
 1. `pairsplit -RS` was executed while communication is unavailable (network failure, MCU power down), or...
 2. `pairsplit -RS` was executed while all pairs in RCU are suspended



- Fix: Reverse Copy Direction
 - After MCU and network recovers, execute resync swap (`pairresync -swaps`).
 - Pair status in RCU consistency group does not change, but pair status in MCU consistency group changes to suspend.
 - Execute resync swap again to change MCU volumes back to P-VOLs.

28

- Situation: Some pairs in the consistency group are reversed copy direction.
 - Volume based resync command has been executed.



- Fix:
 - Issue resync swap command (`pairresync -swaps`) to the individual P-VOL= PSUS or PSUE) and S-VOL= SSWS pair.


■ Situation: P-VOL= COPY or PAIR, S-VOL= suspend

– Reasons:

- RCU power down
- RCU internal error w/o delete pair

– Fix:

- Restart by issuing pairsplit then resync
- **pairsplit -S** then **paircreate** may be necessary




30

■ Situation: P-VOL= PSUE, S-VOL= COPY

– Possible Reasons:

- Failure during **paircreate**

– Fix: Restart by issuing **pairsplit -S**, then issue **paircreate** again



31

Best Practices: Monitoring

- Remote Copy Monitoring
 - **TrueCopy software** – Monitor pair status, consistency group status
 - **Universal Replicator** – Universal Replicator resists extended link failure by storing data in journal volumes. Monitor pair status, journal group status
 - M-JNLG status ACTIVE: Read JNL request arrived
 - R-JNLG Status ACTIVE: Read JNL response (Journal Copy) arrived

- **Remote Copy Monitoring** (*continued*)
 - TrueCopy and Universal Replicator
 - At remote site, check secondary volume consistency.
 - Consistency is good if:
 - All pairs were in **PAIR** status at the time of the disaster.
 - There is no volume based pair operation.

Best Practices: TrueCopy Sync Consistency Group


- Use CCI TrueCopy Synchronous Consistency Preservation Function.
- Use **-fg** parameter with Consistency Group ID when defining the TrueCopy Sync group.
 - Provides multi-volume write order consistency for TrueCopy Synchronous volumes (see note)
 - Create CCI script using these commands to control TrueCopy Sync Consistency Group
- Use UR MxN Consistency Groups to ensure consistency of data across multiple journal groups and multiple storage systems.

34

Note: For TrueCopy Synchronous copy groups with no C/T group ID specified, the update order is only preserved by copy pairs on the hardware. When Fence Level is Never, if a failure occurs on a copy destination volume within the copy group, group consistency might be lost because copy pairs other than the failed copy pair within the same copy group continue processing. When applied to a planned outage, the suspend operation takes place for each copy pair. Consequently, consistency within the copy group might be broken.

Best Practices: Data Protection Configuration

- Generic Recovery Configuration – Remote replication combined with ShadowImage replication
 - ShadowImage at Primary site provides local backup if desired.
 - ShadowImage at remote site provides remote backup copies for rapid recovery, split-mirror backup, decision support, and testing and development.
 - Implementation provides a remote disaster restart and recovery capability, as well as covering remote copy resynchronization vulnerability.



35

In the diagram, **UR** stands for Universal Replicator.

Best Practices: ShadowImage At-Time Split

- Use ShadowImage Consistency Groups AT-TIME split function to provide consistent ShadowImage copy of UR data
 - Backup software
 - Decision Support systems, others
- Match UR primary and/or secondary groups to ShadowImage consistency groups
 - The ShadowImage At-Time Split function applies to ShadowImage pairs that belong to a consistency group.
 - Consistency Group: User-defined set of ShadowImage volume pairs that contain multi-volume data to be operated on with a single command.
 - At-Time Split function: The CCI operation that splits the ShadowImage Consistency Group.

Best Practices Review

- Universal Replicator Remote Copy Monitoring
 - Helps reduce effects of disasters.
 - Provides information about Read JNL to indicate continuous remote copy.
 - M-JNLG: Read JNL request arrived or not
 - R-JNLG: Read JNL response (Journal Copy) arrived or not
 - Does not have heart-beat timeout.
 - Resists extended link failure by storing data in JNL VOL
 - At remote site, check S-VOL consistency.
 - Consistency of S-VOLs is good if:
 - All pairs were **in PAIR status** at the time of the disaster
 - There is no volume based pair operation

- | | | |
|---|---|--|
| <ul style="list-style-type: none">■ CCI on both primary and secondary sites is recommended.■ Monitor remote copy after PAIR status.■ Restart remote copy after suspend. | } | Normal Operations |
| <ul style="list-style-type: none">■ If failover is required:<ul style="list-style-type: none">– At remote site, check consistency of S-VOL– At remote site, issue CCI horctakeover. Restart application on remote storage array | } | Data Protection
Operations at Remote
Site |

Module Review

1. What are the two CCI commands used sequentially to invoke secondary site takeover?
2. What should be the status condition of the secondary volume following the issuance of the **pairsplit -RS**?
3. What is the expected result of the two commands?

8. Three Data Center Operations

Module Objectives

- Upon completion of this module, the learner should be able to:
 - Describe purpose of three data center (3DC) replication
 - Identify the different 3DC configurations, including new two data center (2DC) Pass-thru
 - Describe new 2DC Universal Replicator and TrueCopy Synchronous Configuration
 - Discuss 3DC Operations
 - Describe disaster recovery operations in 3DC environment

2

Purpose of 3DC Replication

- Purpose of 3DC Configuration
 - Provide redundancy of production data
 - Maintain multiple mirrors of the same data at different remote sites
 - Combines two Hitachi replication products
 - TrueCopy Synchronous – low RPO, RTO
 - Universal Replicator – asynchronous replication at any distance
 - Protection in region-wide disasters
 - Low application performance impact
 - ShadowImage can be used to mirror replication data at any site.

Configurations Summary

- Two Data Center Cascade Pass-Thru
 - Reduces cost by allowing Dynamic Provisioning TC/UR data volumes at the local site. No host required if remote command device is used.
- Three Data Center Cascade
 - Replicates data from the primary site to an local site via TrueCopy Synchronous software and then to a third remote location with Universal Replicator software
- Three Data Center Multi-target
 - Simultaneously copies data from a central location to an local site via TrueCopy Synchronous software and to a third remote site via Universal Replicator
- Three Data Center Multi-target with Delta Resync
 - Supports recovery of the remote site from the synchronous copies of journal data at the local site if the primary site has failed
- Four Data Center
 - 3DC multi-target and 3DC cascade combined

4

Configurations: 2DC Cascade PassThru

■ Required Components

- **Primary site – Any VSP, USP V, or USP** model with requisite licenses
 - TrueCopy Synchronous P-VOLs
- **Intermediate site – VSP, USP V, or USP VM** only with Disaster Recovery Extended Bundle – System Option Mode 707 ON
 - TrueCopy Sync writes data to cache normally.
 - Universal Replicator Journal Obtain process creates journal update
 - S-VOLs/UR P-VOLs are Dynamic Provisioning V-VOLs
- **Remote site – Any VSP, USP V, or USP** model with requisite licenses
 - Universal Replicator restore journal group
- CCI at all three sites recommended.
 - If CCI is not installed at Intermediate Site, **Remote Command Device** must be defined in Intermediate Site so that CCI can manage the volumes from the Primary Site (see note).

Pair operation process

1. Primary site application writes data to TrueCopy Synchronous P-VOL.
2. TrueCopy sends data to Intermediate site cache.
3. Intermediate Site send acknowledgement back to Primary Site. TC releases host I/O.
4. In Intermediate Site, UR creates metadata and queues data for destage to master journal volume (Normal Journal Obtain Process).
5. Normal UR Journal Copy moves Journal Updates to RCU in the remote site (Normal Journal Copy).
6. Normal UR Journal Restore operation settles and writes data to UR S-VOLs (Normal Journal Restore).

Suspend and Resync process

1. When the TrueCopy Synchronous pair is suspended, differential data between the primary and secondary volume is managed by normal TC Bitmap Areas s, allowing normal resync of TC Sync pair.
2. When UR pair suspends, the TrueCopy Synchronous pair suspends as well. TC Sync Bitmap Areas become active.
3. When resyncing UR pair, changed tracks noted in the TC Sync bitmap are used to update the UR S-VOL.

- **System Option Mode 707**
 - **OFF:** Default Setting. TrueCopy Synchronous P-VOL data will be written on TrueCopy Synchronous S-VOL/UR P-VOL.
 - TC S-VOL/UR P-VOL can be Dynamic Provisioning volume.
 - UR Journal Obtain process operates normally.
 - Normal 3DC Cascade operation.
 - **ON:** UR intercepts TrueCopy Sync data and creates Journal Update to be queued to UR master journal volume (see note).
 - TC S-VOL/UR P-VOL LDEV must exist but is not used.
- **System option mode 707 must be set before UR journal group creation.**
 - If Mode 707 is changed after creating the Journal Group, there is no effect.
 - Delete the journal group, change the setting, and then create the journal group again.
 - Mode 707 can be seen in Storage Navigator, UR Journal Operations.


6

2DC Pass Thru Operations: System Option Mode 707 ON in Intermediate Site

TrueCopy data will not be stored in the TrueCopy Synchronous S-VOL (UR primary data volume). Instead, UR intercepts the data. TrueCopy reports normal write to secondary system and host I/O is released. UR initiates Journal Obtain process, creating metadata (and CTQ-Marker for MxN Consistency Group). The TrueCopy data (now UR data) is queued for destage to UR master journal volumes.

Normal UR Read Journal and Journal Copy processes move update data to the UR Remote site.

The normal UR Journal Restore process settles the update data by sequence number and restores journal volume in the remote site. According to the sequence numbers attached to the journal data, UR restores the restore journal volume data to the secondary data volume. In this way, the UR secondary data volume can keep the update order consistency with TrueCopy Synchronous P-VOL.



7

2DC Cascade PassThru is similar to 3DC Cascade – open systems only.

- Allows use of Hitachi Dynamic Provisioning volumes for data volumes at Intermediate Site.
- Primary and Remote Sites can contain intermix of any Hitachi enterprise storage systems.
- Intermediate site must be Universal Storage Platform V or VM system.
- This configuration provides synchronous replication to the intermediate site, and asynchronous replication to the remote site.
- Two operational modes, dependent on System Option 707.

Setup Procedure

1. Install TrueCopy Synchronous and UR at appropriate sites.
2. Configure ports and journal groups.
3. Set system option mode 707 on the UR primary storage system (intermediate site).
4. Create UR journal groups at Intermediate and Remote Sites.
5. Create UR pairs in the Intermediate Site. Use Mirror ID 1.
6. Confirm UR PAIR Status.


7. Create a TrueCopy Synchronous pair in the primary site.
8. Confirm that the TrueCopy Synchronous pair status has become **PAIR**.

Note: This procedure is the opposite of the normal 3DC pair creation order.

Configurations: 3DC Cascade

3DC Cascade Configuration — Three sites: Primary, Intermediate, and Remote

- TrueCopy Synchronous software operates between primary and intermediate sites
 - Universal Replicator operates between intermediate and remote sites
 - Storage Navigator required at all sites
 - CCI required at primary and remote sites, highly recommended at intermediate site



8


In the diagram and the following pages,

TrueCopy stands for TrueCopy Remote Replication

UR stands for Universal Replicator

CoW stands for Hitachi Copy-on-Write


Configurations: Cascade Mirror IDs



In this diagram and following pages,
TC stands for TrueCopy Remote Replication
SI stands for ShadowImage Replication


Configurations: Multi-target

- Three storage arrays: Primary, intermediate, and remote storage systems
- TrueCopy Synchronous operates between primary and intermediate sites
- Universal Replicator operates between primary and remote site
- CCI connected to primary, intermediate, and remote site




10

Configurations: Multi-target Bitmaps




Configurations: 4DC

- 4DC Combines 3DC Multi-target and Cascade
 - Four storage arrays: primary, local, two remote subsystems
 - TrueCopy Synchronous software operates between primary and local sites
 - Universal Replicator software operates between primary and remote site
 - CCI connected to *all* sites




The data volumes in the primary site function as TrueCopy P-VOLs and Universal Replicator P-VOLs concurrently.

12



Configurations: Multi-target with Third Link and Delta Resync

- The **Third Link** between the intermediate and remote sites enables quick resumption of remote replication in the event of failure at the primary site.
 - Without Delta Resynchronization, a full volume **paircreate** is required.
 - If Delta Resync function is installed, then only **pairresync** is required.
 - CCI is required at *all* sites.




14

3DC Multi-Target supports two configurations:


- **2-Link:** The Remote Replication software P-VOL in the primary site functions as both Remote Replication software P-VOL and UR P-VOL.
- **3-Link:** A link is provided between the Intermediate site and the remote site. This link greatly increases overall disaster recovery flexibility by providing a quick way to resume 3DC operations should there be a communications failure between the primary and remote sites.

Configurations: Three-Link Multi-target Bitmaps



Configurations: Mirror Unit Numbers

- Common to all 3DC configurations.
- Ensure Mirror Unit ID 1 is coded in **all** Universal Replicator horcm.conf files.



16

MU# for TrueCopy/UR: Defines the mirror unit number (0 - 3) of one of four possible TrueCopy and Universal Replicator bitmap associations for an LDEV. If this number is omitted, it is assumed to be zero (0). The Universal Replicator mirror description is described in the MU# column by adding *h* in order to identify identical LUs as the mirror descriptor for UR. The MU# for TrueCopy must specify *blank* or *0*.

Note: By convention, the MU# applied to Universal Replicator in 3DC environments is *h1* and must be coded in both **horcm.conf** files that define the Universal Replicator pairs.

Configurations

- 3DC Configuration items common to both Cascade and Multi-Target
 - **Uni-directional TrueCopy links** are required between sites containing TrueCopy pairs.
 - **Bi-directional Universal Replicator links** are required between sites containing Universal Replicator pairs.
 - Each host should be network-connected to each other to allow the use of CCI.
 - Prevent host I/O to TrueCopy S-VOL in intermediate site
 - **Mandatory Universal Replicator Fence Level “async”** for CCI
 - Intermix of Hitachi enterprise storage systems
 - Data volumes exist in primary, intermediate, and remote sites.
 - **Disaster Recovery Extended license** is required for Universal Storage Platform V and VM.

- Configuration Considerations
 - No change in pair status conditions due to 3DC.
 - No change in pair operations in 3DC.
 - Exception: Paircreate, Resync, and Resync swap operations that result in prohibited copy combinations are denied.
 - Pair status transition are the same.
 - Maintenance procedure of DKCs are the same.

Operations

- To begin 3DC replication:
 - If TrueCopy pairs exist, but are not members of TC Sync Consistency Group, convert TC pairs to Consistency group with **pairresync -fg** (see note).
 - If necessary, create TrueCopy pairs in TC Sync Consistency Group.
 - If Cascade, then immediately split TrueCopy groups.
 - Create Universal Replicator pairs, matching UR consistency groups with TC consistency groups.
 - If Cascade, resync TrueCopy pairs.

19

Note: CCI **pairresync** command with **-f[g] <fence> [CTGID]** (TrueCopy only):
Changes existing TC Sync volumes to TC Sync Consistency Group.

- Does not require deletion of pair.
- Split TC group with normal pairsplit command, then issue **pairresync -g <grpname> -fg <fence level> <CTG ID>**
- Only valid for normal resync at primary site. Not valid with -swaps or -swapp option from either primary or secondary site


Disaster Recovery with 3DC

- Overview of scripted Disaster Recovery CCI commands
 - **horctakeover** – Useful for temporary loss of primary or secondary site
 - Checks the specified volume's or group's attributes with **paircurchk**
 - Takeover Switch executes the chosen takeover function
 - Temporary loss of Remote Site – P-VOL-takeover
 - Both sites operational – Swap-takeover
 - Temporary loss of Primary Site – S-VOL-takeover
 - **horctakeoff** – New
 - Change 3DC multi-target to 3DC cascade while host applications are running
 - Conditions:
 - Temporary Primary Site failure, failover to Intermediate Site has occurred
 - Execute **horctakeoff**
 - Execute **horctakeover** to switch operations to the Remote Site without affecting Host operations at the Intermediate Site
 - Operates on either individual volume or volume group

20

Situation – 3DC Cascade with disaster at Primary site

- **Split** Universal Replicator pairs
- Issue either **horctakeover** or **pairsplit -RS** on TrueCopy S-VOLs at Intermediate Site
 - Forces TrueCopy S-VOLs to **SSWS** status
- **Resync** Universal Replicator pairs




21

After Primary Site recovers, execute **horctakeover** at Primary Site to change Primary TrueCopy volumes back to P-VOLs.

Situation: 3DC Cascade with disaster at both Primary and Intermediate sites

- Execute **horctakeover** or **pairsplit -RS** on Universal Replicator S-VOLs at Remote site
- Forces Universal Replicator S-VOLs to **SSWS** status




22

After Intermediate Site recovers:

- Execute **horctakeover** to change Intermediate Site Universal Replicator volumes back to P-VOLs.
- Re-establish TrueCopy pairs after Primary Site recovers.

Situation: 3DC Multi-Target with disaster at Primary Site, failover to Intermediate site

- Execute **horctakeover** or **pairsplit -RS** on TrueCopy S-VOLs at Intermediate site
- Forces TrueCopy S-VOLs to **SSWS** status




23

After the Primary Site recovers:

Execute **horctakeover** to change Primary Site TrueCopy volumes back to P-VOLs.

- Situation: 3DC Multi-Target with disaster at Primary site, failover to REMOTE site
 - Execute **horctakeover** or **pairsplit -RS** on Universal Replicator S-VOL at Remote site
 - Forces Universal Replicator S-VOLs to **SSWS** status




- When Primary Site recovers, execute **horctakeover** at Primary site.

24

When Primary Site recovers, execute **horctakeover** from Primary site.

Situation: 3DC Multi-Target with disaster at Intermediate site

- Universal Replicator Pairs are not affected.
- Resync TrueCopy Remote Replication pairs after Intermediate Site recovers.



25

3DC Disaster Recovery

- 3DC Multi-target failback to Primary site
 - **From Local Site**
 - After recovery of former primary site function:
 - At the former primary site, use CCI **horctakeover** operations on the TrueCopy Volumes.
 - Resync or re-create the Universal Replicator pairs as needed.
 - **From Remote Site**
 - At the former primary site, use CCI **horctakeover** operations on the Universal Replicator volumes.
 - Resync or re-create TrueCopy pairs as needed.

26

Module Review

1. True or False. Three data center replication supports TrueCopy Asynchronous software.
2. True or False. TrueCopy Synchronous software can be cascaded behind Universal Replicator software.
3. True or False: The TrueCopy Initial Copy pair can be started after the Universal Replicator Initial Copy.

9. Delta Resynchronization


Module Objectives

- Upon completion of this module, the learner should be able to:
 - Describe concepts and specifications of Hitachi Universal Replicator Multi-target Delta Resync function
 - Configure Multi-target Delta Resync
 - Describe in detail the sequence of Delta Resync operations
 - Outline the new Command Control Interface functionality added to support Delta Resync


2

Concepts: Multi-target 3DC Review


- Review: Hitachi TrueCopy® Synchronous software and Hitachi Universal Replicator software in 3DC multi-target configuration



- A third link can be provided between local and remote sites to ensure that a remote replication can continue if the loss of primary site occurs.
- Universal Replicator pairs can be established using the third link when primary site fails.
 - Requires full Initial Copy to create




- Third link Universal Replicator volumes are *pre-created* with Delta option and remain in **HOLD** status until resynced in the event of primary Universal Replicator failure.
- Supports Remote Command Device



5

■ VSP Remote Command Device



6

Remote Command Devices: When the delta resync operation has been performed and pair status is changed to PAIR, the delta resync P-VOL must be updated from the host for longer than five minutes. This is required to insure internal communications between the intermediate and remote sites.

However, you can work around this five-minute-plus update requirement by setting up command devices and remote command devices. With remote command devices set up, communications between the two sites is performed automatically, and the delta resync is ready to use when the operation is run.

This requires setting up two command devices and two remote command devices on each site — the local, intermediate, and remote sites — as explained in the following general guidelines. Consult the *Hitachi Universal Volume Manager User Guide* for more complete information about remote command devices.

1. Create four command devices each on the local, intermediate, and remote sites.
2. Set up and dedicate two external ports and two target ports on each site for the command/remote command devices. Configure paths between external ports and target ports. For details about the external ports, see *Hitachi Universal Volume Manager User Guide*. For instructions on setting paths, see *Provisioning Guide for Open Systems*.

3. On each site, map a command device via a target port to a device on one of the other sites. The device on the other site should be mapped to as a remote command device, using an external port on that system.
4. Repeat the previous step so that two command devices on each site are mapped to a remote command device on each of the other two sites.

- **Delta Journal Operation**
 - Normal resync operates by transferring changed tracks noted in the differential bitmaps.
 - Delta resync operates by **journal copying**.
- If TrueCopy pairs are in **PAIR** at the time of Primary Site failure:
 - Journal data of Delta Resync Universal Replicator Pairs in Local Site matches Journal Data in Primary Site Replicator Pairs up to the point of failure.
 - Takeover operation at local site conditions TrueCopy S-VOLs to be used as P-VOLs. These volumes are already the P-VOLs of the Delta Resync Universal Replicator pairs.
 - Applications coming up at local site results in new Delta Journal Data.
 - When Delta Resync occurs, only this new Journal Data in local site must be sent to remote site.

7

Journal data in the Local Site are copied to the Universal Replicator secondary (Remote) site by journal copy. Only the journal data, which are not yet restored to the secondary data volume in the **Universal Replicator** secondary site, are copied in chronological order. When the journal copy completes, journal restore takes place in the Universal Replicator secondary site.

Note: When the total capacity of stored journal data exceeds 80% of the TrueCopy Synchronous secondary site's Delta Resync journal volume, old journal data will be automatically deleted. If the TrueCopy pairs are in pair state at the time of Primary Site failure, or they are synchronized periodically, the total capacity of the journal data which is not restored to the Universal Replicator secondary site will not exceed 80% of the Delta Resync journal volume.

Therefore, if the Delta Resync journal data volume utilization exceeds 80% the secondary data volume will not be able to be restored completely. In that case, according to the setting of the journal group option, all tracks of the Delta Resync pairs in the Local Site will be copied to the Universal Replicator secondary data volumes in the Remote Site, or delta resync operation finishes without any processing.

Usually, if the pair between TrueCopy Synchronous primary site and secondary site is synchronized periodically, the total capacity of the journal data that is not restored to the Universal Replicator secondary site will not exceed 80% of the journal volume. Though, for example if the Universal Replicator pair is suspended and the pair has

not been resynchronized for a long time, journal data of more than 80% of the journal volume capacity may be stored before they are restored to Universal Replicator secondary data volume. In such cases, you may not perform delta resync operation properly.


WARNING: Even if the capacity of the journal data does not exceed 80% of the journal volume, note that journal data will or may be destroyed in the following cases.

- When you restore the TrueCopy Synchronous pair, then update the P-VOL
- When you restore the Universal Replicator pair between the primary site and the Universal Replicator secondary site, then update the P-VOL
- When the retry processing occurs because of a delay of the P-VOL update
- When the update of the TrueCopy Synchronous S-VOL is delayed


- If TrueCopy volumes are normally run in **SPLIT** status with periodic resync:
 - When the TrueCopy pairs are resynced, all changed data sent to the TrueCopy S-VOLS in the local site is reflected to the Delta JNL volumes.
 - By default, if Delta JNL volumes utilization goes above 80%, old JNL data is discarded. In the event of Primary site failure at this point, Delta Resync cannot be performed. All data at the local site must be sent to the remote site.
 - For best results, resync the TrueCopy pairs often to keep the Delta Journal volume utilization low.

8

- 3DC and Delta Resync Mirror IDs



9



Specifications

Delta Universal Replicator between local site and remote site

#	Item	Specification
1	Number of pairs	Normal Universal Replicator specification of 8192 data volumes per Journal Group. The total number of pairs between local site and remote site is model-dependent (16K for Universal Storage Platform, 32K for Universal Storage Platform V and VM).
2	JNL Groups	Normal Universal Replicator specification
3	Link between Local Site and Remote Site	Normal Universal Replicator bidirectional links required
4	Extenders between Local and Remote Site	Normal Universal Replicator specification

If remote site is updated earlier than local site, Delta resync cannot be performed. To avoid the problem, journals arriving at remote site are delayed by one minute before writing to the Universal Replicator S-VOL.

Pair status review



	SMPL	The volume is not currently assigned to a pair. When the initial copy is started by a paircreate operation, the volume status changes to Copy.
	COPY	The initial copy operation is in progress. Data in the primary data volume is not synchronized with data in the secondary data volume. When the initial copy is complete, the status will change to PAIR.
	PAIR	The volume is paired with another volume. The two volumes are fully synchronized. All updates from the host to the primary data volume are duplicated at the secondary data volume.
	PSUS	The pair was split by the user (pairsplit-S). The primary data volume and the secondary data volume are not synchronized.
	PSUE	The pair has been split by the primary storage system or the secondary storage system due to an error or failure. The primary data volume and the secondary data volume are not synchronized.
	Suspending	The primary data volume and the secondary data volume are not synchronized. This pair is in transition from the PAIR or COPY status to the PSUS/PSUE status to the SMPL status.
	Deleting	The primary data volume and the secondary data volume are not synchronized. This pair is in transition from the PAIR, COPY, or PSUS/PSUE status to the SMPL status.

Added for Delta Resync

	HOLD	A Universal Replication pair for delta resync operation is created.
	HLDE	An error occurred with the Universal Replication pair for delta resync operation.

■ Journal Group status – Delta Resync Status Conditions

#	Pair/JNLG	Status	Meaning
1	Pair status	HOLD	The journal data for delta resync operation is stored.
2		HLDE	The journal data for delta resync operation cannot be created due to failure. HLDE: HoLD Error
3	JNLG status	HOLD	The journal data for delta resync operation is stored.
4		HLDE	The journal data for delta resync operation cannot be created due to the failure.




13

Configuration

■ Preliminary Conditions:

- Full 3DC Multi-Target configuration is defined with all TrueCopy and Universal Replicator volumes in PAIR status.
- Necessary Universal Replicator links are defined between local site and remote site.
- Remote DKC is defined for both local and remote sites.
- HDS Technical Personnel will ensure that SVP Mode 506 is set ON for all storage systems involved.



14

To create a Universal Replicator pair for delta resync operation, the following items are required.


- Create the pair in 3DC multi-target configuration.
- Use TrueCopy Synchronous S-VOL in **PAIR** status as the primary data volume.
- Use Universal Replicator data volume in **PAIR** status as the secondary data volume.
- Use mirror ID 2 (0 is used by TrueCopy and 1 is used by primary Universal Replicator).
- The system option mode 506 must be set to **ON** at all sites.

In addition to those requirements, all Universal Replicator pairs in the journal group must satisfy the following requirements when you create more than one Universal Replicator pair for delta resync operation.

- Use the same mirror ID for all pairs.
- Use the same restore journal group for all pairs.

1. To Create Delta Resync Pairs:


- a. Create required Journal Groups at local site.
- b. Check that all TrueCopy and Universal Replicator volumes are in PAIR status.
- c. Create Delta Resync pairs using the new Journal Groups in the local site and the existing Journal Groups in the Remote Site.
 - Make sure to select a unique Consistency Group ID.
- d. Check that all volumes are in HOLD status.



15

2. Configure Port Attributes for Universal Replicator Replication on both primary and remote storage systems.**a. Verify Initiator and RCU Target Creation**

- Click on the **RCU Target** component within the tree view. You should see your RCU targets there.
- Click on the **Initiator** component in the tree view.
 - See Initiators here.



16

3. Add DKC to both Local and Remote Storage Systems

- a. Enter the serial number of the remote DKC.
- b. **Controller ID = 5** (Universal Storage Platform and Network Storage Controller)
- c. Select the ports used for Universal Replicator Links.
- d. Click the **Option** button.

Note: The Port column signifies the Initiator ports and the Pair-Port column contains the associated RCU Target port

Storage Navigator

Add RCU (Fibre)

RCU S/N:

LDKC:

Controller ID:

Logical Adr.(RCU CU):

SSID:

MCU Port	RCU Port

Option Cancel

17

When you assign Logical Paths, use the port allocations you set as Initiator and RCU Target. Make sure an Initiator and RCU Target are assigned together. If two Initiators are grouped together, this will cause an error.

Display Features

DKC S/N: Allows you to enter the five-digit serial number of the remote storage system.

Controller ID: Allows you to enter the controller ID (storage system family ID) of the remote storage system.

Note: The controller ID for a Universal Storage Platform is 4.

Path Gr. ID: Allows you to enter the path group ID. Path group IDs are used for identifying groups of logical paths. One path group can contain up to eight logical paths.

Note: In the current version, you cannot enter path group IDs. Also, you cannot clear the **Default** check box. The number of path groups per one remote subsystem is always 1.

M-R Path: Allows you to specify logical paths from initiator ports on the local subsystem to RCU target ports on the remote subsystem.

Port: Displays a list of initiator ports on the local subsystem. Select an initiator port from this drop-down list, or type in a Port Number.

Pair-Port: Displays a list of all ports on the remote subsystem. Select an RCU target port on the remote subsystem from this drop-down list, or type in a port number.


Note: When specifying a port, you can use the keyboard to enter the port number. When you enter the port number, you can abbreviate the port number into two characters. For example, you can enter **1A** instead of **CL1-A**. You can use uppercase and lowercase letters.

Option: Opens the DKC Option panel

Cancel: Cancels the settings you made on the Add DKC panel and then closes the panel.

4. Create Journal Group on Local Site

- Select the PG, ExPG, or Control Unit Number of the candidate LDEV.
- Select LDEV.
- Click **Add**.
 - Add up to 64 volumes (only Open-V)
- Click **Set**.
- Repeat this step for the remote storage system.



18

The Edit JNL Volumes panel displays similar information about:

JNL Volumes – Journal Volumes

Free Volumes – Not registered in journal groups

For each category of volumes:

Parity Group: Indicates the parity group or the external volume group where a journal volume belongs.

Note: If the letter 'E' is displayed at the beginning of a group, the group is an external volume group. In the current version, however, the panel does not display external volumes.

LDKC:CU:LDEV: The LCKC number is displayed to the left of the colon, the CU number is displayed between the colons, and the LDEV number is displayed to the right.

Note: If a pound sign (#) is displayed at the end of a volume, the volume is an external volume. In the current version, however, the panel does not display external volumes. Consider the implications carefully before selecting an external volume as a Journal Volume.

Capacity: Indicates the capacity of a journal volume in gigabytes.

Emulation: Indicates the emulation type of the volume.

Operation: Displays one of the following:

Blank: This column usually blank

Add: Indicates a volume to be added to a journal group

Delete: Indicates a volume to be deleted from a journal group

JNL Volume Buttons

Add: To register volumes in a journal group, select the volumes from **Free Volumes** and click **Add**.

Delete: To delete volumes from a journal group. Select the volumes from **JNL Volumes** and click **Delete**.

Parity Group/CU change

To change the display in the **Free Volumes** list:

Parity Group: Select to display volumes belonging to a parity group. Specify a parity group number in the text boxes to the right, and then click the **Show** button.

PG(Ext.): Select to display *external* volumes belonging to a parity group. Specify a parity group number in the text boxes to the right, and then click the **Show** button.

CU: Select to display volumes belonging to a CU. Then select a CU from the list to the right.

■ Setting Journal Option in Local Site

– Delta Resync Failure


- **Entire** – With this setting, when delta resync fails, **Entire** copy is automatically executed.
- **None** – No operation. In this case, it is necessary to delete the pair between Local Site and Remote Site, and to create the pair again. Recommended.

Change JNL Option	
Inflow Control :	☒ Yes ☐ No
Data Overflow Watch (0 to 600 sec) :	60
Copy Pace :	Low
Unit of Path Watch Time :	☒ min. ☐ hour ☐ day
Path Watch Time (1 to 99 min., 1 to 23 hour, 1 to 30 day) :	5
Forward Path Watch Time :	No
Use of Cache :	Use
Speed of Line (Mbps) :	256
Delta resync Failure :	Entire
Set Cancel	

19

■ Create Delta Universal Replicator pair

- Use TrueCopy Synchronous S-VOL in PAIR status as the primary data volume.
- Use Universal Replicator S-VOL in PAIR status as the secondary data volume.
- Use mirror ID 2 (1 will already be used by primary 3DC Universal Replicator).



20

Delta Universal Replicator pair creation command from Storage Navigator

Paircreate

Port-G-ID-LUN

P-VOL : CL7-F - 00 - 000

S-VOL : CL1-A 00 000

Select Other S-VOL(s) :

M-JNL Mirror ID R-JNL

Mirror : 00 00 01

CT Group : 00

DKC S/N, CTRLID, Path

DKC : 11111, 04, Fibre

Select Delta_

Initial Copy : Entire

Priority : 32


Error Level : Group

M-JNL Information :
Current Mirror(s) : 0 Total Mirror(s) : 0

Set Cancel

Set Mirror ID 2 for
Delta Resync pairs

Manually set CT Group ID
to unique value.



21

Initial Copy: Allows you to specify whether to start the initial copy operation after the volume pair is created. The default is **Entire**:

Entire: The initial copy operation will start after the volume pair is created. When the initial copy operation executes, all data on the primary data volume will be copied to the secondary data volume.

None: The initial copy operation does not start after the volume pair is created. The primary storage system starts copying of update data as needed.


Note: The user must ensure that the primary data volume and secondary data volume are already identical when using None.

Delta: An initial copy operation will not start after the volume pair is created. The status of the volume pair will change to HOLD which means that the pair is for delta resync operation.

Note: Manually set CT Group ID to unique value.

▪ Resynchronization of Delta Universal Replicator pairs in the following cases:

- During HOLD ERROR (HLDE) recovery
- During Delta resync




Delta Resync pairresync command options

#	Aim	Mgmt Software	Command	Description
1	Recovery from HLDE	Storage Navigator	Pairresync	Select Return to standby as the resync mode on the pairresync panel.
2				
4	Delta Resync	Storage Navigator	Pairresync	Select Delta as the resync mode on the pairresync panel.
5				

22

▪ Recovery from HLDE status: Can only be done from SVP

1. Check that status is HLDE.
2. Resync the Delta Pair with **Resync Mode** option **Return to Standby**.
 - Status should change to HOLD.



23

Sequence of Delta Resync

- To execute delta resync:
 - Check that the following conditions are satisfied:
 - Split Primary Universal Replicator pairs. Status of primary Delta Universal Replicator pair is PSUS, PSUE, or HOLD.
 - Issue either takeover operation or CCI command **pairsplit –RS** on TrueCopy S-VOLs at Local site. Status of TrueCopy S-VOLs must be SSWS or PAIR with reversed TrueCopy direction
 - Issue the Delta Resync Operation from Storage Navigator.
- System Option Mode 506 controls whether or not Delta Resync will result in a full-volume copy if there are no Delta Resync JNL Data to be sent (see note).
- If JNL Update data are present, Delta Resync Universal Replicator pair will change directly to PAIR.
- If Journal at Local site fills, Delta Resync will fail (see note).

24

Note: The delta resync operation steps include first using journal copy to copy the journal data in the Local site to the UR secondary site. **Only the journal data that is not yet sent to the UR secondary site are copied in chronological order.** So, if no changed journal data is present at the local site, no delta journal is available, and the delta resync fails. It will default to full volume copy in this case.

In delta resync operation, the status of the UR pair changes to **PAIR** (not **COPY**). This is because the delta resync operation sends journal updates (not changed tracks from the differential bitmap). Therefore, delta resync operation requires less time to recover the UR pair after a failure occurs.

When the total capacity of stored journal data exceeds 80% (PFUL Status) of the journal volume of TrueCopy Synchronous secondary site, old journal data is automatically deleted. Therefore, the secondary data volume is not completely restored to the UR secondary site. In that case, either the entire primary data volume is copied to the secondary data volume (ALL JOURNAL Copy) or delta resync operation finishes without any processing.

System Option mode 506 Universal Replicator, Universal Replicator for z/OS

Enables Delta Resync with no host update I/O by copying only differential JNL instead of copying all data.

- Mode 506 = ON (default):

Without update I/O: Delta Resync is enabled. With update I/O: Delta Resync is enabled.

- Mode 506 = OFF: Without update I/O: Total data copy of Delta Resync is performed. With update I/O: Delta Resync is enabled.

Note: Even when mode 506 is set to ON, the Delta Resync may fail and only the total data copy of the Delta Resync function is allowed if the necessary journal data does not exist on the primary subsystem used for the Delta Resync operation.


■ Delta Resync From Local Site

When status is HOLD, Resync Mode is **Delta**.

Pairresync


Range :	Group
Priority :	76
DKC :	111114,Fibre
Resync Mode :	Delta
Error Level :	Group

Set **Cancel**



25

- After Delta resync operation, Universal Replicator between NEW primary site and remote site transitions to PAIR status. Configuration is 2DC.
 - You can change the system back to 3DC multi-target configuration after removing failures from the primary site and issuing the **resync -swaps**. In the resulting 3DC multi-target configuration, the local site will become the primary site.



26

■ Delta Resync Failure Condition

– Storage Navigator

- If **Delta resync failure** of JNLG option is set to **Entire**, entire copy from local site to remote site is automatically started.
- If **Delta resync failure** of JNLG option is set to **None**, you have to delete the pair between local site and remote site, then create the pair again.

CCI Support for Delta Resync

- Delta Resync Universal Replicator **paircreate**

#	Mgmt Software	Command	Description
1	Storage Navigator	Paircreate	Select Delta as the Initial Copy option on paircreate operation panel.
2	CCI	paircreate	Specify nocsus option

- When CCI is used, **paircreate** command can be issued from primary, local, or remote sites:
 - Primary site example: (See Student Notes for FHORC usage.)
`paircreate -g G1 -gs G2 -FHORC 2 -nocsus -f async <ctgid> -jp <id> -js <id>`
 - Local site example: (Usual method)
`paircreate -g G3 -vl -nocsus -f async <ctgid> -jp <id> -js <id>`
 - Remote site example:
`paircreate -g G3 -vr -nocsus -f async <ctgid> -jp <id> -js <id>`

Note: **paircreate** with **nocsus** option is the **only** supported CCI function for Delta Resync. All other Delta Resync operations **must** be performed with Storage Navigator.

28

FHORC [MU#] or -FCA [MU#]

This option is used to create the cascading configuration with -g <group> and -gs <group> option from the local node (takeover node).

-g <group> is used for specifying the cascading P-VOL, and also -gs <group> option is used for specifying the cascading S-VOL.

This operation ignores -vl or vr option, because S-VOL will be specified with -gs <group> option.

-gs <group>

This “s” option is used to specify a group name for cascading S-VOL (defined in the configuration definition file).

Review

- 3DC Multi-target Delta Resync Checklist:
 - Using Storage Navigator:
 1. Define Port Attributes for the Delta Resync links from the **Local CU** to the **Remote CU**.
 2. Add a Remote DKC definition to the **Local** site.
 - This activates the links added above.
 3. Create Journal Groups in the **Local** site (to be used for the Delta Resync pairs).
 4. Create Pairs using correct Mirror ID and Initial Copy **Delta**.
 - Delta pairs can be created with CCI if desired.
 5. Issue Delta Resync from Storage Navigator only.
 - There is no support to resync a Delta pair in **HOLD** status with CCI.

29

Module Review

1. The mirror ID of the Delta Resync copy group must be different from that of the Universal Replicator copy group. What Mirror ID do we use for Delta Resync?
2. True or False: The M-JNL Group of the Delta Resync group are the TrueCopy Synchronous S-VOLs.
3. What must be the R-JNL Group IDs of the Delta Resync pairs?

10. Data Transport Technologies

Module Objectives


- Upon completion of this module, the learner should be able to:
 - Identify supported topologies for Hitachi Fibre Channel (FC) replication links
 - Identify long distance data transport technologies
 - Examine the role of buffer credits in flow control
 - Identify the different SAN extension options and compare their operational characteristics:
 - Dark fiber
 - Synchronous Optical Network (SONET)
 - Synchronous Digital Hierarchy (SDH)
 - Coarse Wavelength Division Multiplexing (CWDM)
 - Dense Wavelength Division Multiplexing (DWDM)
 - Fibre Channel over Internet Protocol (FCIP)

2

Supported Link Topologies

- Universal Replicator Fibre Channel Connections (topologies)
 - Direct connect
 - Extended connect through dedicated switches
 - Extended connection through DWDM
 - Extended connection through WAN converter

- Direct Connect: Port Setting = FABRIC OFF, FC-AL (see note)
 - Simplest Universal Replicator configuration
 - No switches necessary. Runs on dark fiber
 - Distance Limitations:
 - Shortwave multimode Fibre Channel – 500 Meter
 - Longwave single-mode Fibre Channel – 10km



4

Note: Hitachi implements the Point-to-Point configuration as a special case of a two-node loop called Direct Connect. Fibre Channel settings for Direct Connect are FABRIC OFF, FC-AL (same as Arbitrated Loop). Standard point-to-point Fibre Channel (for example, without a converter or extender) uses a short-wave optical connection between N-ports at distances up to 500 meters.

The Universal Replicator software link connection is dedicated and not be mingled with the regular SAN. The application writing data to the Universal Replicator software P-VOL can't be allowed to "see" the secondary volume of the pair (and visa versa). To do so would cause an LVM error. Therefore, some method of preventing the host operating systems must be used to prevent this. Unless otherwise noted, this requirement exists for all configurations

Each Fibre Channel link consists of an optical fibre pair. Even though there are two fibres in the pair, they are used to move data in a one-way fashion. As a pair, the bulk of the traffic is intended to be in a single direction (for example, Initiator to RCU Target) with the second fibre used just for acknowledgements.


When fibre-channel multimode shortwave connections are used, two switches are required for distances greater than 0.5 km (1,640 feet), and distances up to 1.5 km (4,920 feet, 0.93 miles) are supported.

When fibre channel interface single-mode longwave connections are used, two switches are required for distances greater than 10 km (6.2 miles), and distances up to 30 km (18.6 miles) are supported.

▪ Dedicated pairs of FC Switches


– Distance Limitations

- Shortwave multimode Fibre Channel – 1.5km
- Longwave single-mode Fibre Channel – 30km
- “Hard-zone” switches to prevent host access to links



5

- Dense Wave Division Multiplexer (DWDM)
 - Medium Distance to Long Distance FC
 - If DWDM converters do not provide buffer crediting
 - Use FC switches for buffer credits
 - Hard-zone switches to prevent host access to links




6

This topology involves the use of a passive optical multiplexing device known as Dense Wave Division Multiplexer (DWDM). In this case, the regular storage area network (SAN) and Universal Replicator software traffic can share the DWDM link as long as no cross connections to Universal Replicator software ports are possible.

Note: When using DWDM or other passive devices (at distances of more than a few km) that do not perform data store-and-forward or some sort of adequate buffer-crediting, it is a requirement to insert a Fibre Channel switch (as shown) on each side of the passive DWDM link to provide a buffer credit function.

■ Switched Converters – Any Distance

- Converters can be FCIP, ATM, Frame Relay, and more.
- Hard-zone switches to prevent host access to links.
- WAN optimizers can provide additional functionality such as compression.



7

Note: All Universal Replicator software link traffic must be in an isolated zone. Hosts must not be able to access any ports which connect to Initiator ports.

This interconnect option involves the use of a box that converts Fibre Channel to something else suitable for very long distance switched circuit such as T3, ATM, or switched packet (FCIP) transmission.

The requirement for a switch (between the array and converter), is to translate from N or NL port to E-Port. The regular SAN can share the same converter boxes only in a way that prevents hosts owning the primary volumes from accessing Universal Replicator software ports. Beyond this, the need for LUN Security still applies.

Long-Distance Transport Technology

- T Links (also called Digital Signal Links)
 - T1: 1.544Mbit/sec
 - T3 (comprised of 28 T1 lines): 44.736Mbit/sec
 - Fractional T3: 3 to 45Mbit/sec
 - Fractional T1: 256/ 384/ 512/ 768Kbps

T1 is a high speed digital network (1.544Mbit/sec) implemented in the early 1960s by AT&T to support long-haul pulse-code modulation (PCM) voice transmission. There are also T1-C, T2, T-3 and T-4 networks. T1-C operates at 3.152Mbit/sec. T-2, implemented in the early 1970s to carry one Picturephone or 96 voice channels, operates at 6.312Mbit/sec. T-3 operates at 44.736Mbit/sec and T-4 operates at 274.176Mbit/sec.

T3 (also known as a DS-3) is equal to approximately 672 regular voice-grade telephone lines, which is fast enough to transmit full-motion, real-time video, and very large databases over a busy network. A T3 line is typically installed as a major networking artery for large corporations and universities with high-volume network traffic. A T3 is the second fastest, non-optical connection offered in North America. A T3 line is comprised of 28 T1 lines, each operating at total signaling rate of 1.544Mbit/sec.

Fractional T3

Fractional to full DS3 or T3 circuits run from speeds of 3Mbit/sec up to 45Mbit/sec. For customers that require more bandwidth than a T1 can provide, but do not require the capacity of a full T3 line, a fractional T3 connection is ideal. A fractional T3 is similar to a full T3, only with some of the channels turned off. Unfortunately, the T3 loop is still required for this service.

Fractional T1

Essentially a T1 line with some of the channels turned off. Typical speeds for fractional lines are 256, 384, 512 and 768Kbps. Most providers that offer full connections also offer fractional service.

▪ Important Data Rates

Digital Signal Designator	Data Rate	DSO Multiple	T-Carrier	E-Carrier
DS0	64Kbps	1	-	-
DS1	1.544Mbit/sec	24	T-1	-
DS1C	3.152Mbit/sec	48	-	-
DS2	6.312Mbit/sec	96	-	-
DS3	44.736Mbit/sec	672	T-3	-
DS4/NA	139.264Mbit/sec	2176	-	-
DS4	274.176Mbit/sec	4032	-	-

9

- Digital signal X is a term for the series of standard digital transmission rates or levels based on DS0, a transmission rate of 64Kbps, the bandwidth normally used for one telephone voice channel.
- Both the North American T-carrier system and the European E-carrier systems of transmission operate using the DS series as a base multiple. The digital signal is carried inside the carrier system.
- DS0 is the base for the digital signal X series.
- DS1, used as the signal in the T-1 carrier, is 24 DS0 (64Kbps) signals transmitted using pulse-code modulation (PCM) and time-division multiplexing (TDM).
- DS2 is four DS1 signals multiplexed together to produce a rate of 6.312Mbit/sec.
- DS3, the signal in the T-3 carrier, carries a multiple of 28 DS1 signals or 672 DS0s or 44.736Mbit/sec.
- DS4/North America: Four DS3 lines
- DS4: Eight DS3 lines

- **Frame Relay**

- Encapsulates IP data packets within frame relay packets
- Efficient protocol-independent WAN transport medium
- Lower cost and higher performance than T1 or T3
- Typical speeds -dialup, T1/ T3

10

Frame relay has been a popular Wide Area Network protocol. A company running Ethernet can send the Ethernet protocol across a carrier's frame relay network and have it come out at the destination location in the Ethernet format. Frame relay encapsulates the data packets being sent inside of the frame relay packet, then breaks the frame packet apart once it arrives at the destination location. Frame relay is very good at efficiently handling high-speed data over wide-area networks; specifically LAN to LAN communications. It offers lower costs and higher performance for those applications in contrast to the traditional T1 or T3 services. As the frame relay network is a shared, switched network, there is no need for dedicated private lines, although special-purpose local loops (either DS0, T1 or T3 level connections) connect each location to a frame switch.

Frame Relay can be deployed using typical connection speeds including dialup, DS0, T1, and T3. Most companies use T1 loops which allow for port speeds from 56k to 1.5Mbit/sec. If a port speed of 56k will suit your needs, there is no need to purchase a T1 local loop. A DS0 (one channel of a T1) will provide up to 64Kbps of throughput for far less than a T1 would cost.

Technology	Speed	Physical Medium	Application
GSM mobile telephone service	9.6 to 14.4 Kbps	RF in space (wireless)	Mobile telephone for business and personal use
High-Speed Circuit-Switched Data service (HSCSD)	Up to 56 Kbps	RF in space (wireless)	Mobile telephone for business and personal use
Regular telephone service (POTS)	Up to 56 Kbps	twisted pair	Home and small business access
Dedicated 56Kbps on frame relay	56 Kbps	Various	Business e-mail with fairly large file attachments
DSO	64 Kbps	All	The base signal on a channel in the set of Digital Signal levels
General Packet Radio System (GPRS)	56 to 114 Kbps	RF in space (wireless)	Mobile telephone for business and personal use
ISDN	BRI : 64 Kbps to 128 Kbps PRI : 23 (T-1) or 30 (E1) assignable 64-Kbps channels plus control channel; up to 1.544 Mbps (T-1) or 2.048 (E1)	BRI: Twisted-pair PRI: T-1 or E1 line	BRI: Faster home and small business access PRI: Medium and large enterprise access
IDSL	128 Kbps	Twisted-pair	Faster home and small business access
AppleTalk	230.4 Kbps	Twisted pair	Local area network for Apple devices; several networks can be bridged; non-Apple devices can be connected
Enhanced Data GSM Environment (EDGE)	384 Kbps	RF in space (wireless)	Mobile telephone for business and personal use
satellite	400 Kbps (Direct PC and others)	RF in space (wireless)	Faster home and small enterprise access
frame relay	56 Kbps to 1.544 Mbps	Twisted-pair or coaxial cable	Large company backbone for LANs to ISP ISP to Internet infrastructure
DS1/T-1	1.544 Mbps	Twisted-pair, coaxial cable, or optical fiber	Large company to ISP ISP to Internet infrastructure

11

Universal Mobil Telecommunications Service (UMTS)	Up to 2 Mbps	RF in space (wireless)	Mobile telephone for business and personal use (available in 2002 or later)
E-carrier	2.048 Mbps	Twisted-pair, coaxial cable, or optical fiber	32-channel European equivalent of T-1
T-1C (DS1C)	3.152 Mbps	Twisted-pair, coaxial cable, or optical fiber	Large company to ISP ISP to Internet infrastructure
IBM token ring 802.5	4 Mbps (also 16 Mbps)	Twisted-pair, coaxial cable, or optical fiber	Second most commonly-used local area network after Ethernet
DS2/T-2	6.312 Mbps	Twisted-pair, coaxial cable, or optical fiber	Large company to ISP ISP to Internet infrastructure
Digital Subscriber Line (DSL)	512 Kbps to 8 Mbps	Twisted-pair, (used as a digital, broadband medium)	Home, small business, and enterprise access using existing copper lines
E-2	8.448 Mbps	Twisted-pair, coaxial cable, or optical fiber	Carries four multiplexed E-1 signals
cable modem	512 Kbps to 52 Mbps	Coaxial cable (usually uses Ethernet); in some systems, telephone used for upstream requests	Home, business, school access
Ethernet	10 Mbps	10BASE-T (twisted-pair); 10BASE-2 or -5 (coaxial cable); 10BASE-F (optical fiber)	Most popular business local area network (LAN)
IBM token ring 802.5	16 Mbps (also 4 Mbps)	Twisted-pair, coaxial cable, or optical fiber	Second most commonly-used local area network after Ethernet
E-3	34.368 Mbps	Twisted-pair or optical fiber	Carries 16 E-1 signals
DS3/T-3	44.736 Mbps	Coaxial cable	ISP to Internet infrastructure Smaller links within Internet infrastructure

12

<u>OC-1</u>	51.84 Mbps	Optical fiber	ISP to Internet infrastructure Smaller links within Internet infrastructure
High-Speed Serial Interface (HSSI)	Up to 53 Mbps	HSSI cable	Between router hardware and WAN lines Short-range (50 feet) interconnection between slower LAN devices and faster WAN lines
<u>Fast-Ethernet</u>	100 Mbps	100BASE-T (twisted pair); 100BASE-T (twisted pair); 100BASE-T (optical fiber)	Workstations with 10 Mbps Ethernet cards can plug into a Fast Ethernet LAN
Fiber Distributed-Data Interface (FDDI)	100 Mbps	Optical fiber	Large, wide-range LAN usually in a large company or a larger ISP
T-3D (DS3D)	135 Mbps	Optical fiber	ISP to Internet infrastructure Smaller links within Internet infrastructure
E-4	139.264 Mbps	Optical fiber	Carries 4 E3 channels Up to 1,920 simultaneous voice conversations
<u>OC-3/SDH</u>	155.52 Mbps	Optical fiber	Large company backbone Internet backbone
<u>E-5</u>	565.148 Mbps	Optical fiber	Carries 4 E4 channels Up to 7,680 simultaneous voice conversations
<u>OC-12/STM-4</u>	622.08 Mbps	Optical fiber	Internet backbone
<u>Gigabit Ethernet</u>	1 Gbps	Optical fiber (and "copper" up to 100 meters)	Workstations/networks with 10/100 Mbps Ethernet plug into Gigabit Ethernet switches
OC-24	1.244 Gbps	Optical fiber	Internet backbone
<u>SciNet</u>	2.325 Gbps (15 OC-3 lines)	Optical fiber	Part of the <u>vrns</u> backbone
<u>OC-48/STM-16</u>	2.488 Gbps	Optical fiber	Internet backbone
<u>OC-192/STM-64</u>	10 Gbps	Optical fiber	Backbone
<u>OC-256</u>	13.271 Gbps	Optical fiber	Backbone

13

Best Practices include:

- If using IP as the backbone, do a study of the customer environment and figure out the load on the pipe and whether the existing infrastructure will be able to handle the data.
- Compare the cost and feasibility of using IP/ATM/DWDM as the connectivity media.

Flow Control – Basics

- A variety of mechanisms designed into hardware components of the network to:
 - Increase the distance between two end-points in a storage network
 - Optimize the round-trip traffic between a local and remote data center
 - Use compression to reduce bandwidth consumption
 - Provide network protocol optimizations such as additional inline buffer credits to extend the reach of a Fibre Channel SAN, without reducing the line rate

Flow Control – Buffer Credits

■ Buffer Credit Usage

- Each port must have a buffer available for each Fibre Channel frame that is sent across the cable or fibre.
- The time taken to send a frame across increases as distances increase.
- In general, one buffer credit is required to transmit 2KB frames across 2km to sustain 1Gb/sec of throughput.
- This translates to five buffer credits required to transmit 2KB frames across 10km to sustain 1Gb/sec throughput.

15

Each FC device has a buffer. The size determines how many non-stop frames can be sent to that device. Each FC device tells other FC devices how large a buffer it has. When received frames are processed and moved out of the buffer space, the receiving device tells other devices it has available buffer space again. The term for this is buffer credits.


The following tables summarize buffer credits required to sustain 1 Gb/sec and 2Gb/sec throughputs over varying distances.

For sustained throughput of 1Gb/sec:

<u>Buffer Credits</u>	<u>Distance</u>
1 Credit	2km
5 Credits	10km
25 Credits	50km
50 Credits	100km
500 Credits	1,000km

For sustained throughput of 2Gb/sec:

<u>Buffer Credits</u>	<u>Distance</u>
1 Credit	1km
5 Credits	5km
10 Credits	10km
50 Credits	50km
100 Credits	100km
1,000 Credits	1,000km



SAN Extension Technology: Dark Fiber

- Native Fibre Channel sent directly over dark fiber infrastructure (unused fiber laid by Telcos)
- Eliminates the need for additional media and transmission conversion equipment
- Distance dependent on:
 - Type of fiber deployed
 - Longwave Single Mode – 10km
 - Shortwave Multi-mode – 500m
 - Use of buffer-to-buffer credits to extend the connection.
- Fast, low cost deployment

17

SAN Extension Technology: SONET / SDH

- Synchronous Optical Network (SONET)
- Synchronous Digital Hierarchy (SDH)
- Now supports Ethernet and DWDM
- Guaranteed performance and quality of service (QoS)
- Performance Issues:
 - 20–25 microseconds delay at each transmission node.
 - Each additional intermediate through-node introduces another 10-microsecond delay.
 - Supported distance limit is about 300km.

18

The maximum allowable delay for synchronous replication tasks is generally considered to be one millisecond – 500 microseconds each way. Considering that fiber propagation introduces an approximate five microsecond delay for each 0.6 miles (1km) traveled, the aggregate delay of processing Fibre Channel over SONET/SDH can quickly add up, limiting the distance between the source and target of a replication process to around 100km.

SAN Extension Technology: CWDM


- Coarse Wavelength Division Multiplexing (CWDM)
 - Uses existing dark fiber
 - Uses Optical Add/Drop Multiplexer filters at each end of the network
 - Multiplex and de-multiplex multiple CWDM wavelengths
 - Provides up to eight channels of traffic over a single fiber pair
 - Carries Fibre Channel, FICON, ESCON, others
 - Attenuation reduces the distance between end points to 66km.

SAN Extension Technology: DWDM

- Dense Wavelength Division Multiplexing (DWDM)
 - Optical networking solutions use DWDM technology and supply the infrastructure for high availability of both synchronous and asynchronous data mirroring.
 - DWDM provides a higher concentration of fiber and increased distance with a single-mode connection.
 - It does not have to be dedicated solely to the disk mirroring application.
 - Protocol independent:
 - Supports Fibre Channel, ESCON, FICON, Gigabit Ethernet over FC, others
 - Write acceleration feature allows the extension of synchronous replication distance limitation.

20

■ Overview of DWDM



- Merges *optical* traffic onto one common fiber
- Allows high flexibility in expanding bandwidth
- Reduces costly mux/demux function, reuses **existing optical signals**
- Affected by signal loss (attenuation)

21

DWDM is a technology that puts data from different sources together on an *Optical Fibre*, with each signal carried at the same time on its own separate light wavelength. In addition DWDM provides:

Separate wavelengths or channels of data can be multiplexed into a light stream transmitted on a single optical fiber.

Each channel is demultiplexed at the end of the transmission back into the original source. Different data formats being transmitted at different data rates can be transmitted together.

Signal Loss

Most general-purpose optical fiber being installed today exhibits loss of 4 to 6dB per km (a 60% to 75% loss per km) at a wavelength of 850nm. When the wavelength is changed to 1300nm, the loss drops to about 3 to 4dB (50% to 60%) per km. At 1550nm, it is even lower. Premium fibers are available with loss figures of 3dB (50%) per km at 850nm and 1dB (20%) per km at 1300nm. Losses of 0.5dB (10%) per km at 1550nm are not uncommon. These losses are primarily the result of random scattering of light and absorption by actual impurities in the glass.

The implication is that the greater the loss the shorter the fiber optic ring will have to be. New fiber optic technologies are being developed to reduce the loss incurred by the fiber itself.

- DWDM

- Deployment in Geographically Dispersed SANs (GeoSANs)
 - Performance
 - Current DWDM devices employ Write Acceleration, buffering crediting, caching and other similar features.
 - Extends Synchronous replication distances
 - Improves performance
 - Synchronous versus asynchronous replication differences become important as the distance increases (link latency increase).
 - Write acceleration helps mitigate latency increase.
 - Provides Guaranteed Quality of Service:
 - Bandwidth allocated to each channel is fixed, not shared.
 - As additional data, ports or wavelengths are added the bandwidth allocated to each port stays constant.

22

Synchronous versus Asynchronous Replication

There is a distinct difference in the performance between synchronous and asynchronous applications, that is, whether or not the application must wait for an acknowledgement to a data transfer command before resuming processing. Even though optical networks can offer transmission at wire speeds, when traversing tens of miles even at the speed of light, propagation delay and device turn-around-time become the new limiting factors.

SAN Extension Technology: FCIP

- FCIP: Fibre Channel over IP – “Tunnels” Fibre Channel over IP-based Networks
 - Allows interconnection of Fibre Channel SANs with TCP/IP used as the underlying wide-area transport over LANs, MANs, and WANs.
 - IP provides congestion control and in-order delivery of data.
- Operation:
 - Frames originating on one SAN are wrapped in IP packets and forwarded to the destination SAN.
 - At the receiving end, the IP header is removed and native Fibre Channel frames are delivered to the fabric.
 - The only devices that *talk* IP are the FCIP gateways.
 - FCIP tunneling requires both IP and Fibre Channel management applications.

23

FCIP


Fibre Channel SANs can be interconnected to meet the needs for remote storage access. However, by combining IP networking with SAN technology you can extend the interconnectivity of the SANs across much longer distances. FCIP provides the transport for traffic going between specific Fibre Channel SANs over LANs, MANs, and WANs.

FCIP is used to tunnel Fibre Channel traffic between two geographically separate Fibre Channel SANs. Frames originating on one SAN are wrapped in the IP packets and forwarded to the destination SAN. At the receiving end the IP header is removed and native Fibre Channel frames are delivered to the fabric. A Fibre Channel fabric switch then makes the decision about which end device the frame is intended for. In terms of discovery, the only devices that have IP addresses are the FCIP gateways themselves. IP discovery is thus limited to the FCIP gateways, while Fibre Channel discovery and management is still required for the storage end devices. Since FCIP tunneling requires both IP and Fibre Channel management applications, additional overhead is necessary for a tunneled solution.

■ FCIP Protocol

– Transparent Operation for Local and Remote SANs

- FCIP gateways (converters) are the only devices that need to be aware of FCIP encapsulation.
- It appears like Fibre Channel to the SAN, and IP to the Local / Metro / Wide Area network.
- It allows usage of high-performance backbone-quality IP pipes.




24

This protocol is entirely transparent for existent Fibre Channel SANs and involves usage of the infrastructure of a modern MAN/WAN network. Some application problems that can be successfully solved using the FCIP protocol are remote backup, data recovery, and a shared data access. With high-speed MAN/WAN communications, you can also use synchronous data doubling and a shared distributed access to data storage systems.

■ FCIP

- Latest FCIP technology does not allow Class F admin traffic across IP Network.



25

Principal Switch considerations

- One principal switch per Fibre Channel fabric
- In a multi-site, single Fibre Channel fabric deployment and the principal switch can be only at one site at any moment

■ FCIP Review

– Considerations for IP transport

- Multiple, geographically separated storage systems that need to be connected for remote copy applications
- Cost concern for DWDM, other network options
- Preference of IP over Fibre Channel for backbone
- Mature technology – confidence in operation
- Remote site data vaulting, data recovery, backup
- Requires T3 or faster IP network
- Switch administration traffic issues fixed in latest generation technology

26

SAN Extension Technology: Comparison

■ Comparison of the Commonly Available Transport Media

Connectivity	Pros	Cons	Used for
Dark Fiber	- Bandwidth - Highest Quality of Service	- Cost - Availability - Complexity	SAN extension and Synchronous replication
DWDM	- Bandwidth - High Quality of Service	Cost	SAN extension and Synchronous replication
Optical Carrier Networks	- Quality (packet loss and latency) - Availability - Cost efficiency		Could be shared or dedicated for different network services
Ethernet (IP) Networks	- Lowest cost - Shared with other data services - Highest availability	- Require Fibre protocol conversion - Highest protocol overhead - Latency jitter due to routing	Most widely used network services

27

Module Review

1. What are the basic variations of Universal Replicator Fibre Channel connection topologies?
2. Describe how FCIP handles Fibre Channel frames.
3. Describe DWDM transport medium.


11. Hitachi Replication Manager Overview

Module Objectives

- Upon completion of this module the learner should be able to:
 - Navigate the Replication Manager User Interface
 - Discuss the features of Replication Manager
 - Describe Replication Manager operations
 - Explain Replication Manager Pair Configuration Wizard and concept of Task Management
 - Describe Change Status Wizard
 - Describe Replication Manager replication monitoring functions
 - Configure Alerts and Event Logs
 - Discuss the application support feature

2

Overview of Replication Manager User Interface



Replication Manager GUI is consistent with other Hitachi Storage Command Suite products

- Explorer Navigation Drawers
 - ♦ Much simpler
- Consistent Navigation tree
- “Bread Crumb” Trail for browsing resources
- Dashboard Menu
 - ♦ Displays a list of Storage Command Suite products on the same management server.
 - ♦ Products can be launched using the Go link.
- Application Area
 - ♦ Displays information for the option selected in Explorer menu

- Replication Manager Multiple drawers with options

- **Resources** Drawer

- Displays application servers or hosts, subsystems, and Pair Management servers.

- **Shared Views** Drawer

- Displays Site information
 - Sites information allows you to group resources, such as hosts and subsystems, by purpose so that you can manage resources efficiently.

- **My Groups** Drawer

- Displays My Copy Groups information
 - Copy Groups display the configuration and status of Copy Pairs in a single window.
 - This can be used to group Copy Pairs to be monitored by a user.

4

- Multiple drawers with options (continued)

- **Administration** Drawer

- Used for administrative purposes such as creating users, assigning permissions, setting passwords, creating and managing Resource groups, and managing event logs.

- **Tasks** Drawer

- Displays the scheduled tasks and generated workflows for Copy Pair settings.

- **Alerts** Drawer


- Allows you to check alerts and edit alert settings.

- **Settings** Drawer

- Check and edit user profile.
 - Set Replication Manager software license information.
 - Refresh the configuration to update the information managed by Replication Manager software and set the intervals for updating the information.

5

- Four types of views allows users to understand replication environment depending on the perspective:
 - Hosts View** — Perspective of hosts using the pairs.
 - Subsystems View** — Perspective of subsystems containing the pairs.
 - Pair Configurations View** — Perspective of hosts managing the pairs.
 - Applications View** — Perspective of applications (MS Exchange) being managed.



Replication Manager provides the following four functional views that allow you to visualize pair configurations and status of the replication environment from different perspectives:

- Hosts view:** A view of hosts using the pairs. This view lists open hosts and mainframe hosts and allows you to confirm pair status summaries for each host.
- Subsystems view:** A view of subsystems containing the pairs. This view lists open and mainframe subsystems and allows you to confirm pair status summarized for each subsystem. A subsystem serving both mainframe and open system pairs is recognized as two different resources in order to differentiate open copy pairs and mainframe copy pairs.
- Pair Configurations view:** A view of hosts managing the pairs. This view lists open and mainframe hosts managing copy pairs with CCI or BCM and allows you to confirm pair status summarized for each host. This view also provides a tree structure along with the pair management structure.
- Applications view:** A view of applications (MS Exchange servers) being managed. This view lists the application and data protection status. This view also provides a tree structure along with the servers and their associated objects (Storage Groups, Information Stores and Mount Points).

■ Consistent Interface


- Support for both open and mainframe volumes
- Same interface for different replication technologies and storage subsystems

The **Subsystem** view displays both mainframe and open devices. In this way, permissions to manage each are maintained independently.

Subsystem Summary									
Subsystem Type		USP_V							
Serial Number		45159							
Cache Size		16.00 GB							
Platforms									
Platform	Licensed Copy Type						RAW Capacity	Allocated Capacity	Information Source
	SI	TCA / TCE / SDR	TCS	QS / COW	UR				
Open	✓	✓	✓	✓	✓	✓	1.07 TB	16.00 GB	local HDvM

Features

- “Single Pane of Glass”
 - Integrated console for multisite replication pairs
 - Consolidated monitoring for Copy Pair status and remote copy metrics
- Visual Representation of Replication Structure
 - Copy groups, sites, and all volume pairs



8

Visual Representation of Replication Structure


Copy Groups: A group of Copy Pairs created for management purposes, as required by a particular task or job. By specifying a Copy Group, you can perform operations such as changing the Pair Status of multiple Copy Pairs at once.

Using My Copy Groups feature, a user can register a Copy Group into My Copy Groups that are most important to monitor and see how Copy Groups are related and check Copy Pair statuses in a single window. My Copy Groups is also the default screen after you login to the Hitachi Replication Manager interface.

Sites: With Hitachi Replication Manager, you can define logical sites in the GUI just as you would define physical sites (actual data centers). It allows you to manage resources more efficiently if you setup separate sites because it is easier to locate a required resource among many resources displayed in the GUI.

Features: Mainframe

- Allows mainframe pair management
 - Copy Group configuration and the Pair Status information has been returned to Replication Manager software from the Mainframe Agent.



9

Add the Mainframe Agent as an Information Source to Replication Manager

- The Mainframe Agent executes as a *service*, waiting for a pooling event by Replication Manager.
- The linkage for Device Manager is not necessary for this usage.
- In order to set up the polling operation, the Mainframe Agent must be added as an information source to Replication Manager.
- To add as an Information Source:
 1. Select the **Administration** drawer.
 2. Click **Information Source** option.
 3. Select **BC Manager/Mainframe Agent** in the Navigation area.
 4. Add the IP address, port number, and the name of the Mainframe Agent.

Hitachi Replication Manager

Logged in as: system Close Logout

Explorer

- Resources
- Shared Views
- Sites**
- My Groups
- Administration
- Tasks
- Alerts
- Settings

Dashboard

- HDvM GO
- HPvM GO
- HTvM
- HGLM
- HRvM
- Mode: Normal
- 0 C/T 0
- 0 SF 0
- 0 JNL 0
- 0 Pool 0
- 0 JNLG 0
- HTSM
- HSSM

Sites

- PrimarySite
 - Hosts
 - 2950-acad-02
 - VM-HCMD-LAB-02
 - Subsystems
 - USP_VM@172.17.41.74
 - Pair Configurations
 - 2950-acad-02
 - VM-HCMD-LAB-02
 - Remote Site
 - Hosts
 - 2950-acad-03
 - VM-HCMD-LAB-04
 - Subsystems
 - USP_VM@172.17.41.74
 - Pair Configurations
 - 2950-acad-03
 - VM-HCMD-LAB-04

Hosts

Sites > PrimarySite > Hosts

Rows/ Page: 100 Page 1 of 1

Host	IP Address	Information Source
2950-acad-02	172.17.250.14	local HDvM
VM-HCMD-LAB-02	172.17.250.37	local HDvM

Navigation Area - shows site structure and resources

Application Area - shows the same information as Resources drawer

Refresh Tree Add Hosts Remove Hosts Refresh Hosts

Trusted sites

10

Features: Manage Security

- Replication Manager software provides the following security functions:
 - Sets password policy to prevent users from specifying easy-to-guess passwords
 - Enables automatic locking of user accounts if successive login attempts fail
 - Displays a warning banner in the user login window



11

■ Check for Replication Technology Licenses

The screenshot shows the Hitachi Replication Manager web interface. The left sidebar has a 'Resources' drawer expanded, showing 'Subsystems'. The main content area is titled 'Open' and shows details for a subsystem. The 'Copy Licenses' tab is selected, displaying a table of licenses.

Copy Type	Product Name	Key Type	Term	Used Capacity	Licensed Capacity
SI	ShadowImage	Permanent	-	21.00 GB	100.00 TB
TCA / TCE / SDR	TrueCopy Asynchronous	Permanent	-	0.00 GB	100.00 TB
TCS	TrueCopy	Permanent	-	0.00 GB	100.00 TB
QS / COW	Copy-on-Write Snapshot	Permanent	-	21.00 GB	100.00 TB
UR	Universal Replicator	Permanent	-	0.00 GB	100.00 TB

12

1. In **Explorer** menu, click **Resources** drawer and select **Subsystems**. Subsystems information is displayed in the Navigation area.
2. Expand a subsystem to open subsystem information and click **Open**. The subsystem detailed information is displayed in the Application area.
3. Click **Copy Licenses** tab and verify if license keys are installed for the replication technology to be used.

Operations: Create Command Device

The screenshot displays the Hitachi Replication Manager (HRpM) interface. In the foreground, the 'Add Command Devices - Microsoft Internet Explorer' window is open, showing the 'Add Command Devices Wizard' with steps: 1. Select Devices, 2. Confirm, 3. Finish. The 'Candidate Devices' table lists various LDEVs with their capacities and parity groups. The 'Command Devices' table shows newly configured command devices. In the background, the 'Open Subsystem Summary' window is visible, showing details for the subsystem USP_V@172.17.41.74 (10071). A red box highlights the 'Cmd Devs' tab in the background window, and a red arrow points from this tab to the 'Add Command Devices' button in the foreground wizard.

Open Subsystem Summary

RAW Capacity: 12.84 TB
Allocated Capacity: 71.01 GB

LUNs	Cmd Devs	Pools	JNLGs	Copy Licenses
0:28	Disable	1-4-1	0	
0:41	Disable	1-4-1	0	
0:42	Disable	1-4-1	0	
0:43	Disable	1-4-1	0	
0:44	Disable	1-4-1	0	
0:61	Disable	1-1-1	0	
0:64	Disable	1-1-1	0	
0:65	Disable	1-1-1	0	
0:67	Disable	1-1-1	0	
0:AD	Disable	E1-1	0	

Add Command Devices - Microsoft Internet Explorer

HRpM - Add Command Devices Wizard

1. Select Devices 2. Confirm 3. Finish

Select devices you want to change to command devices. Perform following to select devices:
1. Select filtering parameters ("CPU" and "Capacity"), and click "Refresh".
2. On "Candidate Devices", mark the checkbox beside volumes that you want to set command device, and click "Add".
For canceling setting of command devices, perform following:
On "Command Devices", mark the checkbox beside devices that you want to unregister, and click "<Delete".
Click "Next" when the selection has been finished.

Filters for Candidate Devices: CPU: All Capacity: All Refresh

LDEV	Capacity	Parity Group	SLPR	CLP
0:00	1.00 TB	X1-1	0	
0:01	1.00 TB	X2-1	0	
0:10	100.00 GB	1-7-1	0	
0:11	100.00 GB	1-5-1	0	

Command Devices

Newly configured command devices


LDEV	Capacity	Parity Group	SLPR
0:63	2.00 GB	1-1-1	0

13

Any LDEV not in use by a host or other Universal Storage Platform can be a command device.

- To add Journal Groups

- In the **Explorer** menu, click the **Resources** drawer and select **Subsystems**. The Subsystems information is displayed in the Navigation area.
- Expand a subsystem to open up the subsystem information and then click **Open**. The subsystem detailed information is displayed in the Application area.
- Click the **JNLGs** tab and then click **Create JNLG** on the bottom of the window. The Create Journal Group Wizard starts.



Operations: Configure Remote Paths

Remote Path Configuration

The screenshot shows the 'Open Subsystem Summary' window for the subsystem 'USP@172.17.41.166 (10079)'. The 'Remote Paths' tab is selected and highlighted with a red box. The table below lists the configured remote paths.

	Remote Subsystem	MCU	RCU	Label	Path Gr. ID/SSID	Status	Num of Paths
☐	USP_V@172.17.41.64 (10069)	Any	Any	Production	Default	Normal	1
☐	USP_V@172.17.41.64 (10069)	0B	35	Test	B435	Normal	1
☐	USP_V@172.17.41.64 (10069)	31	31		B431	Normal	1
☐	USP_V@172.17.41.64 (10069)	33	33		B433	Normal	1
☐	USP_VM@172.17.41.105 (25033)	Any	Any		Default	Normal	1
☐	USP_VM@172.17.41.105 (25033)	Any	Any		*30	Normal	1
☐	USP_VM@172.17.41.105 (25033)	32	30		0017, 0018	Warning	1
☐	USP_VM@172.17.41.105 (25033)	33	31		F531	Normal	1


- To Add Pools, perform the following steps:
 - In the **Explorer** menu, click the **Resources** drawer and select **Subsystems**. The Subsystems information is displayed in the Navigation area.
 - Expand a subsystem to open up the subsystem information and click **Open**. The subsystem detailed information is displayed in the Application area.
 - Click the **Pools** tab and click **Create Pool** on the bottom of the window. The Create Pool Wizard starts.

The screenshot shows the 'Open Subsystem Summary' window for the subsystem 'USP_V@10.17.102.180 (45155)'. The 'Pools' tab is selected and highlighted with a red box. The table below lists the configured pools.


	Pool	Used by	Status	Pool Usage	Total Capacity	Number of Associated Pairs	Last
☐	0	QS / COW	Normal	0%	219.75 GB	0	2009

Operations: V-VOL Configuration Support

- Replication Manager software provides a screen to display and configure V-VOLs for Universal Storage Platform V and VM.
- It allows centrally managing V-VOLs in multiple storage system environments.




V-VOL configuration



Use Device Manager to map created V-VOLs to hosts.

- Provides a wizard with a graphical view to define the pair structure.
- Pair Configuration Wizard
 - You can create the replication topologies and manage the assignment and removal of pairs to Copy Groups for replication. The wizard assists by using:
 - Topological view
 - Provides intuitive image of the relationships between Copy Groups
 - Allows specification of complex pair structures easily
 - Filtering Volumes
 - Site, hosts, storage systems, and logical groups can be used as filtering attributes to find candidate volumes.



Candidate List

Criteria Results

Volume Type: Secondary

Select Volume:

Site: Others

Subsystem: USP_V@172.17.41.64 (10069)

Filter By:

Host Name: All

Logical Group:


Label:

Apply Clear

Pair Configuration Wizard: Task Management

- In the **Task List** section:
 - Each task corresponds to one Copy Group.
 - For each task, **Pair Operations** are displayed.
 - Modify file
 - Edits the configuration definition file, but does not create a Copy Pair
 - Modify file + create
 - Edits the configuration definition file and creates a Copy Pair
 - Modify file + delete
 - Edits the configuration definition file and dissolves a Copy Pair
 - The default Execution Request time is set to **Immediately**.
 - You can change it by using the Edit icon of the displayed task.
- Click Edit icon of the existing task to open Edit Task dialog box.

- The **Task Options** tab displays the following **Execution Request Time** options:
 - Execute Immediately
 - The task starts when the Pair Configuration Wizard finishes.
 - Execution Date
 - Scheduled execution, specify execution date and time.
 - Modify Pair Configuration File Only (Do not create pair)
 - Only modifies the configuration file.



20

Modify Pair Configuration file Only


- This option creates the *Simplex pairs* (just the configuration file and no actual pairs on the subsystem).
- Replication Manager software allows you to create pairs from the *Simplex status* using the Change Status Pair Wizard.

Advantage of using this option:

- You do not have to manually edit the HORCM files.
- Backward compatibility with RAID Manager CCI.

Configuration: Task Management — Workflows

- To complete a task at a later date, save as a Workflow.
- When all the corresponding tasks are complete, a workflow is automatically deleted.



21

Each screen of the Pair Configuration Wizard has a **Save** button to launch the Save Workflow dialog box.

The workflow allows a user to save the data, for example, matching the pairs that the user is working on for continuation at a later time. If the user needs to create multiple pairs but is interrupted before submitting the configuration task, the user can save the configuration as a workflow. The volumes that the user has selected and saved in the workflow are noted in the GUI and are reserved and another user is prevented from selecting the same volumes.


Saved Workflow

Displays a workflow previously saved. To save a workflow, use the **Save** button on the Pair Configuration Wizard.

Completed Workflow

A Completed Workflow corresponds to the completed wizard.


- Edit Workflow
 - Opens the Pair Configuration Wizard from the saved screen.
 - Only the saved workflow can be edited (completed workflow cannot be edited).
- Delete Workflow
 - All corresponding tasks should be completed or already deleted before deleting workflow.



22

Change Pair Status Wizard

- Change the status of each Copy Pair individually, multiple Copy Pairs at once, or select a Copy Group to change the pair status.
- Launch the Change Pair Status Wizard in multiple ways:
 - In the **Resources** drawer, select either **Hosts** or **Pair Configurations**.
 - In the Shared Views drawer, select **Sites** and then **Host** or **Pair Configurations**.




- Operations for Changing the Copy Pair Status
 - Create
 - Creates a Copy Pair based on the configuration definition file (the initial copying of a Copy Pair in simplex status is performed)
 - Split
 - Splits a Copy Pair
 - Resync
 - Performs data resynchronization from the primary volume to the secondary volume
 - Restore
 - Performs data resynchronization from the secondary volume to the primary volume
 - Delete
 - Dissolves a Copy Pair (the Copy Pair is placed in simplex status)

Monitoring: My Copy Groups

- Replication Manager software displays the volume associations, configurations, and Copy Pair status of the Copy Groups selected by the storage administrator.
- Using My Copy Groups, users can:
 - Monitor multiple Copy Groups in one window
 - Focus on specific Copy Groups when monitoring Copy Groups
 - Register multiple Copy Groups to My Copy Groups
 - Check detailed status information of each Copy Group
 - Refresh Pair Status of all the displayed Copy Groups
 - Check details of a specific Copy Group by clicking the group's link
- You can register a maximum of 100 Copy Groups.

- My Copy Groups allows users to understand replication configuration and status with topological view.
- The summarized pair status for each Copy Group enables users to grasp the replication status.



Monitoring: Refresh Configuration and Status

- Configuration Refresh: From Device Manager database and agents
 - 1. Configuration Information
 - Volumes
 - Copy Pairs and Copy Groups
 - Command Devices
 - DMLUs
 - Journal Volumes
 - Pool Volumes
 - Remote Paths
 - Copy Type licenses (including usage)

27

- Configuration Refresh: From Device Manager database and agents (continued)
 - 2. Copy Pair Status Information
 - Copy Pair status and summarized Copy Pair status
 - Copy Pair states
 - Copy progress
 - Performance information
 - Consistency Time (C/T) delta values and buffer usage on Copy Group basis*
 - Buffer usage (Sidefiles, Journal Groups, and Pools)
- The configuration information and Copy Pair status information can be refreshed manually or automatically at a preset interval.
 - If you change the replication configuration, make sure that you refresh the configuration information before you refresh the Copy Pair status.


28

* Consistency Time (C/T) Delta gives you the time required to write data of the primary volume to the secondary volume based on the write delay time.

- When refreshing the Copy Pair status, the sidefile usage on a Copy Group basis and journal volume usage on a Copy Group basis are included as target buffer usages of the refresh operation.
- The following information is refreshed automatically every five minutes, regardless of the type of refresh operation that has been set (manual or automatic):
 - ♦ Journal group status
 - ♦ Journal volume usage for each Journal Group
 - ♦ Pool status
 - ♦ Pool volume usage for each Pool

Monitoring: Refreshing Configuration Information

- Configuration information is acquired from the databases of Device Manager servers registered as information sources and from the Device Manager agents.
- When Device Manager servers are installed in multiple sites and are used as information sources, configuration information is acquired by using the Device Manager agents for each of those sites.



29

You can add the Device Manager servers installed on the management servers at remote sites or instances of Business Continuity Manager or Mainframe Agent installed on the hosts of sites (including the local site) as information sources. You can register a total of 100 Device Manager servers, instances of Business Continuity Manager, and instances of Mainframe Agent.

You do not need to register the management server at the local site on which Replication Manager is installed (the Device Manager server) as an information source. Because this Device Manager server is automatically registered during installation, it is displayed as local Hitachi Device Manager in the list of information sources.

- The following information is automatically refreshed every 5 minutes, regardless of the refresh settings for configuration information:
 - Journal group status
 - Journal volume usage on a journal group basis
 - Pool status
 - Pool volume usage on a pool basis

Alerts

- Alert notification is useful for enabling a quick response to a hardware failure or for determining the cause of degradation in transfer performance.
- In Replication Manager software, you can set conditions such as the Copy Pair status and performance thresholds for each Copy Group, Journal Group, sidefile, C/T Delta (Write Delay Time) or Pool for preventing buffer overflow errors and insufficient copy licenses, thereby facilitating the continuity of normal operation.
 - A maximum of 1000 conditions can be created.
 - Alerts can be issued by way of email or SNMP which allows monitoring of the replication environment even after logging out of Replication Manager software.
 - The best practice is to set the alert settings before starting the Replication Manager operations, either during initial setup or during reconfiguration.

31

- Set detailed information about the alert, such as a monitoring condition or an action for performance monitoring:

- **Alert Name**

- **Monitoring Performance Journal per Journal Group > %**

- Select performance information used as an alert monitoring condition, and then specify the threshold for monitored target.
- Performance information that you can monitor depends on resource and copy type of monitored target.

HRpM - Create Alert Setting Wizard

1. Introduction ▶ 2. Select Monitoring Type ▶ 3. Select Alert Setting ▶ **4. Edit Alert Action** ▶ 5. Confirm ▶ 6. Finish

Edit an alert action in the alert settings.
Perform either of the following to edit an alert action:
a. Edit an alert action by using a template:
1. If you have already created a template, select an available template, and then click "Apply Alert Setting". Then, you can edit the parameters.
2. Click "Next" to confirm a new alert setting.
b. Edit an alert action without a template:
1. Select the copy pair status that you want to monitor with Replication Manager.
2. If you want to receive an alert by email, select "Alert by Email", and then enter the required parameters.
3. If you want to receive an alert by an SNMP trap, select "Alert by SNMP Trap", and then enter the required parameters.
4. Click "Next" to confirm a new alert setting.
If you want to save the modified parameters as a template, click "Save Alert Setting as Template" before clicking "OK".

*Alert Name: Journal overview

Available Template: [Select] [Apply Alert Setting]

*Monitoring Performance Journal per Journal Group > 70 %

☒ Alert with Email

Mail Server: [Select]

To: Storage_Admin

From: Replication Manager

Subject: Replication Manager Alert

Message: The Journal Group usage reached the threshold - 70%

☐ Enable SNMP Authentication

User ID: [Text]

Password: [Text]

☐ Alert with SNMP Trap

Host: [Text]

Community: public

Message: [Text]

- **Alert with Email or Alert with SNMP Trap**

32


Performance information that you can monitor depends on the resource type and copy type of the monitored target. Following is the list of options:

- **Sidefile per Copy Group** — This option monitors the sidefile usage for each Copy Group. Enter a value in the range from 0 to 99 to specify the threshold (%) for the monitored target. You can select option when the resource type of the monitored target is a Copy Group type and the Copy Type is TrueCopy Asynchronous.
- **Journal per Copy Group** — This option monitors the Journal Group usage for each Copy Group. Enter a value in the range from 0 to 99 to specify the threshold (%) for the monitored target. You can select this option when the resource type of the monitored target is a Copy Group type and the Copy Type is Universal Replicator.
- **C/T Delta per Copy Group** — This option monitors the C/T delta value for each Copy Group. Enter a value in the range from 0 to 8,640,000 to specify the threshold (seconds) for the monitored target. You can select this option when the resource type of the monitored target is a Copy Group type and the Copy Type is Hitachi TrueCopy Asynchronous, Hitachi TrueCopy Extended Distance, or Hitachi Universal Replicator.
- **Journal per Journal Group** — This option monitors the journal volume usage for each Journal Group. Enter a value in the range from 0 to 99 to specify the

threshold (%) for the monitored target. You can select this option when the resource type of the monitored target is an open system Journal Group type.

- **Pool per Pool** – This option monitors the Pool usage for each Pool. Enter a value in the range from 0 to 99 to specify the threshold (%) for the monitored target. You can select this option when the resource type of the monitored target is an open system Pool type.

■ Set Alerts – Email or SNMP trap.



■ Examples

- Journal Group status and thresholds
- Copy Pair status changes
- Replication License capacity usage

33

In the diagram and the following pages:

HRpM stands for Replication Manager software

HDvM stands for Device Manager software

BCM stands for Business Continuity Manager software

M/F stands for Mainframe


Mgt stands for Management

- All registered alerts can be managed from the Alerts view.
 - Click **Alerts** drawer in **Explorer** menu.
 - Click **Alert Setting List** tab to view registered alerts.

Alerts							
Alerts							
Alert Summary							
Marking Type Auto							
Alert List							
Alert Setting List							
Rows/Pages: 100 of 1							
Date / Time	Alert Name	Resource	Completed/ Not Completed	Alert Action		Detected	
				Email	SNMP	Metrics	Value
2010-04-11 17:59:08	shadowimage_alert	EXG_GEN0	Not Completed	✓		Pair Status	II
2010-04-11 17:54:09	shadowimage_alert	EXG_GEN0	Completed	✓		Pair Status	✓
2010-03-17 18:58:40	TCSAlert	TCSGroup	Completed			Pair Status	✓
2010-03-17 18:23:11	TCSAlert	TCSGroup	Completed			Pair Status	✓
2010-03-17 18:15:13	TCSAlert	TCSGroup	Completed			Pair Status	✓
2010-03-17 17:58:56	TCSAlert	TCSGroup	Completed			Pair Status	✓
2010-03-17 17:55:39	TCSAlert	TCSGroup	Completed			Pair Status	✓
2010-03-17 16:42:53	TCSAlert	TCSGroup	Completed			Pair Status	✓
2010-03-17 15:47:09	TCSAlert	TCSGroup	Completed			Pair Status	✓

Export Alerts

■ Example of Email Alert



"DETECTED PAIRS" shows individual pair information for a copy group.


35

Detected Pairs is displayed only when an alert is generated for a copy group and the alert Automarking feature is enabled (**Marking Type: Auto**).

To limit the volume of information when an enormous number of pairs is involved, the display is limited to ten pairs along with the following message:

More than 10 pairs were detected.

- Every time Replication Manager software notifies an alert, a new item is added on the list with **Unmarked** status and it stops notifying the same alert.
- After solving a problem that causes the error or unexpected status, select alert and click **Mark as Reviewed**.
 - Replication Manager software resumes monitoring for the same alert.



36

Managing Alerts: Metrics for Remote Copy

- Consistency Time Delta
 - To facilitate the evaluation and forecast of service level for remote replication, Replication Manager provides C/T Delta per Copy Group.
 - By checking this metric, you can determine how far behind the remote volume is from the primary volume.
 - Create an alert that triggers if the C/T Delta per Copy Group > # of seconds.

The screenshot shows the Hitachi Replication Manager web interface. On the left is a tree view of 'Pair Configurations' including '2950-acad-02', '2950-acad-03', and 'VM-HCMD-LAB-02'. The main panel displays the 'Copy Group Summary' for '2950-acad-03 > horcm451.conf > URx9'. The 'C/T Delta' is 0sec, highlighted with a red box. Below this, 'Sidefile / Journal Usage' is 0%. The 'Primary' server is '2950-acad-02' with config file 'horcm450.conf'. The 'Secondary' server is '2950-acad-03' with config file 'horcm451.conf'. A 'Pairs' table at the bottom shows one pair with ID 'P_028-035_0077', type 'UR', and status 'Active'.

C/T Delta		SI
0sec		SI

0%	TCA / TCE / SDR
----	-----------------


Server	2950-acad-02	TCS
Config File	horcm450.conf	QS / COW

Server	2950-acad-03	UR
Config File	horcm451.conf	Visible / Total
		1 / 1

Copy Pair	Copy Type	Pair Status	Copy Progress
P_028-035_0077	UR	Active	Inactive

37


- To Export the History of C/T Delta
 - Click **Export History**. In the Export History dialog box, for **Target Metrics** select the **C/T Delta per Copy Group**.
 - Provide the export range and the format.
 - This gives you the time delta between two sites, and you can use this historical report for RPO.



38

Managing Events: Event Logs

- To View Events
 - In **Explorer** menu, open **Administration** drawer and select **Event Logs**.
 - All the events are displayed in Application area.
 - Sort using the **Severity** column to focus on error messages.



39

Event Log

Date / Time - Displays the event occurrence time

Severity - The following are the levels of severity:


- Error
- Warning
- Information

Message ID - Displays the message ID of the event

Message - Displays the event message

In an event message, the message character string varies, depending on names or values included in the message. If the length of this variable character string exceeds 255bytes (UTF-8 conversion), only the first 255 bytes are displayed.

- Click **Export** in the Application area to export the event logs to HTML or CSV file.
 - Specify a time period for the information to be exported and the export format.
 - The event log data is exported to a file.



40

An exported log file may be requested by Hitachi Global Support Center (GSC).

Error Message Example

- Error Generated in the Events List
 - **KAVN01216-E** — An attempt to read installed product version information from the registry has failed.
 - Open the *Hitachi Replication Manager Messages* manual and search for the error message ID.
 - Results of search
 - Cause: Device Manager software might not have been correctly installed.
 - Action: Reinstall Device Manager software and then install Replication Manager software.

- An **application replica** is a snapshot of an application server that is saved to a series of secondary volumes on an immediate or scheduled basis.
- Requires installation of **Application Agent** on the application server
- Supported **Applications**
 - MS-Exchange
 - MS-SQL
- Replication Manager is used to create, manage, and (manually) restore replicas.
 - Support for auto and manual selection of backup targets.
 - Targets can be rotated.

Application Support

■ Applications view

The screenshot displays the Hitachi Replication Manager web interface. The left sidebar contains a tree view under 'Applications' with nodes for 'Exchange', '172.17.72.35', 'SQL Server', and 'ClusteredSQL'. The main pane shows details for host '172.17.72.35', including a 'Host Summary' section and a 'Storage Groups' table. The 'Storage Groups' table lists 'DF800E STORAGE GROUP' and 'SECOND STORAGE GROUP' with their respective 'Status' and 'Latest Replica' information. Annotations with arrows point to the application tree and the storage groups table.

Hitachi Replication Manager

Logged in as: system Close Logout

File Go Help

Explorer

Resources

Hosts

Subsystems

Pair Configurations

Applications

Shared Views

My Groups

Administration

Tasks

Alerts

Settings

Dashboard

HDvM GO

HPvM GO

HTvM

HGLM

HRpM

Mode: Normal

0 C/T 0

0 SF 0

0 JNL 0

1 Pool 0

0 JNLG 0

License 0

HTSM

HSSM

Applications

Exchange

172.17.72.35

SQL Server

ClusteredSQL

172.17.72.35

Applications > Exchange > 172.17.72.35

Host Summary

Name 172.17.72.35

Host ID 172.17.72.35

Storage Groups

Replica History

Rows/Page: 100 Page: 1 of 1

Storage Group	Status	Latest Replica
DF800E STORAGE GROUP	?	■
SECOND STORAGE GROUP	?	■

List of Application servers in the tree

List of storage groups and buttons for replica operations

Pair Management Create Replica Restore Replica

43

The Applications view lists the applications managed and the current protection status.

- HRpM allows for discover and configure of Agent through intuitive GUI.
 - New “Application Agent” in “Information Source” view provides the screens for the settings.

Information Sources

HDvM Server
Open Pair Management

Application
Replica Management

BC Manager
Mainframe Pair Management

HRpM Server

44

Application Agent
Information Source > Application Agent

Application Agent Setting

Rows/Page: 100 Page: 1 of 1

Name	Host ID	Port	Host Type	Configuration Type	Component Version
scsav6945	172.17.69.45	24041	Database/File server	Physical host (Stand-Alone)	6.3.0-00
scsav6946	172.17.69.46	24041	Backup server	Physical host (Stand-Alone)	6.3.0-00

Add Agent Remove Agents Setup Agent Refresh Agent

Discover Application Agent by specifying the host information

Application Agent Setting

If you have registered the Application Agent without selecting this to acquire the configuration, select "Application Agent" under "Info" and then execute "Refresh Agent".

*Name:

*Host ID:

*Host:

*Host Type: ☒ Database/File Server ☐ Backup Server

*Configuration Type: ☒ Physical Host (Stand-Alone) ☐ Virtual Host (Cluster)

Cluster Settings:

Application Type:

Virtual Host:

Cluster Nodes:

☒ Acquire application configuration managed by the Application Agent during the registration.

* Required field

OK Cancel

Application Support: Create Replica Wizard

- HRpM provides a new wizard for creating replica.
- Result of **Create Replica Wizard** is registered as task and **Task View** provides the detail of the task execution.

The screenshot displays the HRpM interface with three overlapping windows. The 'Tasks' window shows a list of tasks, with 'TASK0009' highlighted. The 'Edit Task - TASK0009' window is open, showing the 'Schedule' section with 'Execute immediately' selected. The 'Task History - TASK0009' window shows a table of execution results.

Task List

Task ID	Workflow	MF/Onen	Host	Target	Copy Type	Task Type	Status
TASK0009		Open	scsaw6945	FIRST STORAGE GROUP	Local	create replica	Success
TASK0008		Open	scsaw6945	FIRST STORAGE GROUP	Local	create replica	Failure
TASK0007		Open	scsaw6945	FIRST STORAGE GROUP	Local	create replica	Success
TASK0006		Open	scsaw6945	FIRST STORAGE GROUP	Local	create replica	Success
TASK0005		Open	scsaw6945	FIRST STORAGE GROUP	Local	create replica	Success
TASK0004		Open	scsaw6945	FIRST STORAGE GROUP	Local	create replica	Success
TASK0003		Open	scsaw6945	FIRST STORAGE GROUP	Local	create replica	Success

Edit Task - TASK0009

Schedule

When task has to be executed immediately, please select "Execute immediately".
When task has to be executed recursively, please select "Execute on schedule".
And, error information of previous execution can be referred from task history screen.

Execution Mode: ☐ Execute immediately ☐ Execute on schedule

Type:

Time (on Application Agent): : :

Task History - TASK0009

Task Summary

Application: Exchange
Host: scsaw6945
Task Type: create replica

Result List

Start Time	Status	Time for Completion
2008-10-13 16:18:13	Success	00:03:23
2008-10-15 15:55:19	Failure	00:03:25


"Edit Task" allows to update the execution schedule.

"Task History" allows to view the history of execution results

45

- Sites Enhancements

- “Applications” are also added in the Sites view.
 - Sites view allows most of the HRpM managements under site categorization.



46

Module Review

1. True or False? Site structure is shared by all users.
2. True or False? You can schedule the task of creating a Copy Pair and executing it automatically.
3. True or False? A resource can be added to multiple sites.

12. MxN Consistency Groups

Module Objectives

- Upon completion of this course, the learner should be able to:
 - Identify the licenses needed to use MxN Consistency Group (MxN CTG)
 - Define the concepts of MxN Consistency Group function
 - Describe the Architecture of MxN Consistency Groups with examples
 - Manage MxN Consistency Group with CCI
 - Define the current restrictions on the use of MxN Consistency Groups

Licensing

- Licensing considerations for MxN Consistency Group
 - **Local site** — Primary storage systems contain:
 - Customer data — HUR P-VOLs (source) and journal volumes
 - Disaster Recovery bundle for P-VOLs usable capacity only
 - Licenses Hitachi Universal Replicator
 - Disaster Recovery Extended bundle for P-VOLs usable capacity only
 - Licenses 3DC, MxN Consistency Group
 - No license capacity required for journal volumes
 - **Remote site** — Secondary storage systems contain:
 - Customer data — HUR S-VOLs (target) and journal volumes
 - Disaster Recovery bundle for S-VOLs usable capacity only
 - Licenses Hitachi Universal Replicator
 - Disaster Recovery Extended bundle for S-VOLs usable capacity only
 - Licenses 3DC, MxN Consistency group
 - No license capacity required for journal volumes


Concepts

- Allows CCI to maintain S-VOL data consistency across multiple Journal Groups
- Supported configurations:
 - Extends a consistency group across Journal Groups (up to four)
 - Journal Groups may exist in separate storage systems (up to 4) per site
 - Consistency of Consistency Group S-VOLs based on Consistency Q-Markers (timestamps)
- Intermix of VSP, USP V and USP VM *only*
 - Requires USP V microcode level 60-04-04-00/00 or higher
 - CCI version 01-23-03/02 or higher
- 2DC only


4

Concepts: Configuration Examples


■ Case 1: 1x1 Configuration with four Journals in one Consistency Group



■ Case 2: 4x4 Configuration




■ **Case 3: 4x1 Configuration**



7

■ **Case 4: 1x4 Configuration**



8

Concepts

- “Freeze & Run” architecture provides accurate time slice in order to deliver exact consistency.
 - This solution requires a periodic delay (default 3 sec intervals) of host I/O activities to insert “Consistency Q-marker” (CTQ-Marker).
- Primary site CCI Instance:
 - Issues “Freeze – Mark – Run” to all primary controllers in the HUR MxN configuration
- Remote site CCI Instance:
 - Requests and monitors the CTQ-Marker from all remote site controllers.
 - Decides what is the latest common CTQ-Marker from all the remote site controllers.
 - When the group is split, CCI requests all remote site controllers to settle data at the common CTQ-Marker.


CCI Processes for Managing

- CCI Instance delivers the CTQ-Marker to the defined MxN Consistency Group. If CCI Instances are shut down, then UR will stop committing changes to S-VOL. If shut down CCI, suspend the MxN Consistency groups first.
- **Freeze** – Halts all Channel Processors which are receiving application I/O for the relevant Consistency Group
- **Mark** – When all Channel Processors are in freeze mode, a command is sent to increment the extended consistency timestamp marker (CTQ-Marker).
 - This is not the same as the usual HUR Q-Marker and is updated much more slowly.
 - Q-Marker increases once per received I/O, so may increment tens of thousands of times a second.
 - CTQ-Marker increases once per cycle (by default one second).
- **Run** – Releases Channel Processors after the extended consistency CTQ-Marker is generated
 - All transactions received after the freeze will use the new CTQ-Marker.
 - With the factory estimate of 3% disruption at a 1 second cycle time, we can estimate that it takes about 30ms to execute a freeze/run operation.
- **Wait** – CCI waits before starting the next cycle to insert the next CTQ-Marker.
 - Default — three seconds

Architecture


Basic Design

- Primary CCI instance issues **Freeze**, inserts **CTQ-Marker** and **Run** to all MCUs.
- All MCUs send **CTQ-Marker** to connected remote site controllers.
- Remote CCI manages **CTQ-Markers** in secondary **MxN Consistency Groups**.
- When pairsplit command is issued, Remote CCI determines latest common **CTQ-Marker** among all RCUs for all secondary **MxN Consistency Groups**.
- In the example below, **QM2** is the most recent CTQ-Marker present in both remote systems.



11

Architecture: Order of Data Update



Managing: Sample HORCM Files

```

#/****** HORCM0 on production *****/
HORCM_MON
#ip_address service poll(10ms) timeout(10ms)
NONE horcm0 1000 3000

#/****** For HORCM_CMD *****/
HORCM_CMD
#dev_name
#UnitID #0(Serial# 64034), #1(Serial# 64045)
\\.\CMD-64034:/dev/rdsd
\\.\CMD-64045:/dev/rdsd

#/****** For HORCM_LDEV *****/
HORCM_LDEV
#dev_group dev_name Serial# LDEV# MU#
ora data0 64034 400
ora data1 64034 401
ora data2 64045 400
ora data3 64045 401

#/****** For HORCM_INST *****/
HORCM_INST
#dev_group ip_address service
ora RHOST horcm0

#/****** For UR of multiple DKC *****/
HORCM_CTQM
#dev_group interval(10ms) mode
ora 300
    
```

```

#/****** HORCM0 on Remote *****/
HORCM_MON
#ip_address service poll(10ms) timeout(10ms)
NONE horcm0 1000 3000

#/****** For HORCM_CMD *****/
HORCM_CMD
#dev_name
#UnitID #0(Serial# 64035), #1(Serial# 64046)
\\.\CMD-64035:/dev/rdsd
\\.\CMD-64046:/dev/rdsd

#/****** For HORCM_LDEV *****/
HORCM_LDEV
#dev_group dev_name Serial# LDEV# MU#
ora data0 64035 400
ora data1 64035 401
ora data2 64046 400
ora data3 64046 401

#/****** For HORCM_INST *****/
HORCM_INST
#dev_group ip_address service
ora PHOST horcm0

#/****** For UR of multiple DKC *****/
HORCM_CTQM
#dev_group interval(10ms) mode
ora 300
    
```

New HORCM_CTQM parameter (see note)

13

(1) HORCM_CTQM parameter

CCI adds "HORCM_CTQM" as a keyword in horcm.conf where

group is the MxN CT Group name as specified in the HORCM_LDEV parameter. Each group requires an entry.

interval is the interval for the CTQ-Marker. (1 sec default)

mode is the run mode for the CTQ-Marker (timestamp). The default run mode is freeze/run. This does not normally need to be specified. If "run" is specified, then the CTQ-Marker is issued without freeze.

Managing: Group ID

- MxN Consistency Group ID

Example:

```
#paircreate -g ora -f async 00 -v1 -jp 00 -js 01  
#paircreate -g ora1 -f async 00 -v1 -jp 02 -js 03
```

- First command creates **MxN Consistency Group** with ID **00** using M-JNL 00 and R-JNL **01**.
- Second command adds M-JNL **02** and R-JNL **03** to MxN Consistency Group ID **00**.

Note that the added M-JNL and R-JNL Group may be in different storage systems than the first group, depending on how the HORCM_LDEV parameter is defined in the configuration files.

Managing: Sample HORCM Files

```

# /***** HORCM0 on production *****/
HORCM_MON
#ip_address service poll(10ms) timeout(10ms)
NONE horcm0 1000 3000

# /***** For HORCM_CMD *****/
HORCM_CMD
#dev_name
#UnitID #0(Serial# 64034), #1(Serial# 64045)
\\.\CMD-64034:/dev/rdisk
\\.\CMD-64045:/dev/rdisk

# /***** For HORCM_LDEV *****/
HORCM_LDEV
#dev_group dev_name Serial# LDEV# MU#
ora data0 64034:1 400
ora data1 64034:1 401
ora data2 64045:2 400
ora data3 64045:2 401

# /***** For HORCM_INST *****/
HORCM_INST
#dev_group ip_address service
ora RHOST horcm0

# /***** For UR of multiple DRC *****/
HORCM_CTQM
#dev_group interval(10ms) mode
ora 300

```

```

# /***** HORCM0 on Remote *****/
HORCM_MON
#ip_address service poll(10ms) timeout(10ms)
NONE horcm0 1000 3000

# /***** For HORCM_CMD *****/
HORCM_CMD
#dev_name
#UnitID #0(Serial# 64046)
\\.\CMD-64046:/dev/rdisk

# /***** For HORCM_LDEV *****/
HORCM_LDEV
#dev_group dev_name Serial# LDEV# MU#
ora data0 64046:1 400
ora data1 64046:1 401
ora data2 64046:2 402
ora data3 64046:2 403

# /***** For HORCM_INST *****/
HORCM_INST
#dev_group ip_address service
ora PHOST horcm0

# /***** For UR of multiple DRC *****/
HORCM_CTQM
#dev_group interval(10ms) mode
ora 300

```

New method of specifying Journal Group IDs. It allows one CCI Consistency Group to contain more than one Journal Group (see note).

15

Specifying different JID for MxN Consistency Group

Specify Journal ID along with Serial Number in HORCM_LDEV parameter.

Note: If JID (Journal ID) is specified in horcm.conf as mentioned above, then the **paircreate** command need not specify Journal ID (-jp <jid> -js <jid>) option. This allows a MxN Consistency Group to contain multiple Journal Groups (up to four). Those Journal Groups can be in separate storage systems if desired.

If JID (Journal ID) is not specified in horcm.conf, then Journal ID (-jp <jid> -js <jid>) option of the paircreate command is used. In that case, one Journal Group can contain multiple CCI groups with one Consistency Group ID. Again, the Consistency group can span storage systems.

Managing: Group ID

- MxN Consistency Group ID (*continued*)
 - Horcm files contain Journal Group IDs:
 - First Journal M-JNL 01, R-JNL 01
 - Second Journal M-JNL 02, R-JNL 02

paircreate -g ora -f async 01 -vl -IH8

The command creates group *ora* with **MxN Consistency Group ID 01** that contains M-JNL groups 01 and 02, replicating to R-JNL groups 01 and 02 respectively.

- MxN Consistency Group ID / Journal Group ID
 - The MxN Consistency Group ID ***must*** be unique.
 - Journal Group IDs within a MxN Consistency Group ***must*** be unique.
 - However, associated M-JNL and R-JNL IDs can be the same.
 - **Best practice:** Select *unique* M-JNL and R-JNL group IDs across all participating storage systems.
 - This eliminates Journal Group ID conflicts when creating MxN Consistency Groups.
 - Examples: (see note)
 - MxN Consistency Group 00 may contain:
 - S/N 23456, M-JNL ID = 00, R-JNL ID = 00
 - S/N 34567, M-JNL ID = 01, R-JNL ID = 01
 - MxN Consistency Group 02 contains:
 - S/N 23456, M-JNL ID = 02, R-JNL ID = 03
 - S/N 23456, M-JNL ID = 04, R-JNL ID = 05

17

Notes:

UR supports Consistency Group IDs 00-FF

ShadowImage and TrueCopy support Consistency Group IDs 00-7F

Managing: Storage Navigator Display

- Example shows two sets of Journal Groups in Consistency Group 00.

The top screenshot shows the 'Universal Replicator Pair Operation' window. The 'Subsystem' list on the left includes CL1-A, CL3-A, CL5-A, CL7-A, CL1-B, CL3-B, CL5-B, CL7-B, CL1-E, CL3-E, CL5-E, CL7-E, CL1-F, CL3-F, CL5-F, and CL7-F. The main table displays the following data:

VOL	Status	JNLG-MirrorID	S/N (LDKC)	CTRLID
CL1-A - 00 - 000 (00:00:00)	COPY	00-01	10156(00)	5(USF)
CL1-A - 00 - 001 (00:00:01)	COPY	01-01	10156(00)	5(USF)

The bottom screenshot shows the same window with the 'Subsystem' list expanded to show CL1-A. The main table displays the following data:

Paired VOL	Pair JNLG	CTG	Err Lv.	Sync.	CLPR
CL1-A - 00 - 000 (00:00:00)	02	00	Group	014 %	00:CLPR0
CL1-A - 00 - 001 (00:00:01)	03	00	Group	013 %	00:CLPR0
					00:CLPR0
					00:CLPR0
					00:CLPR0
					00:CLPR0
					00:CLPR0

At the bottom of the window, it shows 'Used Volume : 0.01 (999.0)TB' and 'Total Pairs : 2'.

Managing: Sample pairedisplay Results

- Enhanced **pairedisplay** command
 - v jnl** parameter displays the **MxN Consistency Group** and **CTQ-Marker**

```
# pairedisplay -g ora -v jnl -fxe
```

JID	MU	CTG	JNLS	AP	U(%)	Q-Marker	Q-CNT	D-SZ (BLK)	Seq#	Num	LDEV#	CTQM
000	2	00	PJSN	1	0	0000bb1b	0	198578688	64014	1	200	-
000	2	00	SJSN	1	0	0000bb1b	0	99283968	64035	1	400	0000187f
001	2	00	PJSN	1	0	00000011	0	198578688	64014	1	300	-
001	2	00	SJSN	1	0	00000011	0	99283968	64035	1	500	0000187f

- Note that LDEV column displays the LDEV number of the first Journal Volume in the Group.
- Note that CTGID is always expressed as a decimal number, while JID (Journal Group ID) is a hex number.

Enhanced **pairedisplay** command: **-v jnlt**

- Timer settings for async remote replication
 - DOW: Data Overflow Watch** setting in Journal Group Options panel
 - PBW: Path Blockade Watch:** Not settable
 - APW: Path Watch Time** setting in Journal Group Options Panel

```
# pairedisplay -g ora -v jnlt
```

JID	MU	CTG	JNLS	AP	U(%)	Q-Marker	Q-CNT	D-SZ (BLK)	Seq#	DOW	PBW	APW
001	1	0	PJSN	1	0	00000c76	20	1633672703	64034	60	300	40
002	1	0	SJSN	1	0	00000c62	20	1633672703	64035	60	300	40
001	1	0	PJSN	1	0	00000c3a	7	211506164	64045	60	300	40
002	1	0	SJSN	1	0	00000c33	7	260319089	64046	60	300	40

Managing: Sample pairedisplay Results

- **-v ctg** Option: Displays Inflow Control and timer settings of each group (see note)

Example:

```
# pairedisplay -g ora -v ctg
CTG P/S   Status AP U(%) Q-Marker QM-Cnt SF(%) Seq# IFC OT/s CT/m RT/m
001 P-VOL PAIR  2  0   00000080 3     50   63528 ON  90   5    5
001 S-VOL PAIR  -  0   0000007d -     50   63528 -   -    -    -
```

CTG: Displays the MxN Consistency Group ID.

P/S : The attribute of a volume in first LDEV of the specified group.

Status: The status of the paired volume in first LDEV of the specified group.

AP: Displays the number of Active Path in HORC links on PVOL, also displays the number of active path in UR links on PVOL and SVOL, 'Unknown' is shown as '- '.

U(%): For UR: The usage rate of the current journal data

Q-Marker: In P-VOL, the latest sequence # of the MCU P-Vol when the write command was received. In S-VOL, the latest sequence # on RCU. This item is valid at PAIR state.

QM-Cnt: The number of Q-Markers within Consistency Group (or MxN Consistency Group if defined) of the Unit. If no data is being replicated,

UR/TrueCopy: Sends a dummy recordset at regular interval, therefore QM-Cnt always shows "2" or "3" even if Host has NO replication data. This item is valid at PAIR state.

SF(%): The usage of cache setting as the sidefile regardless of UR and TC Async.

Seq#: The serial number of the RAID storage system.

IFC: Displays INFLOW CONTROL Setting in Storage Navigator Journal Group Options.

OT/s: The “offloading timer” (in unit of Sec.) setting to Consistency group for UR / TC Async. In UR, this is the same as “DOW” item shown by `raidvchkscan -v jnlt` or `pairedisplay -v jnlt`.

CT/m: The “Copy Pending timer” (TC Async only)

RT/m: The “RCU Ready timer” (TC Async only)

Managing: pairsplit Command

- **pairsplit** command — Internal behavior is different from 1x1 UR.
- **pairsplit -r or -rw** option
 - Issues **Freeze** to MxN Consistency Group on each **MCU**
 - Issues Suspend & Run to create **CTQ-Markers** for MxN Consistency Group on each **MCU**
 - Commits a minimum matching point of the CTQ-Marker on **RCU** (in other words, do Journal Restore with CTQ-Marker)
 - Repeats above until detecting an End Of Marker (EOM) on all RCU via MCU
 - Issues End of Suspend to terminate a suspending state
 - [Exception] If link failure is detected during **pairsplit**, then the commit operation is aborted, keeping current CTQ-Marker level, and a suspending state terminates without waiting for the EOM. Results in PSUE Status for the group.

22

- **pairsplit -S** option – Delete pair from local site
 - Issues Freeze to MxN Consistency Group on each MCU
 - Issues Delete & Run for MxN Consistency Group on each MCU
 - Commits a minimum matching point of the CTQ-Marker on RCU via MCU (in other words, do Journal Restore with CTQ-Marker)
 - Repeats above until detecting an EOM marked on MCU on all RCU via MCU
 - Issues End of Delete to terminate a deleting state, after committed with EOM on all RCUs

▪ **pairsplit -RS** option – Suspend swap

- Issues Swap Suspend to make a suspending state for MxN Consistency Group on each **RCU**.
- Searches/commits a minimum matching point of the CTQ-Marker on RCU (in other words, do Journal Restore with CTQ-Marker).
- Repeats above until detecting an EOM on all RCU.
- Issues End of Suspend to terminate a suspending state, after committed with an EOM on all RCU.
- Ending MxN Consistency Group status will be **SSWS**.

24

▪ **pairsplit -R** option – Forces secondary volume to simplex

- Issues Delete to make a deleting state for MxN Consistency Group on each RCU.
- Commits a minimum matching point of the CTQ-Marker on RCU (in other words, do Journal Restore with CTQ-Marker).
- Repeats above until detecting an EOM on all RCUs.
- Issues End of Delete to terminate a deleting state, after committed with an EOM on all RCUs.
- Ending status will be Simplex for all volumes in the group. Will terminate MxN CT group operations.

▪ **pairsplit -RB** option – Forces secondary volume from SSWS to SSUS

25

Restrictions


- Four journal groups per MxN Consistency Group.
- 8192 LDEVs per MxN Consistency Group.
- In the case of multiple MxN Consistency Groups per HORCM instance:
 - No more than 256 LDEVs per MxN Consistency Group is recommended.
 - Multiple CCI instances may be required.
- 3DC configuration not yet supported.
- CCI Instance delivers the CTQ-Marker to the defined MxN Consistency Group. If CCI Instances are shut down, then UR will stop committing changes to S-VOL. If you shut down CCI, suspend the MxN Consistency groups first.

Module Review

1. Define the purpose of the Open UR MxN Consistency Group function.
2. Identify the licenses needed.
3. Identify supported MxN configurations.
4. State the “best practice rule” for MxN Consistency Group configurations.
5. State the Preliminary Conditions required for MxN Consistency Group management with CCI.

Your Next Steps





Training Course Glossary

[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—

ACC — Action Code. A SIM System Information Message. Will produce an ACC, which requires an engineer to the correct fix procedures in the ACC directory in the MM (Maintenance Manual).

ACE — Access Control Entry. Stores access rights for a single user or group within the Windows security model

ACL — Access Control List. Stores a set of ACEs, so that it describes the complete set of access rights for a file system object within the Microsoft Windows security model

ACP — Array Control Processor. Microprocessor mounted on the disk adapter circuit board (DKA) that controls the drives in a specific disk array. Considered part of the back end; it controls data transfer between cache and the hard drives.

ACP Domain — Also Array Domain. All of the array-groups controlled by the same pair of DKA boards, or the HDDs managed by one ACP PAIR (also called BED).

ACP PAIR — Physical disk access control logic. Each ACP consists of two DKA PCBs. To provide eight loop paths to the real HDDs

Actuator (arm) — Read/write heads are attached to a single head actuator, or actuator arm, that moves the heads around the platters.

AD — Active Directory

ADC — Accelerated Data Copy

Address — A location of data, usually in main memory or on a disk. A name or token that identifies a network component. In local area networks (LANs), for example, every node has a unique address.

ADP — Adapter

ADS — Active Directory Service

AIX — IBM UNIX

AL — Arbitrated Loop. A network in which nodes contend to send data, and only one node at a time is able to send data.

AL-PA — Arbitrated Loop Physical Address

AMS — Adaptable Modular Storage

APAR — Authorized Program Analysis Reports

APF — Authorized Program Facility. In z/OS and OS/390 environments, a facility that permits the identification of programs that are authorized to use restricted functions.

API — Application Programming Interface

APID — An ID to identify a command device

Application Management — The processes that manage the capacity and performance of applications.

ARB — Arbitration or *request*

ARM — Automated Restart Manager

Array Domain — Also ACP Domain. All functions, paths, and disk drives controlled by a single ACP pair. An array domain can contain a variety of LVI or LU configurations.

Array Group — Also called a parity group, is a group of hard disk drives (HDDs) that form the basic unit of storage in a subsystem. All HDDs in a parity group must have the same physical capacity.

Array Unit — A group of Hard Disk Drives in one RAID structure. Same as Parity Group.

ASIC — Application specific integrated circuit

ASSY — Assembly

Asymmetric virtualization — See Out-of-band virtualization.

Asynchronous — An I/O operation whose initiator does not await its completion before proceeding with other work. Asynchronous I/O operations enable an initiator to have multiple concurrent I/O operations in progress. Also called Out-of-band virtualization.

ATA — Advanced Technology Attachment. A disk drive implementation that integrates the controller on the disk drive itself. Also known as IDE (Integrated Drive Electronics) Advanced Technology Attachment.

Authentication — The process of identifying an individual, usually based on a username and password.

AUX — Auxiliary Storage Manager

Availability — Consistent direct access to information over time

[-back to top-](#)

—B—

B4 — A group of 4 HDU boxes that are used to contain 128 HDDs

Back end — In client/server applications, the client part of the program is often called the front end and the server part is called the back end.

Backup image—Data saved during an archive operation. It includes all the associated files, directories, and catalog information of the backup operation.

BADM — Basic Direct Access Method

BASM — Basic Sequential Access Method

BATCTR — Battery Control PCB

BC — Business Class (in contrast with EC, Enterprise Class)

BCM — Hitachi Business Continuity Manager

BCP — Base Control Program

BCPii — Base Control Program internal interface

BDW — Block Descriptor Word

BED — Back end director. Controls the paths to the HDDs.

Bind Mode — One of two modes available when using FlashAccess™, in which the FlashAccess™ extents hold read data for specific extents on volumes (see Priority Mode).

BLKSIZE — Block size

BMC — Baseline Management Controller

BOS — Basic Operating System

BOS V — Basic Operating System Virtualization

BPAM — Basic Partitioned Access Method

BPM — Business Process Management

BST — Binary Search Tree

BTU — British Thermal Unit

Business Continuity Plan — Describes how an organization will resume partially or completely interrupted critical functions within a predetermined time after a disruption or a disaster. Sometimes also called a Disaster Recovery Plan.

[-back to top-](#)

—C—

CA — Continuous Access software (see HORC) or Continuous Availability (or Computer Associates)

Cache — Cache Memory. Intermediate buffer between the channels and drives. It has a maximum capacity of 64GB (32GB x 2 areas). It is available and controlled as two areas of cache (cache A and cache B). It is fully battery-backed (48 hours).

Cache hit rate — When data is found in the cache, it is called a cache hit, and the effectiveness of a cache is judged by its hit rate.

Cache partitioning — Storage management software that allows the virtual partitioning of cache and allocation of it to different applications.

CAD — Computer-Aided Design

Capacity — Capacity is the amount of data that a drive can store after formatting. Most data storage companies, including HDS, calculate capacity based on the assumption that 1KB = 1,024 bytes, 1MB = 1,024 kilobytes, 1GB = 1,024 megabytes, and 1TB = 1,024 gigabytes.

CAPEX — Capital expenditure is the cost of developing or providing non-consumable parts for the product or system. For example, the purchase of a photocopier is the CAPEX, and the annual paper and toner cost is the OPEX. (See OPEX).

CAS — Column address strobe is a signal sent to a dynamic random access memory (DRAM) that tells it that an associated address is a column address. CAS- column address strobe sent by the processor to a DRAM circuit to activate a column address.

CBU — Capacity Backup

CCHH — Common designation for Cylinder and Head

CCI — Command Control Interface

CDP — Continuous Data Protection

CDWP — Cumulative disk write throughput

CE — Customer Engineer

CEC — Central Electronics Complex

CentOS — Community Enterprise Operating System

Centralized management — Storage data management, capacity management, access security management, and path management functions accomplished by software.

CF — Coupling Facility

CFCC — Coupling Facility Control Code

CFW — Cache Fast Write

CH — Channel

CH S — Channel SCSI

CHA — Channel Adapter. Provides the channel interface control functions and internal cache data transfer functions. It is used to convert the data format between CKD and FBA. The CHA contains an internal processor and 128 bytes of edit buffer memory.

CHA/DKA — Channel Adapter/Disk Adapter

CHAP — Challenge-Handshake Authentication Protocol

CHF — Channel Fibre

CHIP — Client-Host Interface Processor. Microprocessors on the CHA boards that process the channel commands from the hosts and manage host access to cache.

CHK — Check

CHN — CHannel adapter NAS

CHP — Channel Processor or Channel Path

CHPID — Channel Path Identifier

CHSN or C-HSN — Cache Memory Hierarchical Star Network

CHT — Channel tachyon, a Fibre Channel protocol controller

CICS — Customer Information Control System

CIFS protocol — Common internet file system is a platform-independent file sharing system. A network file system accesses protocol primarily used by Windows clients to communicate file access requests to Windows servers.

CIM — Common Information Model

CKD — Count-key Data. A format for encoding data on hard disk drives; typically used in the mainframe environment.

CKPT — Check Point

CL — See Cluster

CLI — Command Line Interface

CLPR — Cache Logical PaRtition. Cache can be divided into multiple virtual cache memories to lessen I/O contention.

Cluster — A collection of computers that are interconnected (typically at high-speeds) for the purpose of improving reliability, availability, serviceability or performance (via load balancing). Often, clustered computers have access to a common pool of storage, and run special software to coordinate the component computers' activities.

CM — Cache Memory, Cache Memory Module. Intermediate buffer between the channels and drives. It has a maximum of 64GB (32GB x 2 areas) of capacity. It is available and controlled as two areas of cache (cache A and cache B). It is fully battery-backed (48 hours).

CM DIR — Cache Memory Directory

CM-HSN — Control Memory Hierarchical Star Network

CM PATH — Cache Memory Access Path. Access Path from the processors of CHA, DKA PCB to Cache Memory.

CM POK — Cache Memory Package

CM/SM — Cache Memory/Shared Memory

CMA — Cache Memory Adapter

CMD — Command

CMG — Cache Memory Group

CNAME — Canonical NAME

CNS — Cluster Name Space or Clustered Name Space

CNT — Cumulative network throughput

CoD — Capacity on Demand

Concatenation — A logical joining of two series of data, usually represented by the symbol “|”. In data communications, two or more data are often concatenated to provide a unique name or reference (e.g., S_ID | X_ID). Volume managers concatenate disk address spaces to present a single larger address spaces.

Connectivity technology — a program or device's ability to link with other programs and devices. Connectivity technology allows programs on a given computer to run routines or access objects on another remote computer

Controller — A device that controls the transfer of data from a computer to a peripheral device (including a storage system) and vice versa.

Controller-based Virtualization — Driven by the physical controller at the hardware microcode level versus at the application software layer and integrates into the infrastructure to allow virtualization across heterogeneous storage and third party products

Corporate governance — Organizational compliance with government-mandated regulations

CoW — Copy-on-Write snapshot or software

CP — Central Processor (also called Processing Unit or PU)

CPC — Central Processor Complex

CPM — Cache Partition Manager. Allows for partitioning of the cache and assigns a partition to a LU; this enables tuning of the system's performance.

CPS — Cache Port Slave

CPU — Central Processing Unit

CRM — Customer Relationship Management

CruiseControl — Now called Hitachi Volume Migration software

CSS — Channel Subsystem

CS&S — Customer Service and Support

CSTOR — Central Storage or Processor Main Memory

C-Suite — The C-suite is considered the most important and influential group of individuals at a company. Referred to as “the C-Suite within a Healthcare provider.”

CSV — Comma Separated Value

CSW — Cache Switch PCB. The cache switch (CSW) connects the channel adapter or disk adapter to the cache. Each of them is connected to the cache by the Cache Memory Hierarchical Star Net (C-HSN) method. Each cluster is provided with the two CSWs, and each CSW can connect four caches. The CSW switches any of the cache paths to which the channel adapter or disk adapter is to be connected through arbitration.

CTE — Common Technology Engine. CommVault® is one branded CTE.

CTG — Consistency Group

CTN — Coordinated Timing Network

CU — Control Unit (refers to a storage subsystem. The hexadecimal number to which 256 LDEVs may be assigned.

CUDG — Control Unit DiaGnostics. Internal system tests.

CUoD — Capacity Upgrade on Demand

CV — Custom Volume

CVS — Customizable Volume Size. Software used to create custom volume sizes. Marketed under the name Virtual LVI (VLVI) and Virtual LUN (VLUN).

CXRC — Coupled z/OS Global Mirror

[-back to top-](#)

—D—

DA — Device Adapter

DACL — Discretionary access control list (ACL). The part of a security descriptor that stores access rights for users and groups.

DAD — Device Address Domain. Indicates a site of the same device number automation support function. If several hosts on the same site have the same device number system, they have the same name.

DAMP — Disk Array Management Program. Renamed to Storage Navigator Modular (SNM).

DAS — Direct Attached Storage

DASD — Direct Access Storage Device

Data Blocks — A fixed-size unit of data that is transferred together. For example, the X-modem protocol transfers blocks of 128 bytes. In general, the larger the block size, the faster the data transfer rate.

Data Integrity — Assurance that information will be protected from modification and corruption.

Data Lifecycle Management — An approach to information and storage management. The policies, processes, practices, services and tools used to align the business value of data with the most appropriate and cost-effective storage infrastructure from the time data is created through its final disposition. Data is aligned with business requirements through management policies and service levels associated with performance, availability, recoverability, cost, and what ever parameters the organization defines as critical to its operations.

Data Migration — The process of moving data from one storage device to another. In this context, data migration is the same as Hierarchical Storage Management (HSM).

Data Migrator — Software allows for using existing storage resources more efficiently by overseeing the migration of data to less expensive secondary storage. This includes file system data from a Microsoft® Windows® or UNIX host. By shrinking the size of your primary storage, Data Migrator reduces the backup window and prevents over-capacity. Also referred to as Hierarchical Storage Management (HSM).

Data Pool — A volume containing differential data only.

Data Protection Suite platform — Designed to automate management of data by changing the approach from static to dynamic. Built on the CTE. See Hitachi Data Protection Suite.

Data Stream — CommVault's patented high performance data mover used to move data back and forth between a data source and a MediaAgent or between two MediaAgents.

Data Striping — Disk array data mapping technique in which fixed-length sequences of virtual disk data addresses are mapped to sequences of member disk addresses in a regular rotating pattern.

Data Transfer Rate (DTR) — The speed at which data can be transferred. Measured in kilobytes per second for a CD-ROM drive, in bits per second for a modem, and in megabytes per second for a hard drive. Also, often called simply data rate.

DBMS — Data Base Management System

DCA — Data Cache Adapter

DCR — Dynamic Cache Residency. See FlashAccess™

DDL — Database Definition Language

DDM — Disk Drive Module

DDNS — Dynamic DNS

DE — Data Exchange Software

Device Management — Processes that configure and manage storage systems

Data Pipe or Data Stream — The connection set up between the MediaAgent, source or destination server is called a *Data Pipe* or more commonly a *Data Stream*.

DFS — Microsoft Distributed File System

DFSMS — Data Facility Storage Management Subsystem

DFSMS SDM — Data Facility Storage Management Subsystem System Data Mover

DFSMSdftp — Data Facility Storage Management Subsystem Data Facility Product

DFSMSdss — Data Facility Storage Management Subsystem Data Set Services

DFSMSHsm — Data Facility Storage Management Subsystem Hierarchical Storage Manager

DFSMSRmm — Data Facility Storage Management Subsystem Removable Media Manager

DFSMSStvs — Data Facility Storage Management Subsystem Transactional VSAM Services

DFW — DASD Fast Write

DICOM — Digital Imaging and Communications in Medicine

DIMM — Dual In-line Memory Module

Direct Access Storage Device — A type of storage device, in which bits of data are stored at precise locations, enabling the computer to retrieve information directly without having to scan a series of records.

Direct Attached Storage — Storage that is directly attached to the application or file server. No other device on the network can access the stored data.

Director class switches — Larger switches often used as the core of large switched fabrics

Disaster Recovery Plan (DRP) — A plan that describes how an organization will deal with potential disasters. It may include the precautions taken to either maintain or quickly resume mission-critical functions. Sometimes also referred to as a Business Continuity Plan.

Disk Administrator — An administrative tool that displays the actual LU storage configuration.

Disk Array — A linked group of one or more physical independent hard disk drives generally used to replace larger, single disk drive systems. The most common disk arrays are in daisy chain configuration or implement RAID (Redundant Array of Independent Disks) technology.

A disk array may contain several disk drive trays, and is structured to improve speed and increase protection against loss of data. Disk arrays organize their data storage into Logical Units (LUs), which appear as linear block paces to their clients. A small disk array, with a few disks, might support up to eight LUs; a large one, with hundreds of disk drives, can support thousands.

DKA — Disk Adapter. Also called an array control processor (ACP); it provides the control functions for data transfer between drives and cache. The DKA contains DRR (Data Recover and Reconstruct), a parity generator circuit. It supports four fibre channel paths and offers 32KB of buffer for each fibre channel path.

DKC — Disk Controller Unit. In a multi-frame configuration, the frame that contains the front end (control and memory components).

DKCMN — Disk Controller Monitor. Monitors temperature and power status throughout the machine.

DKF — Fibre disk adapter. Another term for a DKA.

DKU — Disk Array Frame or Disk Unit or Disk Unit. In a multi-frame configuration, a frame that contains hard disk units (HDUs).

DKUPS — Disk Unit Power Supply

DLIBs — Distribution Libraries

DKUP — Disk Unit Power Supply

DLM — Data Lifecycle Management

DMA — Direct Memory Access

DM-LU — Differential Management Logical Unit.
DM-LU is used for saving management information of the copy functions in the cache.

DMP — Disk Master Program

DNS — Domain Name System

DOC — Deal Operations Center

Domain — A number of related storage array groups

DOO — Degraded Operations Objective

DP — Dynamic Provisioning (pool)

DP-VOL — Dynamic Provisioning Virtual Volume

DR — Disaster Recovery

DRAC — Dell Remote Access Controller

DRAM — Dynamic random access memory

DRP — Disaster Recovery Plan

DRR — Data Recover and Reconstruct. Data Parity Generator chip on DKA.

DRV — Dynamic Reallocation Volume

DSB — Dynamic Super Block

DSF — Device Support Facility

DSF INIT — Device Support Facility Initialization (for DASD)

DSP — Disk Slave Program

DTA — Data adapter and path to cache-switches

DTR — Data Transfer Rate

DVE — Dynamic Volume Expansion

DW — Duplex Write

DWL — Duplex Write Line or Dynamic Workspace Linking

Dynamic Link Manager — HDS software that ensures that no single path becomes overworked while others remain underused. Dynamic Link Manager does this by providing automatic load balancing, path failover, and recovery capabilities in case of a path failure.

[-back to top-](#)

—E—

EAV — Extended Address Volume

EC — Enterprise Class (in contrast with BC, Business Class)

ECC — Error Checking and Correction

ECC.DDR SDRAM — Error Correction Code Double Data Rate Synchronous Dynamic RAM Memory

ECM — Extended Control Memory

ECN — Engineering Change Notice

E-COPY — Serverless or LAN free backup

EMIF — ESCON Multiple Image Facility

ENC — Enclosure or ENClosure Controller. The units that connect the controllers in the DF700 with the Fibre Channel disks. They also allow for online extending a system by adding RKAs

EOF — End of Field

EOL — End of Life

EPO — Emergency Power Off

EREP — Error REporting and Printing

ERP — Enterprise Resource Management

ESA — Enterprise Systems Architecture

ESB — Enterprise Service Bus

ESC — Error Source Code

ESD — Enterprise Systems Division

ESCD — ESCON Director

ESCON — Enterprise Systems Connection. An input/output (I/O) interface for mainframe computer connections to storage devices developed by IBM.

ESDS — Entry Sequence Data Set

ESS — Enterprise Storage Server

ESW — Express Switch or E Switch. Also referred to as the Grid Switch (GSW).

Ethernet — A local area network (LAN) architecture that supports clients and servers and uses twisted pair cables for connectivity.

ETR — External Time Reference (device)

EVS — Enterprise Virtual Server

EXCP — Execute Channel Program

ExSA — Extended Serial Adapter

[-back to top-](#)

—F—

Fabric — The hardware that connects workstations and servers to storage devices in a SAN is referred to as a "fabric." The SAN fabric enables any-server-to-any-storage device connectivity through the use of Fibre Channel switching technology.

Failback — The restoration of a failed system share of a load to a replacement component. For example, when a failed controller in a redundant configuration is replaced, the devices that were originally controlled by the failed controller are usually failed back to the replacement controller to restore the I/O balance, and to restore failure tolerance. Similarly, when a defective fan or power supply is replaced, its load, previously borne by a redundant component, can be failed back to the replacement part.

Failed over — A mode of operation for failure tolerant systems in which a component has failed and its function has been assumed by a redundant component. A system that protects against single failures operating in failed over mode is not failure tolerant, since failure of the redundant component may render the system unable to function. Some systems (e.g., clusters) are able to tolerate more than one failure; these remain failure tolerant until no redundant component is available to protect against further failures.

Failover — A backup operation that automatically switches to a standby database server or network if the primary system fails, or is temporarily shut down for servicing. Failover is an important fault tolerance function of mission-critical systems that rely on constant accessibility. Failover automatically and transparently to the user redirects requests from the failed or down system to the backup system that mimics the operations of the primary system. Also called path failover.

Failure tolerance — The ability of a system to continue to perform its function or at a reduced performance level, when one or more of its components has failed. Failure tolerance in disk subsystems is often achieved by including redundant instances of components whose failure would make the system inoperable, coupled with facilities that allow the redundant components to assume the function of failed ones.

FAIS — Fabric Application Interface Standard

FAL — File Access Library

FAT — File Allocation Table

Fault Tolerant — Describes a computer system or component designed so that, in the event of a component failure, a backup component or procedure can immediately take its place with no loss of service. Fault tolerance can be provided with software, embedded in hardware, or provided by some hybrid combination.

FBA — Fixed-block Architecture. Physical disk sector mapping.

FBA/CKD Conversion — The process of converting open-system data in FBA format to mainframe data in CKD format.

FBUS — Fast I/O Bus

FC — Fibre Channel or Field-Change (microcode update) or Fibre Channel. A technology for transmitting data between computer devices; a set of standards for a serial I/O bus capable of transferring data between two ports

FC RKAJ — Fibre Channel Rack Additional.

Acronym referring to an additional rack unit(s) that houses additional hard drives exceeding the capacity of the core RK unit of the Thunder 9500V/9200 subsystem.

FC-0 — Lowest layer on fibre channel transport. This layer represents the physical media.

FC-1 — This layer contains the 8b/10b encoding scheme.

FC-2 — This layer handles framing and protocol, frame format, sequence/exchange management and ordered set usage.

FC-3 — This layer contains common services used by multiple N_Ports in a node.

FC-4 — This layer handles standards and profiles for mapping upper level protocols like SCSI an IP onto the Fibre Channel Protocol.

FCA — Fibre Adapter. Fibre interface card. Controls transmission of fibre packets.

FC-AL — Fibre Channel Arbitrated Loop. A serial data transfer architecture developed by a consortium of computer and mass storage device manufacturers, and is now being standardized by ANSI. FC-AL was designed for new mass storage devices and other peripheral devices that require very high bandwidth. Using optical fiber to connect devices, FC-AL supports full-duplex data transfer rates of 100MBps. FC-AL is compatible with SCSI for high-performance storage systems.

FCC — Federal Communications Commission

FCIP — Fibre Channel over IP, a network storage technology that combines the features of Fibre Channel and the Internet Protocol (IP) to connect distributed SANs over large distances. FCIP is considered a tunneling protocol, as it makes a transparent point-to-point connection between geographically separated SANs over IP networks. FCIP relies on TCP/IP services to establish connectivity between remote SANs over LANs, MANs, or WANs. An advantage of FCIP is that it can use TCP/IP as the transport while keeping Fibre Channel fabric services intact.

FCP — Fibre Channel Protocol

FC-P2P — Fibre Channel Point-to-Point

FC-SW — Fibre Channel Switched

FCU — File Conversion Utility

FD — Floppy Disk or Floppy Drive

FDR — Fast Dump/Restore

FE — Field Engineer

FED — Channel Front End Directors

Fibre Channel — A serial data transfer architecture developed by a consortium of computer and mass storage device manufacturers and now being standardized by ANSI. The most prominent Fibre Channel standard is Fibre Channel Arbitrated Loop (FC-AL).

FICON — Fiber Connectivity. A high-speed input/output (I/O) interface for mainframe computer connections to storage devices. As part of IBM's S/390 server, FICON channels increase I/O capacity through the combination of a new architecture and faster physical link rates to make them up to eight times as efficient as ESCON (Enterprise System Connection), IBM's previous fiber optic channel standard.

Flash ACC — Flash access. Placing an entire LUN into cache.

FlashAccess — HDS software used to maintain certain types of data in cache to ensure quicker access to that data.

FLGFAN — Front Logic Box Fan Assembly

FLOGIC Box — Front Logic Box

FM — Flash Memory. Each microprocessor has FM. FM is non-volatile memory that contains microcode.

FOP — Fibre Optic Processor or fibre open

FPC — Failure Parts Code or Fibre Channel Protocol Chip

FPGA — Field Programmable Gate Array

Frames — An ordered vector of words that is the basic unit of data transmission in a Fibre Channel network.

Front end — In client/server applications, the client part of the program is often called the front end and the server part is called the back end.

FRU — Field Replaceable Unit

FS — File System

FSA — File System Module-A

FSB — File System Module-B

FSM — File System Module

FSW — Fibre Channel Interface Switch PCB. A board that provides the physical interface (cable connectors) between the ACP ports and the disks housed in a given disk drive.

FTP — File Transfer Protocol. A client-server protocol that allows a user on one computer to transfer files to and from another computer over a TCP/IP network.

FWD — Fast Write Differential

[-back to top-](#)

—G—

GARD — General Available Restricted Distribution

Gb — Gigabit

GB — Gigabyte

Gb/sec — Gigabit per second

GB/sec — Gigabyte per second

GbE — Gigabit Ethernet

GBIC — Gigabit Interface Converter

GDG — Generation Data Group

GDPS — Geographically Dispersed Parallel Sysplex

GID — Group Identifier within the Unix security model.

gigE — Gigabit Ethernet

GLM — Gigabyte Link Module

Global Cache — Cache memory is used on demand by multiple applications. Use changes dynamically, as required for READ performance between hosts/applications/LUs.

GPFS — General Parallel File System

Graph-Track™ — HDS software used to monitor the performance of the Hitachi storage subsystems. Graph-Track™ provides graphical displays, which give information on device usage and system performance.

GSC — Global Support Center

GSS — Global Solutions Services

GSSD — Global Solutions Strategy and Development

GSW — Grid Switch Adapter. Also known as E Switch (Express Switch).

GUI — Graphical User Interface

GUID — Globally Unique Identifier

GWS — Grid Switch

[-back to top-](#)

—H—

H1F — Essentially the Floor Mounted disk rack (also called Desk Side) equivalent of the RK. (See also: RK, RKA, and H2F).

H2F — Essentially the Floor Mounted disk rack (also called Desk Side) add-on equivalent similar to the RKA. There is a limitation of only one H2F that can be added to the core RK Floor Mounted unit. See also: RK, RKA, and H1F.

HA — High Availability

HATP — Hitachi Assurance Test Program

HBA — Host Bus Adapter. An I/O adapter that sits between the host computer's bus and the Fibre Channel loop and manages the transfer

of information between the two channels. In order to minimize the impact on host processor performance, the host bus adapter performs many low-level interface functions automatically or with minimal processor involvement.

HCA — Host Channel Adapter

HCAP — Hitachi Content Archive Platform

HCD — Hardware Configuration Definition

HCmD — Hitachi Command Director software.

HCP — Hitachi Content Platform

HCS — Hitachi Command Suite (formerly Hitachi Storage Command Suite)

HD — Hard Disk

HDA — Head Disk Assembly

HDD — Hard Disk Drive. A spindle of hard disks that make up a hard drive, which is a unit of physical storage within a subsystem.

HDDPWR — Hard Disk Drive Power

HD DS — Hitachi Data Discovery Suite

HDev — Hidden devices. Hitachi Tuning Manager Main Console may not display some drive letters in its resource tree, and information such as performance and capacity is not available for such invisible drives. This problem occurs if there is a physical drive with lower Physical Drive number assigned that is in “damaged (SCSI Inquiry data cannot be obtained)” or “hidden by HDLM” status.

HDI — Hitachi Dynamic Ingestor

HDI Remote — HDI Cloud Storage & Computing

HDLM — Hitachi Dynamic Link Manager software

HDP — Hitachi Dynamic Provisioning

HDR — Hitachi Dynamic Replicator

HDS — Hitachi Data Systems

HDT — Hitachi Dynamic Tiering

HDU — Hard Disk Unit. A number of hard drives (HDDs) grouped together within a subsystem.

HDvM — Hitachi Device Manager; also referred to as Device Manager (DM)

Head — See read/write head

HENP — Hitachi Essential NAS Platform

Heterogeneous — The characteristic of containing dissimilar elements. A common use of this word in information technology is to describe a product as able to contain or be part of a “heterogeneous network,” consisting of different manufacturers' products that can interoperate.

Heterogeneous networks are made possible by standards-conforming hardware and software

interfaces used in common by different products, thus allowing them to communicate with each other. The Internet itself is an example of a heterogeneous network.

HFSM — Hitachi File Service Manager

HGLM — Hitachi Global Link Manager

Hi Track System — Automatic fault reporting system.

HIHSM — Hitachi Internal Hierarchy Storage Management

HiRDB — Hitachi Relational Database

HIS — High Speed Interconnect

HiStar — Multiple point-to-point data paths to cache

Hitachi Data Protection Suite — Unified platform comprising solutions for data backup and recovery, migration, archiving, and replication.

Hitachi TrueCopy Extended Distance Remote Replication Software

Hitachi Storage Capacity Reporter Software — provides end-to-end storage capacity reporting from applications and hosts to heterogeneous storage arrays without the use of host based agents.

HLQ — High-level Qualifier

HLU — Host Logical Unit. A LU that the Operating System and the HDLM recognizes. Each HLU includes the devices that comprise the storage LU.

H-LUN — Host Logical Unit Number. See LUN.

HMC — Hardware Management Console

HMDE — Hitachi Multiplatform Data Exchange

HMRC F — Hitachi Multiple Raid Coupling Feature

HMRS — Hitachi Multiplatform Resource Sharing

HNAS — Hitachi NAS Platform

HODM — Hitachi Online Data Migration

Homogeneous — Of the same or similar kind

HOMRCF — Hitachi Open Multiple Raid Coupling Feature. Shadow Image is the marketing name for HOMRCF.

HORC — Hitachi Open Remote Copy. See TrueCopy.

HORCM — Hitachi Open Raid Configuration Manager

Host — Also called a server. Basically a central computer that processes end-user applications or requests.

Host LU — See HLU

Host Storage Domains — Allows host pooling at the LUN level and the priority access feature

lets administrator set service levels for applications.

HP — Hewlett-Packard Company or High Performance

HPAV — Hyper Parallel Access Volume (PAV)

HPC — High Performance Computing or Hitachi Remote Copy (see TrueCopy)

HRpM — Hitachi Replication Manager is the software tool that configures, monitors and manages Hitachi replication products in both open and mainframe environments for enterprise and modular storage systems from *a single pane of glass*.

HSA — Hardware System Area

HSG — Host Security Group

HSM — Hierarchical Storage Management (see Data Migrator)

HSN — Hierarchical Star Network

HSSDC — High Speed Serial Data Connector

HTnM — Hitachi Tuning Manager

HTsM — Hitachi Tiered Storage Manager

HTTP — Hyper Text Transfer Protocol

HTTPS — Hyper Text Transfer Protocol Secure

Hub — A common connection point for devices in a network. Hubs are commonly used to connect segments of a LAN. A hub contains multiple ports. When a packet arrives at one port, it is copied to the other ports so that all segments of the LAN can see all packets. A switching hub actually reads the destination address of each packet and then forwards the packet to the correct port. Device to which nodes on a multi-point bus or loop are physically connected.

HUR — Hitachi Universal Replicator is an asynchronous, continuous, nondisruptive, host-independent remote data replication solution for disaster recovery or data migration over long distances.

HXRC — Hitachi Extended Remote Copy

HyperPAV — Hyper Parallel Access Volume (PAV)

[-back to top-](#)



IDE — Integrated Drive Electronics Advanced Technology. A standard designed to connect hard and removable disk drives.

I/F — Interface

Index Cache — Provides quick access to indexed data on the media during a browse/restore operation.

I/O — Input/Output. Term used to describe any program, operation, or device that transfers data to or from a computer and to or from a peripheral device.

IBR — Incremental Block-level Replication

IBR — Intelligent Block Replication

ICB — Integrated Cluster Bus

ICF — Integrated Coupling Facility

ID — Identifier

IDR — Incremental Data Replication

iFCP — Internet Fibre Channel Protocol. Allows an organization to extend Fibre Channel storage networks over the Internet by using TCP/IP. TCP is responsible for managing congestion control as well as error detection and recovery services.

iFCP allows an organization to create an IP SAN fabric that minimizes the Fibre Channel fabric component and maximizes use of the company's TCP/IP infrastructure.

IFL — Integrated Facility for LINUX

IID — Initiator ID. Is used to identify LU, whether it is NAS System LU or User LU. If it is 0, that means it is NAS System LU; and if it is 1, then the LU is User LU.

IIS — Internet Information Server

ILM — Information Life Cycle Management

ILO — (Hewlett-Packard) Integrated Lights-Out

IML — Initial Microprogram Load

IMS — Information Management System

In-band virtualization — Refers to the location of the storage network path, between the application host servers in the storage systems. Provides both control and data along the same connection path. Also called symmetric virtualization.

Interface — The physical and logical arrangement supporting the attachment of any device to a connector or to another device.

Internal bus — Another name for an internal data bus. Also, an expansion bus is often referred to as an internal bus.

Internal data bus — A bus that operates only within the internal circuitry of the CPU, communicating among the internal caches of memory that are part of the CPU chip's design. This bus is typically rather quick and is independent of the rest of the computer's operations.

IOCDS — I/O Control Data Set

IODF — I/O Definition file

IOS — I/O Supervisor

IOSQ — Input/Output Subsystem Queue
 IP — Internet Protocol
 IPL — Initial Program Load
 IPSEC — IP security
 ISC — Initial shipping condition or Inter-System Communication
 iSCSI — Internet SCSI. Pronounced eye skuzzy.
 An IP-based standard for linking data storage devices over a network and transferring data by carrying SCSI commands over IP networks. iSCSI supports a Gigabit Ethernet interface at the physical layer, which allows systems supporting iSCSI interfaces to connect directly to standard Gigabit Ethernet switches or IP routers. When an operating system receives a request, it generates the SCSI command and then sends an IP packet over an Ethernet connection. At the receiving end, the SCSI commands are separated from the request, and the SCSI commands and data are sent to the SCSI controller and then to the SCSI storage device. iSCSI will also return a response to the request using the same protocol. iSCSI is important to SAN technology because it enables a SAN to be deployed in a LAN, WAN, or MAN.
 iSER — iSCSI Extensions for RDMA
 ISL — Inter-Switch Link
 iSNS — Internet Storage Name Service
 ISOE — iSCSI Offload Engine
 ISP — Internet service provider
 ISPF — Interactive System Productivity Facility
 ISPF/PDF — Interactive System Productivity Facility/Program Development Facility
 ISV — Independent Software Vendor

[-back to top-](#)

—J—

Java — A widely accepted, open systems programming language. Hitachi's enterprise software products are all accessed using Java applications. This enables storage administrators to access the Hitachi enterprise software products from any PC or workstation that runs a supported thin-client internet browser application and that has TCP/IP network access to the computer on which the software product runs.
 Java VM — Java Virtual Machine
 JBOD — Just a Bunch of Disks
 JCL — Job Control Language
 JMP —Jumper. Option setting method

JMS — Java Message Service
 JNL — Journal
 JNLG — Journal Group
 JRE —Java Runtime Environment
 JVM — Java Virtual Machine
 J-VOL — Journal Volume

[-back to top-](#)

—K—

KSDS — Key Sequence Data Set
 kVA— Kilovolt Ampere
 kW — Kilowatt

[-back to top-](#)

—L—

LACP — Link Aggregation Control Protocol
 LAG — Link Aggregation Groups
 LAN— Local Area Network
 LBA — Logical block address. A 28-bit value that maps to a specific cylinder-head-sector address on the disk.
 LC — Lucent connector. Fibre Channel connector that is smaller than a simplex connector (SC)
 LCDG—Link Processor Control Diagnostics
 LCM— Link Control Module
 LCP — Link Control Processor. Controls the optical links. LCP is located in the LCM.
 LCSS — Logical Channel Subsystems
 LCU — Logical Control Unit
 LD — Logical Device
 LDAP — Lightweight Directory Access Protocol
 LDEV — Logical Device or Logical DEVICE (number). A set of physical disk partitions (all or portions of one or more disks) that are combined so that the subsystem sees and treats them as a single area of data storage. Also called a volume. An LDEV has a specific and unique address within a subsystem. LDEVs become LUNs to an open-systems host.
 LDKC — Logical Disk Controller or Logical Disk Controller Manual
 LDM — Logical Disk Manager
 LDS — Linear Data Set
 LED — Light Emitting Diode
 LFF — Large Form Factor
 LIC — Licensed Internal Code
 LLQ — Lowest Level Qualifier

LM — Local Memory

LMODs — Load Modules

LNKLST — Link List

Load balancing — The process of distributing processing and communications activity evenly across a computer network so that no single device is overwhelmed. Load balancing is especially important for networks where it is difficult to predict the number of requests that will be issued to a server. If one server starts to be swamped, requests are forwarded to another server with more capacity. Load balancing can also refer to the communications channels themselves.

LOC — Locations section of the Maintenance

Logical DKC (LDKC) — Logical Disk Controller Manual. An internal architecture extension to the Control Unit addressing scheme that allows more LDEVs to be identified within one Hitachi enterprise storage system. The LDKC is supported only on Universal Storage Platform V/VM class storage systems.

As of March 2008, only one LDKC is supported, LDKC 00. Refer to product documentation as Hitachi has announced their intent to expand this capacity in the future.

LPAR — Logical Partition (mode)

LR — Local Router

LRECL — Logical Record Length

LRP — Local Router Processor

LRU — Least Recently Used

LSS — Logical Storage Subsystem (equivalent to LCU)

LU — Logical Unit. Mapping number of an LDEV

LUN — Logical Unit Number. One or more LDEVs. Used only for open systems.

LUN Manager — HDS software used to map Logical Unit Numbers (LUNs) to subsystem ports.

LUSE — Logical Unit Size Expansion. Feature used to create virtual LUs that are up to 36 times larger than the standard OPEN-x LUs.

LVDS — Low Voltage Differential Signal

LVI — Logical Volume Image. Identifies a similar concept (as LUN) in the mainframe environment.

LVM — Logical Volume Manager

— M —

MAC — Media Access Control. A MAC address is a unique identifier attached to most forms of networking equipment.

MAID — Massive array of disks

MAN — Metropolitan Area Network

Mapping — Conversion between two data addressing spaces. For example, mapping refers to the conversion between physical disk block addresses and the block addresses of the virtual disks presented to operating environments by control software.

Mb — Megabits

MB — Megabytes

MBA — Memory Bus Adaptor

MBUS — Multi-CPU Bus

MC — Multi Cabinet

MCU — Main Disk Control Unit. The local CU of a remote copy pair. Main or Master Control Unit.

MCU — Master Control Unit.

MediaAgent — The workhorse for all data movement. MediaAgent facilitates the transfer of data between the data source, the client computer, and the destination storage media.

Metadata — In database management systems, data files are the files that store the database information; whereas other files, such as index files and data dictionaries, store administrative information, known as metadata.


MFC — Main Failure Code

MG — Module Group. Two (DIMM) cache memory modules that work together.

MGC — (3-Site) Metro/Global Mirror

MIB — Management Information Base. A database of objects that can be monitored by a network management system. Both SNMP and RMON use standardized MIB formats that allow any SNMP and RMON tools to monitor any device defined by a MIB.

Microcode — The lowest-level instructions that directly control a microprocessor. A single machine-language instruction typically translates into several microcode instructions.



[-back to top-](#)

Microprogram — See Microcode
MIF — Multiple Image Facility
Mirror Cache OFF — Increases cache efficiency over cache data redundancy.
M-JNL — Primary journal volumes.
MM — Maintenance Manual
MMC — Microsoft Management Console
Mode — The state or setting of a program or device. The term mode implies a choice, which is that you can change the setting and put the system in a different mode.
MP — Microprocessor
MPA — Microprocessor adapter
MPIO — Multipath I/O
MPU — Microprocessor Unit
MS/SG — Microsoft Service Guard
MSCS — Microsoft Cluster Server
MSS — Multiple Subchannel Set
MTBF — Mean Time Between Failure
MTS — Multi-Tiered Storage
M-VOL — Main Volume
MVS — Multiple Virtual Storage

[-back to top-](#)

—N—

NAS — Network Attached Storage. A disk array connected to a controller that gives access to a LAN Transport. It handles data at the file level.
NAT — Network Address Translation
NDMP — Network Data Management Protocol is a protocol meant to transport data between NAS devices.
NetBIOS — Network Basic Input/Output System
Network — A computer system that allows sharing of resources, such as files and peripheral hardware devices.
NFS protocol — Network File System is a protocol that allows a computer to access files over a network as easily as if they were on its local disks.
NIM — Network Interface Module
NIS — Network Information Service (originally called the Yellow Pages or YP)
NLS — Native Language Support
Node — An addressable entity connected to an I/O bus or network, used primarily to refer to computers, storage devices, and storage subsystems. The component of a node that connects to the bus or network is a port.

Node name — A Name_Identifier associated with a node
NRO — Network Recovery Objective
NTP — Network Time Protocol
NVS — Non Volatile Storage

[-back to top-](#)

—O—

OEM — Original Equipment Manufacturer
OFC — Open Fibre Control
OID — Object identifier
OLA — Operating Level Agreements
OLTP — On-Line Transaction Processing
On/Off CoD — On/Off Capacity on Demand
ONODE — Object node
OPEX — Operational Expenditure. This is an operating expense, operating expenditure, operational expense, or operational expenditure, which is an on-going cost for running a product, business, or system. Its counterpart is a capital expenditure (CAPEX).
ORM — Online Read Margin
OS — Operating System
Out-of-band virtualization — Refers to systems where the controller is located outside of the SAN data path. Separates control and data on different connection paths. Also called asymmetric virtualization.

[-back to top-](#)

—P—

PACS — picture archiving and communication systems
Parity — A technique of checking whether data has been lost or written over when it is moved from one place in storage to another or when it is transmitted between computers.
Parity Group — Also called an array group. This is a group of hard disk drives (HDDs) that form the basic unit of storage in a subsystem. All HDDs in a parity group must have the same physical capacity.
Partitioned cache memory — Separate workloads in a “storage consolidated” system by dividing cache into individually managed multiple partitions. Then customize the partition to match the I/O characteristics of assigned LUs
PAT — Port Address Translation
PATA — Parallel ATA
Path — Also referred to as a transmission channel, the path between two nodes of a

network that a data communication follows. The term can refer to the physical cabling that connects the nodes on a network, the signal that is communicated over the pathway or a sub-channel in a carrier frequency.

Path failover — See Failover.

PAV — Parallel Access Volumes

PAWS — Protect Against Wrapped Sequences

PBC — Port By-pass Circuit

PCB — Printed Circuit Board

PCHIDS — Physical Channel Path Identifiers

PCI — Power Control Interface

PCI CON — Power Control Interface Connector Board

PD — Product Detail

PDEV — Physical Device

PDM — Policy based Data Migration or Primary Data Migrator

PDS — Partitioned Data Set

PDSE — Partitioned Data Set Extended

Performance — Speed of access or the delivery of information

PFA — Predictive Failure Analysis

PGR — Persistent Group Reserve

PI — Product Interval

PIR — Performance Information Report

PiT — Point-in-Time

PK — Package (see PCB)

PL — Platter. The circular disk on which the magnetic data is stored. Also called motherboard or backplane.

PM — Package Memory

Port — In TCP/IP and UDP networks, an endpoint to a logical connection. The port number identifies what type of port it is. For example, port 80 is used for HTTP traffic.

P-P — Point-to-point; also P2P.

PPRC — Peer-to-Peer Remote Copy

PR/SM — Processor Resource/System Manager

Priority Mode— Also PRIO mode, is one of the modes of FlashAccess™ in which the FlashAccess™ extents hold read and write data for specific extents on volumes. See Bind Mode.

Protocol — A convention or standard that enables the communication between two computing endpoints. In its simplest form, a protocol can be defined as the rules governing the syntax, semantics, and synchronization of communication. Protocols may be

implemented by hardware, software, or a combination of the two. At the lowest level, a protocol defines the behavior of a hardware connection.

Provisioning — The process of allocating storage resources and assigning storage capacity for an application, usually in the form of server disk drive space, in order to optimize the performance of a storage area network (SAN). Traditionally, this has been done by the SAN administrator, and it can be a tedious process. In recent years, automated storage provisioning (also called auto-provisioning) programs have become available. These programs can reduce the time required for the storage provisioning process, and can free the administrator from the often distasteful task of performing this chore manually

PS — Power Supply

PSA — Partition Storage Administrator

PSSC — Perl Silicon Server Control

PSU — Power Supply Unit

PTAM — Pickup Truck Access Method

PTF — Program Temporary Fixes

PTR — Pointer

PU — Processing Unit

P-VOL — Primary Volume

[-back to top-](#)

— Q —

QD — Quorum Device

QoS — Quality of Service. In the field of computer networking, the traffic engineering term quality of service (QoS) refers to resource reservation control mechanisms rather than the achieved service quality. Quality of service is the ability to provide different priority to different applications, users, or data flows, or to guarantee a certain level of performance to a data flow.

QSAM — Queued Sequential Access Method

[-back to top-](#)

— R —

RACF — Resource Access Control Facility

RAID — Redundant Array of Independent Disks, or Redundant Array of Inexpensive Disks. A group of disks that look like a single volume to the server. RAID improves performance by pulling a single stripe of data from multiple disks, and improves fault-tolerance either through mirroring or parity checking and it is a component of a customer's SLA.

RAID-0 — Striped array with no parity

RAID-1 — Mirrored array and duplexing

RAID-3 — Striped array with typically non-rotating parity, optimized for long, single-threaded transfers.

RAID-4 — Striped array with typically non-rotating parity, optimized for short, multi-threaded transfers.

RAID-5 — Striped array with typically rotating parity, optimized for short, multithreaded transfers.

RAID-6 — Similar to RAID-5, but with dual rotating parity physical disks, tolerating two physical disk failures.

RAM — Random Access Memory

RAM DISK — A LUN held entirely in the cache area.

RAS — Reliability, Availability, and Serviceability or Row Address Strobe

RBAC — Role Base Access Control

RC — Reference Code or Remote Control

RCHA — RAID Channel Adapter

RCP — Remote Control Processor

RCU — Remote Control Unit or Remote Disk Control Unit

RDWR — Read/Write

RDM — Raw Disk Mapped

RDMA — Remote Direct Memory Access

RDW — Record Descriptor Word

Read/Write Head — Read and write data to the platters, typically there is one head per platter side, and each head is attached to a single actuator shaft.

RECFM — Record Format Redundant. Describes the computer or network system components, such as fans, hard disk drives, servers, operating systems, switches, and telecommunication links that are installed to back up primary resources in case they fail.

A well-known example of a redundant system is the redundant array of independent disks (RAID). Redundancy contributes to the fault tolerance of a system.

Redundancy — Backing up a component to help ensure high availability.

Reliability — Level of assurance that data will not be lost or degraded over time.

Resource Manager — Hitachi Resource Manager™ utility package is a software suite

that rolls into one package the following four pieces of software:

- Hitachi Cache Residency Manager feature (formerly FlashAccess) (uses cache to speed data reads and writes)
- Hitachi Graph-Track™ performance monitor feature
- LUN Manager (reconfiguration of LUNS, or logical unit numbers)
- Virtual Logical Volume Image (VLMI) Manager (optimizes capacity utilization)

Reliability — An attribute of any computer component (software, hardware, or a network) that consistently performs according to its specifications.

REXX — Restructured eXtended eXecutor

RID — Relative Identifier that uniquely identifies a user or group within a Microsoft Windows domain

RISC — Reduced Instruction Set Computer

R-JNL — Secondary journal volumes.

RK — Rack. Acronym referring to the main “Rack” unit, which houses the core operational hardware components of the Thunder 9500V/9200 subsystem. See also: RKA, H1F, and H2F.

RKA (Rack Additional) — Acronym referring to “Rack Additional”, namely additional rack unit(s) that house additional hard drives exceeding the capacity of the core RK unit of the Thunder 9500V/9200 subsystem. See also: RK, RKA, H1F, and H2F.

RKAJAT — Rack Additional SATA disk tray

RLGFAN — Rear Logic Box Fan Assembly

RLOGIC BOX — Rear Logic Box

RMF — Resource Measurement Facility

RMI — Remote Method Invocation. A way that a programmer, using the Java programming language and development environment, can write object-oriented programming in which objects on different computers can interact in a distributed network. RMI is the Java version of what is generally known as a RPC (remote procedure call), but with the ability to pass one or more objects along with the request.

ROA — Return on Asset

RoHS — Restriction of Hazardous Substances (in Electrical and Electronic Equipment)

ROI — Return on Investment

ROM — Read-only memory

Round robin mode — A load balancing technique in which balances power is placed in the DNS server instead of a strictly dedicated machine as other load techniques do. Round robin works on a rotating basis in that one server IP address is handed out, then moves to the back of the list; the next server IP address is handed out, and then it moves to the end of the list; and so on, depending on the number of servers being used. This works in a looping fashion.

Round robin DNS is usually used for balancing the load of geographically distributed Web servers.

Router — A computer networking device that forwards data packets toward their destinations, through a process known as routing.

RPC — Remote procedure call

RPO — Recovery Point Objective. The point in time that recovered data should match.

RPSFAN — Rear Power Supply Fan Assembly

RRDS — Relative Record Data Set

RS CON — RS232C/RS422 Interface Connector

RSD — Raid Storage Division

R-SIM — Remote Service Information Message

RSM — Real Storage Manager

RTM — Recovery Termination Manager

RTO — Recovery Time Objective. The length of time that can be tolerated between a disaster and recovery of data.

R-VOL — Remote Volume

R/W — Read/Write

[-back to top-](#)

—S—

SA — Storage Administrator

SA z/OS — System Automation for z/OS

SAA — Share Access Authentication. The process of restricting a user's rights to a file system object by combining the security descriptors from both the file system object itself and the share to which the user is connected

SACK — Sequential Acknowledge

SACL — System ACL. The part of a security descriptor that stores system auditing information.

SAN — Storage Area Network. A network linking computing devices to disk or tape arrays and other devices over Fibre Channel. It handles data at the block level.

SANTinel — HDS software that provides LUN security. SANTinel protects data from unauthorized access in SAN environments. It restricts server access by implementing boundaries around predefined zones and is used to map hosts in a host group to the appropriate LUNs.

SAP — System Assist Processor (for I/O Processing)

SARD — System Assurance Registration Document

SAS — SAN Attached Storage. Storage elements that connect directly to a storage area network (SAN) and provide data access services to computer systems. or Serial Attached SCSI. Disk drive configurations for Hitachi Simple Modular Storage 100 systems

SATA — Serial ATA. Serial Advanced Technology Attachment is a new standard for connecting hard drives into computer systems. SATA is based on serial signaling technology, unlike current IDE (Integrated Drive Electronics) hard drives that use parallel signaling.

SBOD — Switched Bunch of Disks

SBSC — Smart Business Storage Cloud

SBX — Small Box (Small Form Factor)

SC — Simplex connector. Fibre Channel connector that is larger than a Lucent connector (LC). or Single Cabinet

SCM — Supply Chain Management

SCP — Secure Copy

SCSI — Small Computer Systems Interface. A parallel bus architecture and a protocol for transmitting large data blocks up to a distance of 15 to 25 meters.

SD — Software Division

SDM — System Data Mover

SDSF — Spool Display and Search Facility

Sector — A sub-division of a track of a magnetic disk that stores a fixed amount of data.

SEL — System Event Log

Selectable segment size — Can be set per partition.

Selectable Stripe Size — Increases performance by customizing the disk access size.

SENC — Is the SATA (Serial ATA) version of the ENC. ENCs and SENCs are complete microprocessor systems on their own and they occasionally require a firmware upgrade.

Serial Transmission — The transmission of data bits in sequential order over a single line.

Server — A central computer that processes end-user applications or requests, also called a host.

Service-level agreement — SLA. A contract between a network service provider and a customer that specifies, usually in measurable terms, what services the network service provider will furnish. Many Internet service providers (ISP) provide their customers with a SLA. More recently, IT departments in major enterprises have adopted the idea of writing a service level agreement so that services for their customers (users in other departments within the enterprise) can be measured, justified, and perhaps compared with those of outsourcing network providers.

Some metrics that SLAs may specify include:

- The percentage of the time services will be available
- The number of users that can be served simultaneously
- Specific performance benchmarks to which actual performance will be periodically compared
- The schedule for notification in advance of network changes that may affect users
- Help desk response time for various classes of problems
- Dial-in access availability
 - Usage statistics that will be provided

Service-Level Objective — SLO. Individual performance metrics built into an SLA. Each SLO corresponds to a single performance characteristic relevant to the delivery of an overall service. Some examples of SLOs include: system availability, help desk incident resolution time, and application response time.

SES — SCSI Enclosure Services

SFF — Small Form Factor

SFI — Storage Facility Image

SFM — Sysplex Failure Management

SFP — Small Form-Factor Pluggable module Host connector. A specification for a new generation of optical modular transceivers. The devices are designed for use with small form factor (SFF) connectors, offer high speed and physical compactness, and are hot-swappable.

ShadowImage® — HDS software used to duplicate large amounts of data within a

subsystem without affecting the service and performance levels or timing out. ShadowImage replicates data with high speed and reduces backup time.

SHSN — Shared memory Hierarchical Star Network

SI — Hitachi ShadowImage® Replication software

SID — Security Identifier. A user or group identifier within the Microsoft Windows security model.

SIGP — Signal Processor

SIM — Service Information Message. A message reporting an error that contains fix guidance information. or Storage Interface Module or Subscriber Identity Module

SIM RC — Service (or system) Information Message Reference Code

SIMM — Single In-line Memory Module

Slz — Hitachi ShadowImage® Replication Software

SLA — Service Level Agreement

SLO — Service-Level Objective

SLPR (Storage administrator Logical PaRtition) — Storage can be divided among various users to reduce conflicts with usage.

SLRP — storage logical partitions

SM — Shared Memory or Shared Memory Module. Stores the shared information about the subsystem and the cache control information (director names). This type of information is used for the exclusive control of the subsystem. Like CACHE, shared memory is controlled as two areas of memory and fully non-volatile (sustained for approximately seven days).

SM PATH— Shared Memory Access Path. The Access Path from the processors of CHA, DKA PCB to Shared Memory.

SMB/CIFS — Server Message Block Protocol/Common Internet File System

SMC — Shared Memory Control

SMF — System Management Facility

SMI-S — Storage Management Initiative Specification

SMP — Symmetric Multiprocessing. An IBM licensed program used to install software and software changes on z/OS systems.

SMP/E — System Modification Program/Extended

SMS — Hitachi Simple Modular Storage or System Managed Storage

SMTP — Simple Mail Transfer Protocol

SMU — System Management Unit

Snapshot Image — A logical duplicated volume (V-VOL) of the primary volume. It is an internal volume intended for restoration.

SNIA — Storage Networking Industry Association. An association of producers and consumers of storage networking products, whose goal is to further storage networking technology and applications.

SNM — Storage Navigator Modular Program

SNM2 — Storage Navigator Modular 2 Program

SNMP — Simple Network Management Protocol. A TCP/IP protocol that was designed for management of networks over TCP/IP, using agents and stations.

SOA — Service Oriented Architecture

SOAP — Simple object access protocol. A way for a program running in one kind of operating system (such as Windows 2000) to communicate with a program in the same or another kind of an operating system (such as LINUX) by using the World Wide Web's Hypertext Transfer Protocol (HTTP) and its Extensible Markup Language (XML) as the mechanisms for information exchange.

Socket — In UNIX and some other operating systems, socket is a software object that connects an application to a network protocol. In UNIX, for example, a program can send and receive TCP/IP messages by opening a socket and reading and writing data to and from the socket. This simplifies program development because the programmer need only worry about manipulating the socket and can rely on the operating system to actually transport messages across the network correctly. Note that a socket in this sense is completely soft; it is a software object, not a physical component.

SOM — System Option Mode

SOSS — Service Oriented Storage Solutions

SPAN — Span is a section between two intermediate supports. See Storage pool

Spare — An object reserved for the purpose of substitution for a like object in case of that object's failure.

SPC — SCSI Protocol Controller

SpecSFS — Standard Performance Evaluation Corporation Shared File system

SPECsfs97 — Standard Performance Evaluation Corporation (SPEC) System File Server (sfs) developed in 1997 (97)

SRDF/A — (EMC) Symmetrix Remote Data Facility Asynchronous

SRDF/S — (EMC) Symmetrix Remote Data Facility Synchronous

SSB — Sense Byte

SSC — SiliconServer Control

SSCH — Start Subchannel

SSD — Solid-state Drive or Solid-State Disk

SSH — Secure Shell

SSID — Storage Subsystem ID or Subsystem Identifier

SSL — Secure Sockets Layer

SSPC — System Storage Productivity Center

SSVP — Sub Service Processor interfaces the SVP to the DKC

SSW — SAS Switch

Sticky Bit — Extended Unix mode bit that prevents objects from being deleted from a directory by anyone other than the object's owner, the directory's owner or the root user

Storage pooling — The ability to consolidate and manage storage resources across storage system enclosures where the consolidation of many appears as a single view.

STP — Server Time Protocol

STR — Storage and Retrieval Systems

Striping — A RAID technique for writing a file to multiple disks on a block-by-block basis, with or without parity.

Subsystem — Hardware or software that performs a specific function within a larger system.

SVC — Supervisor Call Interruption

SVC Interrupts — Supervisor calls

S-VOL — (ShadowImage) Source Volume for In-System Replication or (Universal Replicator) Secondary Volume or Secondary Volume

SVP (Service Processor) — A laptop computer mounted on the control frame (DKC) and used for monitoring, maintenance and administration of the subsystem

Switch — A fabric device providing full bandwidth per port and high-speed routing of data via link-level addressing.

SXP — SAS Expander

Symmetric virtualization — See In-band virtualization.

Synchronous — Operations that have a fixed time relationship to each other. Most commonly used to denote I/O operations that occur in time sequence, i.e., a successor operation does not occur until its predecessor is complete.

[-back to top-](#)

—T—

T.S.C. — Technical Support Center. A chip developed by HP, and used in various devices. This chip has FC-0 through FC-2 on one chip.

Target — The system component that receives a SCSI I/O command, an open device that operates at the request of the initiator

TC — TrueCopy software

TCA — TrueCopy Asynchronous

TCO — Total Cost of Ownership

TCP/IP — Transmission Control Protocol over Internet Protocol

TCP/UDP — User Datagram Protocol is one of the core protocols of the Internet protocol suite. Using UDP, programs on networked computers can send short messages known as datagrams to one another.

TCS — TrueCopy Synchronous

TCz — Hitachi TrueCopy® Remote Replication (Synchronous) software is a continuous, nondisruptive, host independent remote data replication solution for disaster recovery or data migration for distances within the same metropolitan area.

TDCONV — Trace Dump CONVerter. A software program that is used to convert traces taken on the system into readable text. This information is loaded into a special spreadsheet that allows for further investigation of the data. More in-depth failure analysis.

TDMF — Transparent Data Migration Facility

TEP — Tivoli Enterprise Portal

TFS — Temporary File System

TGTLIBs — Target Libraries

THF — Front Thermostat

Thin Provisioning — Thin Provisioning allows space to be easily allocated to servers, on a just-enough and just-in-time basis.

THR — Rear Thermostat

Throughput — The amount of data transferred from one place to another or processed in a specified amount of time. Data transfer rates for disk drives and networks are measured in terms of throughput. Typically, throughputs are measured in kbps, Mbps and Gbps.

TID — Target ID

Tiered storage — A storage strategy that matches data classification to storage metrics. Tiered

storage is the assignment of different categories of data to different types of storage media in order to reduce total storage cost. Categories may be based on levels of protection needed, performance requirements, frequency of use, and other considerations. Since assigning data to particular media may be an ongoing and complex activity, some vendors provide software for automatically managing the process based on a company-defined policy.

Tiered Storage Promotion — Moving data between tiers of storage as their availability requirements change

TISC — The Hitachi Data Systems internal Technical Information Service Centre from which microcode, user guides, ECNs, and so on. can be downloaded.

TLS — Tape Library System

TLS — Transport Layer Security

TMP — Temporary or Test Management Program

TOC — Table Of Contents

TOD (or ToD) — Time Of Day

TOE — TCP Offload Engine

Topology — The shape of a network or how it is laid out. Topologies are either physical or logical.

TPC-R — Tivoli Productivity Center for Replication

TPF — Transaction Processing Facility

Track — Circular segment of a hard disk or other storage media.

Transfer Rate — See Data Transfer Rate.

Trap — A program interrupt, usually an interrupt caused by some exceptional situation in the user program. In most cases, the Operating System performs some action, and then returns control to the program.

TRC — Technical Resource Center

TrueCopy — HDS software that replicates data between subsystems. These systems can be located within a data center or at geographically separated data centers. The 9900V adds the capability of using TrueCopy to make copies in two different locations simultaneously.

TSC — Technical Support Center

TSO — Time Sharing Option

TSO/E — Time Sharing Option/Extended

T-VOL — (ShadowImage) Target Volume for In-System Replication

[-back to top-](#)

—U—

UA — Unified Agent
UBX — Large Box (Large Form Factor)
UCB — Unit Control Block
UFA — UNIX File Attributes
UID — User Identifier within the UNIX security model
Universal Volume Manager software — Provides the virtualization of a multi-tiered storage area network comprised of heterogeneous storage systems.
UPS — Uninterruptible Power Supply — A power supply that includes a battery to maintain power in the event of a power outage.
UR — Universal Replicator
UR P-VOL/S-VOL — Universal Replicator Primary or Secondary Volume
URz — Hitachi Universal Replicator software
USP — Universal Storage Platform™
USP V — (Hitachi) Universal Storage Platform™ V
USP VM — (Hitachi) Universal Storage Platform™ VM
UUID — Universally Unique Identifier
UVM — Universal Volume Manager

[-back to top-](#)

—V—

vContinuum — Using the vContinuum wizard, users can push agents to primary and secondary servers, set up protection, and perform failovers and failbacks.
VCS — Veritas Cluster System
VDEV — Virtual Device
VHDL — VHSIC (Very-High-Speed Integrated Circuit) Hardware Description Language
VHSIC — Very-High-Speed Integrated Circuit
VI — Virtual Interface. A research prototype that is undergoing active development, and the details of the implementation may change considerably. It is an application interface that gives user-level processes direct but protected access to network interface cards.

This allows applications to bypass IP processing overheads (for example, copying data, computing checksums) and system call overheads while still preventing one process from accidentally or maliciously tampering with or reading data being used by another.

Virtualization — The amalgamation of multiple network storage devices into what appears to be a **single** storage unit. Storage virtualization is often used in a SAN, and makes tasks such as archiving, back up, and recovery easier and faster. Storage virtualization is usually implemented via software applications.

VLL — Virtual Logical Volume Image/Logical Unit Number
VLUN — Virtual LUN. Customized volume. Size chosen by user.
VLVI — Virtual Logic Volume Image. Marketing name for CVS (custom volume size).
VM — Virtual Machines
VOJP — (Cache) Volatile Jumper
VOLID — Volume ID
VOLSER — Volume Serial Numbers
Volume — A fixed amount of storage on a disk or tape. The term volume is often used as a synonym for the storage medium itself, but it is possible for a single disk to contain more than one volume or for a volume to span more than one disk.
VSAM — Virtual Storage Access Method
VSD — Virtual Storage Director
VSP — Virtual Storage Platform
VTOC — Volume Table of Contents
VTOCIX — Volume Table of Contents Index
VVDS — Virtual Volume Data Set
V-VOL — Virtual Volume

[-back to top-](#)

—W—

WAN — Wide Area Network
WDIR — Directory Name Object
WDIR — Working Directory
WDS — Working Data Set
WFILE — File Object or Working File
WFS — Working File Set
WINS — Windows Internet Naming Service
WLM — Work Load Manager
WMS — Hitachi Workgroup Modular Storage system
WORM — Write Once, Read Many
WSDL — Web Services Description Language
WTREE — Directory Tree Object or Working Tree

WWN — World Wide Name. A unique identifier for an open-system host. It consists of a 64-bit physical address (the IEEE 48-bit format with a 12-bit extension and a 4-bit prefix). The WWN is essential for defining the Hitachi Volume Security software (formerly SANTinel) parameters, because it determines whether the open-system host is to be allowed or denied access to a specified LU or a group of LUs.

WWNN — World Wide Node Name. A globally unique 64-bit identifier assigned to each Fibre Channel node process.

WWPN — World Wide Port Name. A globally unique 64-bit identifier assigned to each Fibre Channel port. Fibre Channel ports' WWPN are permitted to use any of several naming authorities. Fibre Channel specifies a Network Address Authority (NAA) to distinguish between the various name registration authorities that may be used to identify the WWPN.

[-back to top-](#)

—X—

XAUI — "X"=10, AUI = Attachment Unit Interface

XCF — Cross System Communications Facility

XFI — Standard interface for connecting 10Gb Ethernet MAC device to XFP interface

XFP — "X"=10Gb Small Form Factor Pluggable

XRC — Extended Remote Copy

[-back to top-](#)

—Y—

[-back to top-](#)

—Z—

z/OS — z Operating System (IBM® S/390® or z/OS® Environments)

z/OS NFS — (System) z/OS Network File System

z/OSMF — (System) z/OS Management Facility

zAAP — (System) z Application Assist Processor (for Java and XML workloads)

zFS — (System) zSeries File System

zHPF — (System) z High Performance FICON

zIIP — (System) z Integrated Information Processor (specialty processor for database)

Zone — A collection of Fibre Channel Ports that are permitted to communicate with each other via the fabric

Zoning — A method of subdividing a storage area network into disjoint zones, or subsets of nodes on the network. Storage area network nodes outside a zone are invisible to nodes within the zone. Moreover, with switched SANs, traffic within each zone may be physically isolated from traffic outside the zone.

[-back to top-](#)



Evaluating this Course

Please use the online evaluation system to help improve our courses.

For evaluations handled inside the Learning Center, sign in to:

<https://learningcenter.hds.com/Saba/Web/Main>

Evaluations can be reached by clicking on the My Learning tab, followed by Evaluations & Surveys on the left bar. Click on the “Launch” link to evaluate the class.

The screenshot shows the Learning Center interface. At the top, there is a navigation bar with tabs: Home, My Success Plan, My Learning (highlighted with a red box), My Skills and Competencies, My Network, and My Profile. Below this, on the left, is a sidebar menu with options: In-Progress Learning, Catalog, All Learning Activity, Learning Requests, Approvals Inbox, Evaluations & Surveys (highlighted with a red box), and Surveys. A red arrow points from the 'Evaluations & Surveys' menu item to the 'Evaluations' section in the main content area. The 'Evaluations' section has two tabs: Current and Completed. Below the tabs is a table of evaluations. The table has columns: Evaluation Name, Title, Date Submitted, and Action. There is one row with the following data: WBTEvaluation, Preventing Sexual Harassment (WBT), 01-21-2010, and a Launch button (highlighted with a red box). Above the table are links for Print, Export, and Modify Table. A red arrow points from the 'Launch' button to the right.

Evaluation Name	Title	Date Submitted	Action
WBTEvaluation	Preventing Sexual Harassment (WBT)	01-21-2010	Launch

